# Centre Stage: Centricity-based Audio-Visual Temporal Action Detection

Hanyuan Wang hanyuan.wang@bristol.ac.uk Majid Mirmehdi majid@cs.bris.ac.uk Dima Damen dima.damen@bristol.ac.uk

Toby Perrett toby.perrett@bristol.ac.uk Department of Computer Science Faculty of Engineering University of Bristol Bristol, UK

#### Abstract

Previous one-stage action detection approaches have modelled temporal dependencies using only the visual modality. In this paper, we explore different strategies to incorporate the audio modality, using multi-scale cross-attention to fuse the two modalities. We also demonstrate the correlation between the distance from the timestep to the action centre and the accuracy of the predicted boundaries. Thus, we propose a novel network head to estimate the closeness of timesteps to the action centre, which we call the centricity score. This leads to increased confidence for proposals that exhibit more precise boundaries. Our method can be integrated with other one-stage anchor-free architectures and we demonstrate this on three recent baselines on the EPIC-Kitchens-100 action detection benchmark where we achieve state-of-the-art performance. Detailed ablation studies showcase the benefits of fusing audio and our proposed centricity scores. Code and models for our proposed method are publicly available at https: //github.com/hanielwang/Audio-Visual-TAD.git.

# **1** Introduction

Temporal action detection aims to predict the boundaries of action segments from a long untrimmed video and classify the actions, as a fundamental step towards video understanding  $[B, \square]$ ,  $\square]$ . A typical challenging scenario is in unscripted actions in egocentric videos  $[B, \square]$  which contain dense action segments of various lengths in an unedited video, ranging from seconds to minutes.

Most recently, a few have approached egocentric action detection by modelling their long-range visual dependencies with transformers [53, 53, 59, 53, 59, 53, 59]. However, only using visual information, means a missed opportunity to exploit potentially meaningful aural action cues. As shown in Figure 1(a), sound exhibits discriminating characteristics around the starting point of actions, such as 'open drawer', 'take spoon' and 'scoop yoghurt', which can be useful for boundary regression. Also for action classification, the sound of flowing



Figure 1: Our motivation -(a) Sounds help detect actions both in refining the action boundaries (regression) and in identifying the action within these boundaries (classification), (b) Timesteps closer to the action center generate better proposals with high tIoU.

water can boost confidence in identifying an action as 'turn-on tap' rather than 'turn-off tap', even though their visual content is similar. Unlike methods [17, 13, 11] that directly fuse audio and visual modalities at the same scale through concatenation, addition or gating modules, in this paper we learn these modalities with separate encoders and fuse their representations using a cross-modal attention mechanism at different temporal scales. This allows us to exploit sufficient audio-visual information to detect actions of various duration.

Recent one-stage anchor-free methods [59, 52, 59] operate on egocentric videos by simultaneously predicting boundaries and action categories for each timestep. In contrast to anchor-based methods, anchor-free methods do not require pre-defined anchors to locate actions but directly generate one proposal for each timestep. We have observed that timesteps near the centre of actions tend to produce proposals with more precise boundaries. These proposals have higher temporal Intersection-over-Union (tIoU) values with corresponding ground-truth segments. As shown in Figure 1(b), the closer the current timestep is to the action centre, the greater the tIoU. Inspired by this observation, we introduce a centricity head that predicts a score so as to indicate how close the current timestep is to the action centre. This score is then integral to calculating the confidence scores for ranking candidate proposals, where those with more precise boundaries will be ranked higher. Our approach can be incorporated into most one-stage anchor-free action detectors and achieve significant improvement.

In summary, our key contributions are as follows: (i) we introduce a framework to effectively fuse audio and visual modalities using a cross-modal attention mechanism at various temporal scales, (ii) we propose a novel centricity head to predict the degree of closeness of each frame's temporal distance to the action centre – this boosts a proposal's confidence score and allows for the preferential selection of proposals with more precise boundaries, and (iii) we achieve state-of-the-art results on the EPIC-Kitchens-100 action detection benchmark, demonstrating the effectiveness of audio modality and the benefits of centricity in improving detection performance.

## 2 Related Work

**Temporal action detection** – Current temporal action detection methods can be divided into: (i) two-stage methods that first generate proposals and then classify them, and (ii) one stage methods that predict boundaries and corresponding classes simultaneously. Some two-stage works generate proposals by estimating boundary probabilities  $[\Box 2, \Box 3, \Box 3, \Box 3]$  and

action-ness scores [50]. Many one-stage methods [10], 12, 12, 12, 12] rely on pre-defined anchors to model temporal relations, which often leads to inflexibility and poor boundaries when detecting actions with various lengths. To address this, recent anchor-free methods [12], 13, 19] predict the action category and offsets to the boundaries simultaneously for each timestep using parallel classification and regression heads. Then, candidate proposals constructed by these predictions are filtered to obtain the final results. Our work follows such an anchor-free pipeline.

Inspired by the DETR framework [5], some works input relational queries [53], learned actions [29] or graph queries [53] to a transformer decoder to detect actions. However, with a limited number of queries, these methods struggle to cover a large number of actions in long videos. Alternatively, other works [0, 8, 59, 59] use multi-scale transformer encoders [59, 59] to model temporal dependencies for stronger video representations. For example, ActionFormer [59] applies local self-attention to extract a discriminative feature pyramid, which is then used for classification and regression. Our work falls under this workflow.

Audio-visual learning – Sight and hearing are both vital sensory modes that assist humans in perceiving the world. This can transfer to computational approaches too to learn models from. Numerous works [0, 0, 0, 0, 0, 10, 10, 10, 10] have focused on jointly learning audio and visual representations for tasks such as action recognition  $[\square 4, \square, \square]$ , video parsing  $[\square, \square]$  and event localization [44, 44], 44]. Audio-visual event localization aims to classify each timestep into a limited number of categories [41], relying on clear audio-visual signals and without the need to predict temporal boundaries. In contrast, our action detection task aims to leverage the audio-visual representation to detect temporal boundaries for dense actions with various lengths and unclear audio cues, and then classify them into a wide range of categories. OWL attempts different strategies for fusing audio and visual modalities, but it fuses at a single temporal scale only and classifies pre-generated proposals from [26], rather than detect boundaries. In [22], the authors address this task by extracting intra-modal features, but their proposed framework is designed for simple, weakly-labelled data with sparse actions per video. Our work focuses on large-scale egocentric data comprising dense complex actions of various durations, and we propose a framework to incorporate audio-visual learning and centricity into one-stage anchor-free methods.

### 3 Method

We propose a novel framework for temporal action detection, rooted in audio-visual data, which can be incorporated into one-stage anchor-free pipelines [23, 59, 52, 59] (see Figure 2). Similar to such temporal action detection works, we define the problem as follows. Given an untrimmed video, we extract features for the video and audio modalities and then process them using transformer encoders to obtain the visual and audio representation sequences  $F^{\nu} = \{f_t^{\nu}\}_{t=1}^T$  and  $F^a = \{f_t^a\}_{t=1}^T$ , respectively. Based on these, our approach is to learn to predict a set of possible action instances  $\Phi = \{(s, e, \alpha)_m\}_{m=1}^M$ , where *s* and *e* represent the starting and ending boundaries of an action, and  $\alpha$  represents the predicted action class.

### 3.1 Audio-visual Fusion

In this section, we explore three different strategies to effectively utilise the audio modality and combine it with visual information to improve action detection performance.



Figure 2: An overview of our architecture – Given an untrimmed video, audio and visual features are extracted from video clips, then fed into encoders  $E_a$  and  $E_v$  for generating audio  $F^a$  and visual  $F^v$  feature pyramids. These are fused using cross-attention across N temporal scales to build the audio-visual representation  $F^{av}$ . This representation is passed to the centricity head  $h_C$  for centricity scores, and to prediction heads  $h_b^v$ ,  $h_c^v$ , and  $h_c^a$  for boundary, classification, and audio scores respectively. Finally, predicted scores and boundaries are used to construct candidate proposals, which are then filtered to obtain the final predictions.

**Proposal fusion** – In this strategy (see Figure 3 (a)), at first the visual and audio representations  $(F^v, F^a)$  are produced by encoders  $E_v$  and  $E_a$  respectively. Classification heads  $h_c^v$  and  $h_c^a$  and regression heads  $h_b^v$  and  $h_b^a$  are then used to predict the classification scores  $(p_t^v, p_t^a)$  and boundaries  $(s_t^v, e_t^v)$  and  $(s_t^a, e_t^a)$  for the visual and audio modalities, respectively. Thus, we can obtain a set of candidate proposals for the visual modality  $\Phi_v = \{(s_t^v, e_t^v, p_t^v)\}_{t=1}^T$  and similarly, a set of candidate proposals for the audio modality  $\Phi_a = \{(s_t^a, e_t^a, e_t^a)\}_{t=1}^T$ . Then, we concatenate these two sets as  $\Phi_o = \{(s_t^v, e_t^v, p_t^v), (s_t^a, e_t^a, p_t^a)\}_{t=1}^T$ .

**Classification scores fusion** – Although sounds can be associated with actions for classification purposes, the duration of an action does not necessarily correspond to its audio start and end as recently shown in [13]. Thus, we discard the audio boundaries, integrate the classification scores from both visual and audio modalities, and then use them along with the visual boundaries to generate proposals.

We use an approach similar to [ $\square$ ] to fuse visual and audio classification scores. Specifically, as shown in Figure 3(b), based on  $F^v$  and  $F^a$ , the visual classification head  $h_c^v$ , audio classification head  $h_c^a$  and visual boundary head  $h_b^v$  predict scores  $p_t^v$  and  $p_t^a$ , and frame boundaries  $s_t^v$  and  $e_t^v$ , respectively. We fuse the classification scores  $p_t^v$  and  $p_t^a$  by simple addition and combine them with the visual boundaries  $s_t^v$  and  $e_t^v$  to construct the set of fused candidate proposals  $\Phi_c$ , such that  $\Phi_c = \{(s_t^v, e_t^v, p_t^v + p_t^a)\}_{t=1}^T$ .

**Feature pyramid fusion** – In this strategy (see Figure 3(c)),  $F^{\nu}$  and  $F^{a}$  are put through a cross-attention mechanism [20, 53] (see also grey box of Figure 2) to model their inter-modal dependencies which then results in a single representation vector  $F^{a\nu}$ .

Firstly,  $F^{\nu}$  and  $F^{a}$  are projected into query  $Q = W_{Q}F^{\nu}$ , key  $K = W_{K}F^{a}$ , and value  $V = W_{V}F^{a}$ , where  $F^{\nu}$  serves as a query input and  $F^{a}$  serves as key and value inputs.  $W_{Q}$ ,  $W_{K}$  and  $W_{V} \in \mathbb{R}^{d \times d}$  are learnable weight matrices, where *d* is the embedding dimension. Next, we



Figure 3: Three strategies for fusing modalities - (a) Proposal fusion: visual and audio modalities are separately fed into their respective streams and generate corresponding sets of proposals, which are directly concatenated to obtain the final set of proposals, (b) Classification scores fusion: classification scores from both modalities are combined alongside boundaries from the video to obtain the final set of proposals, (c) Feature pyramid fusion: feature pyramids from the two modalities are fused through cross-attention and fed into parallel heads to predict boundaries and class scores, which then construct the set of proposals.

calculate the audio-visual representation vector

$$F^{av} = softmax(\frac{QK^T}{\sqrt{d}})V, \qquad (1)$$

which is then fed into a classification head  $h_c$  and a regression head  $h_b$  to obtain the classification scores  $p_t^{av}$  and the boundaries  $s_t^{av}$ ,  $e_t^{av}$  for each timestep. Therefore, the set of candidate proposals is  $\Phi_f = \{(s_t^{av}, e_t^{av}, p_t^{av}))\}_{t=1}^T$ .

We ablate these three strategies in Sec. 4.1. Based on the ablations, we chose feature pyramid fusion to generate audio-visual representations across N temporal pyramid scales for assessment by the centricity head to predict corresponding scores (see Sec. 3.2 next). We also selected the classification scores fusion approach to generate stronger audio-visual classification scores to predict action categories and calculate confidence scores.

### 3.2 Audio-visual Centricity Head

We investigated the relationship between the distance of a timestep from the action centre and the tIoU value between its generated proposal and the ground truth. As shown in Figure 4(a), as a timestep gets closer to the action center, its generated proposal has a higher tIoU value. This indicates that timesteps around the centre of an action can generate proposals with more reliable action boundaries. Thus, we propose a simple, yet effective, centricity head based on the audio-visual representation  $F^{av}$  to estimate how close the timestep t is to the centre of the action (as shown in Figure 4 (b)). The centricity head consists of three 1D conv layers with layer normalization and a ReLU activation function.

**Label assignment** – We require the centricity scores  $p_t^{C*}$  calculated from ground-truth data as supervision signals for training. For each timestep *t*, we consider the relative distance  $d_t$ 



Figure 4: The illustration of centricity – (a) The tIoUs between ground-truth and predictions (from RAB [12]) is plotted across various centre distances. Actions are divided into five groups based on segment lengths (in seconds): XS (0, 2], S (2, 4], M (4, 6], L (6, 8], and XL (8, inf). (b) The centricity head takes in the audio-visual feature  $F^{av}$  and produces the centricity scores  $p_t^C$ . The ground-truth centricity score  $p_t^{C*}$  is calculated based on the relative distance  $d_t$  between the time step t and the action center.

between the current timestep t and the centre of the corresponding ground-truth action to map the training labels of centricity scores

$$p_t^{C*} = \exp\left(-(d_t)^2/2\sigma^2\right)$$
, (2)

where  $\sigma$  is a scaling hyperparameter which defines that the closer a timestep is to the action centre, the higher the centricity score. This has previously been explored to predict boundary confidences [21, 42], but not specifically for centricity use cases. The centricity scores are normalized to a range of 0 to 1.

**Training** – We optimize the loss between the ground-truth  $p_t^{C*}$  and the predicted centricity scores  $p_t^{C}$  using Mean Square Error (MSE) loss as

$$L_{\mathcal{C}} = \frac{1}{T'} \sum_{t}^{T'} (p_t^{\mathcal{C}*} - p_t^{\mathcal{C}})^2 , \qquad (3)$$

where T' is the total number of timesteps used for training from all scales of the audio-visual feature pyramid. Our method can be integrated into any one-stage anchor-free frameworks and trained in an end-to-end manner. The total loss is  $L_{total} = L_g + \lambda_1 L_c + \lambda_2 L_b + \lambda_3 L_c$ , where  $L_g$  and  $L_c$  are losses for regression [52] and classification [52] and are the same as in [53, 52, 59].  $L_b$  is the boundary confidence loss from [52].  $\lambda_1$ ,  $\lambda_2$  and  $\lambda_3$  denote the loss balancing weights, and  $\lambda_3$  is set to 0 when the baseline is ActionFormer [53].

### 3.3 Post-processing

For each timestep *t*, the network produces the visual and audio classification scores  $p_t^v$ ,  $p_t^a$ , the corresponding action class label  $\alpha$ , a centricity score  $p_t^C$ , and a pair of starting and ending boundaries  $s_t$ ,  $e_t$  with their corresponding boundary confidences  $p_{s_t}^s$  and  $p_{e_t}^e$  (with these confidences computed as in [12]). The final confidence score for timestep *t* is then a weighted combination of the learnt knowledge, i.e.

$$S = p_t^v + \tau p_t^a + \beta p_t^c + \gamma (p_{s_t}^s + p_{e_t}^e), \qquad (4)$$

where  $\tau$ ,  $\beta$  and  $\gamma$  are fusion weights, and  $\gamma$  is set to 0 when the baseline is ActionFormer[ $\square$ ] or TriDet [ $\square$ ]. Finally, we follow standard practice [ $\square$ ,  $\square$ ,  $\square$ ,  $\square$ ] to rank these candidate

actions based on the final confidence score S and filter them using Soft-NMS [I] to obtain a final set of M predictions  $\Phi = \{(s, e, \alpha)_m\}_{m=1}^M$ .

### **4** Experiments

**Dataset** – We conduct experiments on EPIC-Kitchens-100 [**D**], a large-scale audio-visual dataset that contains 700 unscripted videos with 97 verb and 300 noun classes. On average, there are 128 action instances per video, with significant overlap.

**Evaluation metric –** We use mean Average Precision (mAP) for verb, noun and action tasks at various IoU thresholds {0.1, 0.2, 0.3, 0.4, 0.5} to evaluate our method against others.

**Baselines** – Our approach is integrated with three one-stage anchor-free approaches [59, 59, 19]. ActionFormer [59] is chosen as it is a pioneering anchor-free work that models longrange temporal dependencies using the Transformer for action detection. TriDet [59] extends this by incorporating scalable-granularity perception layers and a Trident head to regress boundaries. Finally, Wang et al. [50] (hereafter RAB) introduces a method to estimate boundary confidences through Gaussian scaling.

Again, following [59, 52, 59], the feature pyramid generated by the transformer encoder has N = 6 levels, with a level scaling factor of 2. For training, we use one Nvidia P100 GPU. We crop the video features with various lengths to 2304. The loss balancing weights in Sec. 3.2 are  $\lambda_1 = 1$ ,  $\lambda_2 = 0.5$ , and  $\lambda_3 = 1.7$ . The weight ratio of the classification loss between verb and noun is set to 2:3. The scaling hyperparameter in Eq. (2) is set to  $\sigma = 1.7$ . During the inference stage, the confidence score weights in Sect. 3.3 are assigned as  $\tau = 0.2$ ,  $\beta = 1$ , and  $\gamma = 0.7$ . For the multi-task classification, we select the top 11 verb and the top 33 noun predictions to combine the candidate actions.

**Main results** – Table 1 shows that the proposed method outperforms recent SOTA approaches [9, 124, 26, 155, 159, 122, 159] on the EPIC-Kitchens-100 action detection benchmark, and achieves significant improvements when added to existing SOTA one-stage multi-scale methods [159, 122, 159]. ActionFormer [159] and TriDet [159] train different models for verb and noun detection and do not detect actions. Instead, we train one model and add a multi-task classification head to predict their results for the action task. It can be seen that our enhanced proposed method with audio fusion and centricity improves performance in every metric. In action detection, mAP improves by 1.35% and 0.97%, respectively. RAB [123] also performs well on egocentric data and follows the same anchor-free pipeline as ours. The improvement achieved on RAB [123] was 1.32% and is also the best result amongst the baselines.

**Qualitative results** – Qualitative plots of 'RAB [ $\blacksquare$ ]+Ours' on the EPIC-Kitchens-100 action detection validation dataset are shown in Figure 5. The bottom two lines showcase the model's ability to detect dense actions with different classes and durations in a video, demonstrating that our approach can effectively utilise the audio modality to learn discriminative representations. The middle three lines show a zoomed-in, detailed look where it is

Method	Venue	Feature	Audio	Avg. mAP@task		
				Verb	Noun	Action
BMN [ <b>9</b> , <b>26</b> ]	IJCV 2022	SF [ <b>1</b> ]	×	8.36	6.53	5.21
OWL [	ArXiv 2022	SF [🎞]	$\checkmark$	11.47	12.63	8.35
BMN+TSN [26]	ICLR 2022	TSN 🖾	×	13.47	12.37	9.71
BMN+TAda2D [12, 16]	ICLR 2022	TAda2D [🗖]	×	16.78	17.39	13.18
ActionFormer [	ECCV 2022	SF [🎞]	×	20.45	20.90	16.63
ActionFormer [11]* + Ours	-	SF [🛄]	$\checkmark$	20.48	22.41	17.98
TriDet [55]*	CVPR 2023	SF [🎞]	Х	20.87	21.04	17.21
TriDet [1]* + Ours	-	SF [🎞]	$\checkmark$	21.94	22.86	18.18
RAB 🛄	AVSS 2022	SF [🛄]	Х	20.71	20.53	17.18
RAB 🛄 + Ours	-	SF [Щ]	$\checkmark$	21.10	23.08	18.50

Table 1: Comparative results on the EPIC-Kitchens-100 action detection validation set - \*Actionformer and TriDet only provide results on verb and noun detection, hence we produce action results by modifying them with a multi-task action classification head [12].



Figure 5: Qualitative results on the EPIC-Kitchens-100 action detection dataset – The top row shows the visual content of selected frames. The middle seven lines display the ground-truth (GT), the predictions of RAB [12], RAB [12]+Ours, TriDet [13], TriDet [13]+Ours, AF [13], AF [13]+ours for a zoomed-in region. The bottom seven lines represent the whole video sequence. These results demonstrate the effectiveness of our method in accurately detecting dense actions.

easier to see that our method better deals with challenging actions, e.g. see action 'put fork' in the  $77^{th}$  second of the first video and the action 'wash hand' in the  $68^{th}$  second of the second video that were missed by the baseline model. This indicates that our centricity head enhances the confidence scores for actions with more precise boundaries, resulting in their preferential ranking and selection during the Soft-NMS processing.

### 4.1 Ablations

All our ablations use RAB [2] as baseline and are performed on EPIC-Kitchens-100 validation dataset [2].

Basalina	Audio	Centri-	Avg. mAP@task			
basenne		city	Verb	Noun	Action	
ActionFormer [	×	×	20.45	20.90	16.63	
ActionFormer [	√	×	19.75	21.46	17.22	
ActionFormer [11] + Ours	×	$\checkmark$	21.50	22.25	17.60	
ActionFormer [11] + Ours	$\checkmark$	$\checkmark$	20.48	22.41	17.98	
TriDet [	Х	×	20.87	21.04	17.21	
TriDet [🛄] + Ours	√	×	21.07	21.75	17.61	
TriDet [🛄] + Ours	×	$\checkmark$	21.65	21.23	17.42	
TriDet [ 🛄 + Ours	√	$\checkmark$	21.94	22.86	18.18	
RAB 🖾	×	×	20.71	20.53	17.18	
RAB 🛄 + Ours	√	×	20.93	21.84	17.88	
RAB [	×	$\checkmark$	21.47	21.43	17.69	
RAB 🖾] + Ours	✓	$\checkmark$	21.10	23.08	18.50	

Table 2: Components analysis on the EPIC-Kitchens-100 action detection validation set.

Classification	Avg. mAP@task			
Weight $\lambda_1$	Verb	Noun	Action	
0.5	22.04	22.70	18.22	
1	21.10	23.08	18.50	
2	22.13	22.52	18.01	
4	22.36	22.11	18.09	
6	22.21	22.75	17.83	
8	22.20	22.92	17.83	



Table 3: Ablation on varying classification weight  $\lambda_1$ .



**Components ablation** – We ablate the contributions of our two main components, audiovisual fusion and centricity head, as seen in Table 2. Both components demonstrate notable improvements in the performance of the baseline methods [59, 52, 59], particularly when they are engaged, singularly or in combination, for action detection.

Loss function weights – We train our network in an end-to-end manner by minimizing the total loss function in Sec. 3.2, where we assign three weights to balance various losses. Table 3 shows that  $\lambda_1 = 1$  is the best value for performing the action task for the classification loss  $L_c$  weight, across the improvements made to the baseline. For the boundary confidence loss  $L_b$  weight  $\lambda_2$ , we follow the baseline RAB's recommended setting [12]. Finally, Table 4 demonstrates that our results are relatively stable when varying centricity loss  $L_c$  weight  $\lambda_3$  between 0.5 - 2, with the best action detection result at  $\lambda_3 = 1.7$ .

**The effect of centricity on confidence score** – Figure 6 shows that as the centre distance increases, the average tIoU values (blue bars) between the ground truth and the proposals generated by timesteps at the centre distance present a notable decreasing trend ( $\downarrow$ 0.16). However, the original confidence scores (dashed blue line) only slightly decreases ( $\downarrow$ 0.03). Adding centricity into the confidence scores (solid blue line) responds better to the expected trend ( $\downarrow$ 0.09). Thus, proposals with more accurate boundaries (higher tIoU values) rank higher based on their confidence scores when centricity is incorporated.

Audio-visual fusion strategies – The three strategies for fusing audio and visual modalities (see Section 3.1) are compared in Table 5. The proposal fusion strategy has the lowest overall action detection performance due to the audio stream's proposals having less precise boundaries. The audio-visual fusion of classification scores strategy improves on Visual-only through multiplication ( $\uparrow 0.19\%$ ) or addition ( $\uparrow 0.51\%$ ). For the feature pyramid fusion, while a direct concatenation achieves a relative increase ( $\uparrow 0.27\%$ ), the cross-attention mechanism provides a more significant learning opportunity ( $\uparrow 0.70\%$ ).

Centricity vs. action-ness - Centricity establishes how close the current timestep is from

10 WANG ET AL: CENTRICITY-BASED AUDIO-VISUAL TEMPORAL ACTION DETECTION

Centricity	Avg. mAP@task				
Weight $\lambda_3$	Verb	Noun	Action		
0.5	21.18	22.82	18.04		
1	21.19	22.02	18.11		
1.5	21.23	23.24	18.24		
1.6	21.71	22.89	18.16		
1.7	21.10	23.08	18.50		
1.8	21.16	22.77	18.13		
2	21.27	22.90	18.22		

Fusion Stratogies		Avg. mAP@task			
r usion strategies	Verb	Noun	Action		
Visual-only	20.71	20.53	17.18		
Audio-only	7.94	6.41	4.57		
Proposals fusion (Figure 3 (a))	20.36	20.38	16.65		
Classif. scores fusion (multiplication) (Fig 3(b))	20.93	21.82	17.37		
Classif. scores fusion (addition) (Fig 3(b))	20.18	21.89	17.69		
Feature pyramid fusion (concatenation) (Fig 3(c))	20.95	21.52	17.45		
Feature pyramid fusion (cross-attention) (Fig 3(c))	20.93	21.84	17.88		

Table 4: Ablation on varying centricity weight  $\lambda_3$ .

Table 5: Comparing different strategies to fuse audio & visual modalities.



Figure 7: Visualization examples of tIoU values, and centricity and action-ness scores – The x-axis represents the re-scaling of the temporal dimension of an action segment to the range of [0-1]. The bars represent the average tIoU between the proposals generated by corresponding timesteps and their ground truth. The green and purple lines are the centricity and action-ness scores, respectively.

the action centre, while action-ness  $[\Box]$  represents the probability of the action occurring. Figure 7 displays individual instances of actions with centricity scores, action-ness scores and the average tloUs between the proposals generated by corresponding timesteps and their ground truth. In the middle of the action, timesteps tend to exhibit peak values of tloUs that gradually decrease on both sides, and a similar trend is also observed in the centricity scores. This suggests that timesteps associated with higher centricity scores are inclined to generate proposals with more precise boundaries. In contrast, the action-ness scores (purple line) tend to drop in the middle of the action.

# 5 Conclusion

We introduced an audio-visual fusion approach and a novel centricity head for one-stage anchor-free action detectors. Our method achieves state-of-the-art results on the large-scale egocentric EPIC-Kitchens-100 action detection benchmark where audio and video streams are available. Detailed ablations demonstrated the benefits of fusing audio and visual modalities and emphasized the importance of centricity scores.

Many questions about the use of multi-modalities in temporal action detection remain unexplored, such as the discrepancies in the training data and the temporal misalignment between different modalities. An extension of our work would involve joint learning from visual, audio, and language modalities to enhance action detection performance, with specific focus on mitigating the misalignment among these three modalities and developing novel fusion techniques to provide discriminative representations.

## References

- Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Selfsupervised learning of audio-visual objects from video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 208–224, 2020.
- [2] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 609–617, 2017.
- [3] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. SoundNet: Learning sound representations from unlabeled video. In Advances in Neural Information Processing Systems (NeurIPS), pages 892–900, 2016.
- [4] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S. Davis. Soft-NMS improving object detection with one line of code. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5562–5570, 2017.
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 213–229, 2020.
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the Kinetics dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision* and Pattern Recognition (CVPR), pages 6299–6308, 2017.
- [7] Shuning Chang, Pichao Wang, Fan Wang, Hao Li, and Zheng Shou. Augmented transformer with adaptive graph for temporal action proposal generation. In *Proceedings of the 3rd International Workshop on Human-Centric Multimedia Analysis*, pages 41–50, 2022.
- [8] Feng Cheng and Gedas Bertasius. TallFormer: Temporal action localization with a long-memory transformer. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 503–521, 2022.
- [9] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision. *International Journal of Computer Vision (IJCV)*, pages 33–55, 2022.
- [10] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision (ICCV), pages 6202–6211, 2019.
- [11] Jiyang Gao, Zhenheng Yang, Kan Chen, Chen Sun, and Ram Nevatia. TURN TAP: Temporal unit regression network for temporal action proposals. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3628–3636, 2017.

#### 12 WANG ET AL: CENTRICITY-BASED AUDIO-VISUAL TEMPORAL ACTION DETECTION

- [12] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4D: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18995–19012, 2022.
- [13] Di Hu, Feiping Nie, and Xuelong Li. Deep multimodal clustering for unsupervised audiovisual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9248–9257, 2019.
- [14] Ziyuan Huang, Shiwei Zhang, Liang Pan, Zhiwu Qing, Mingqian Tang, Ziwei Liu, and Marcelo H Ang Jr. TAda! temporally-adaptive convolutions for video understanding. In *Proceedings of International Conference on Learning Representations (ICLR)*, pages 1–23, 2022.
- [15] Jaesung Huh, Jacob Chalk, Evangelos Kazakos, Dima Damen, and Andrew Zisserman. EPIC-SOUNDS: A Large-Scale Dataset of Actions that Sound. In *IEEE International Conference on Acoustics, Speech, Signal Processing (ICASSP)*, 2023.
- [16] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. EPIC-Fusion: Audio-visual temporal binding for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5492–5501, 2019.
- [17] Evangelos Kazakos, Jaesung Huh, Arsha Nagrani, Andrew Zisserman, and Dima Damen. With a little help from my temporal context: Multimodal egocentric action recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2021.
- [18] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Slow-Fast auditory streams for audio recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 855–859. IEEE, 2021.
- [19] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. Advances in Neural Information Processing Systems (NeurIPS), pages 7774–7785, 2018.
- [20] Jun-Tae Lee, Mihir Jain, Hyoungwoo Park, and Sungrack Yun. Cross-attentional audiovisual fusion for weakly-supervised action localization. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2021.
- [21] Yanghao Li, Cuiling Lan, Junliang Xing, Wenjun Zeng, Chunfeng Yuan, and Jiaying Liu. Online human action detection using joint classification-regression recurrent neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 203–220, 2016.
- [22] 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 *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 11499–11506, 2020.

- [23] Chuming Lin, Chengming Xu, Donghao Luo, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yanwei Fu. Learning salient boundary feature for anchorfree temporal action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3320–3329, 2021.
- [24] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In Proceedings of the 25th ACM International Conference on Multimedia, pages 988– 996, 2017.
- [25] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. BSN: Boundary sensitive network for temporal action proposal generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.
- [26] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. BMN: Boundary-matching network for temporal action proposal generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3889–3898, 2019.
- [27] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision (ICCV), pages 2980–2988, 2017.
- [28] Qinying Liu and Zilei Wang. Progressive boundary refinement network for temporal action detection. In *Proceedings of the AAAI Conference on Artificial Intelligence* (AAAI), pages 11612–11619, 2020.
- [29] 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*, pages 5427–5441, 2022.
- [30] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 344–353, 2019.
- [31] Shentong Mo and Yapeng Tian. Multi-modal grouping network for weakly-supervised audio-visual video parsing. In *Advances in Neural Information Processing Systems* (*NeurIPS*), pages 34722–34733, 2022.
- [32] Arsha Nagrani, Chen Sun, David Ross, Rahul Sukthankar, Cordelia Schmid, and Andrew Zisserman. Speech2Action: Cross-modal supervision for action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), pages 10317–10326, 2020.
- [33] Megha Nawhal and Greg Mori. Activity graph transformer for temporal action localization. *arXiv preprint arXiv:2101.08540*, 2021.
- [34] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision* (ECCV), pages 631–648, 2018.
- [35] Merey Ramazanova, Victor Escorcia, Fabian Caba Heilbron, Chen Zhao, and Bernard Ghanem. OWL (Observe, Watch, Listen): Localizing actions in egocentric video via audiovisual temporal context. arXiv preprint arXiv:2202.04947, 2022.

#### 14 WANG ET AL: CENTRICITY-BASED AUDIO-VISUAL TEMPORAL ACTION DETECTION

- [36] Varshanth Rao, Md Ibrahim Khalil, Haoda Li, Peng Dai, and Juwei Lu. Dual perspective network for audio-visual event localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 689–704, 2022.
- [37] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 658–666, 2019.
- [38] Dingfeng Shi, Yujie Zhong, Qiong Cao, Jing Zhang, Lin Ma, Jia Li, and Dacheng Tao. ReAct: Temporal action detection with relational queries. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 105–121, 2022.
- [39] Dingfeng Shi, Yujie Zhong, Qiong Cao, Lin Ma, Jia Li, and Dacheng Tao. TriDet: Temporal action detection with relative boundary modeling. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18857–18866, 2023.
- [40] Haisheng Su, Weihao Gan, Wei Wu, Junjie Yan, and Yu Qiao. BSN++: Complementary boundary regressor with scale-balanced relation modeling for temporal action proposal generation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 2602–2610, 2021.
- [41] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 247–263, 2018.
- [42] Hanyuan Wang, Majid Mirmehdi, Dima Damen, and Toby Perrett. Refining action boundaries for one-stage detection. In 18th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), pages 1–8. IEEE, 2022.
- [43] Huiyu Wang, Mitesh Kumar Singh, and Lorenzo Torresani. Ego-Only: Egocentric action detection without exocentric pretraining. arXiv preprint arXiv:2301.01380, 2023.
- [44] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 20–36, 2016.
- [45] Yu Wu and Yi Yang. Exploring heterogeneous clues for weakly-supervised audiovisual video parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1326–1335, 2021.
- [46] Yan Xia and Zhou Zhao. Cross-modal background suppression for audio-visual event localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19989–19998, 2022.
- [47] Huijuan Xu, Abir Das, and Kate Saenko. R-C3D: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision (ICCV), pages 5783–5792, 2017.

- [48] Le Yang, Houwen Peng, Dingwen Zhang, Jianlong Fu, and Junwei Han. Revisiting anchor mechanisms for temporal action localization. *IEEE Transactions on Image Processing*, pages 8535–8548, 2020.
- [49] Chen-Lin Zhang, Jianxin Wu, and Yin Li. ActionFormer: Localizing moments of actions with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 492–510, 2022.
- [50] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2914–2923, 2017.