

Hear Me out: Fusional Approaches for Audio Augmented Temporal Action Localization

Anurag Bagchi¹^a, Jazib Mahmood¹^b, Dolton Fernandes¹^c and Ravi Kiran Sarvadevabhatla²^d

¹International Institute of Information Technology, Hyderabad, India

²Center for Visual Information Technology (CVIT), IIT Hyderabad, India

Keywords: Temporal Activity Localization, Graph Convolution Networks, Multi-modal Fusion, Audio.

Abstract: State of the art architectures for untrimmed video Temporal Action Localization (TAL) have only considered RGB and Flow modalities, leaving the information-rich audio modality unexploited. Audio fusion has been explored for the related but an arguably easier problem of trimmed (clip-level) action recognition. However, TAL poses a unique set of challenges. In this paper, we propose simple but effective fusion-based approaches for TAL. To the best of our knowledge, our work is the first to jointly consider audio and video modalities for supervised TAL. We experimentally show that our schemes consistently improve performance for the state-of-the-art video-only TAL approaches. Specifically, they help achieve a new state-of-the-art performance on large-scale benchmark datasets - ActivityNet-1.3 (54.34 mAP@0.5) and THUMOS14 (57.18 mAP@0.5). Our experiments include ablations involving multiple fusion schemes, modality combinations, and TAL architectures. Our code, models, and associated data are available at <https://github.com/skelemao/tal-hmo>.

1 INTRODUCTION

With the boom in online video production, video understanding has become one of the most heavily researched domains. Temporal Action Localization (TAL) is one of the most interesting and challenging problems in the domain. The objective of TAL is to identify the category (class label) of activities present in a long, untrimmed, real-world video and their temporal boundaries (start and end time). Apart from inheriting the challenges from the related problem of *trimmed* (clip-level) video action recognition, TAL also requires accurate temporal segmentation, i.e. to precisely locate the start time and end time of action categories present in a given video.


TAL is an active area of research and several approaches have been proposed to tackle the problem (Xu et al., 2020b; Zeng et al., 2019b; Alwassel et al., 2020; Liu et al., 2021; Liu et al., 2021; Lin et al., 2021). For the most part, existing approaches depend solely on the visual modality (RGB, Optical Flow). An important and obvious source of additional


information – the audio modality – has been overlooked. This is surprising since audio has been shown to be immensely useful for other video-based tasks such as object localization (Arandjelovic and Zisserman, 2017), action recognition (Owens and Efros, 2018; Wu et al., 2016; Long et al., 2018b; Long et al., 2018a; Owens et al.,) and egocentric action recognition (Kazakos et al., 2019).


Analyzing the untrimmed videos, it is evident that the audio track provides crucial complementary information regarding the action classes and their temporal extents. Action class segments in untrimmed videos are often characterized by signature audio transitions as the activity progresses (e.g. the rolling of a ball in a bowling alley culminating in striking of pins, an aquatic diving event culminating with the sound of a splash in the water). Depending on the activity, the associated audio features can supplement and complement their video counterparts if the two feature sequences are fused judiciously (Figure 1).


Motivated by the observations mentioned above, we make the following contributions:

- We propose simple but effective fusion approaches to combine audio and video modalities for TAL (Section 3). Our work is the first to jointly process audio and video modalities for supervised TAL.

^a <https://orcid.org/0000-0002-9893-9643>

^b <https://orcid.org/0000-0002-0076-1315>

^c <https://orcid.org/0000-0002-8926-1156>

^d <https://orcid.org/0000-0003-4134-1154>



Ground Truth	1:48 sec	Having a conversation	2:59 sec	2:59 sec	Bowling	3:17 sec	
Visual prediction	1:48 sec	Bowling	2:10 sec	Having a conversation	2:58 sec	Bowling	3:20 sec
Audio prediction	1:48 sec	Having a conversation	2:55 sec	2:55 sec	Bowling	3:19 sec	
Fusion prediction	1:48 sec	Having a conversation	3:00 sec	3:00 sec	Bowling	3:18 sec	

Figure 1: An example illustrating a scenario where audio modality can help improve performance over video-only temporal action localization.

- We show that our fusion schemes can be readily plugged into existing state-of-the-art video-based TAL pipelines (Section 3).
- To determine the efficacy of our fusional approaches, we perform a comparative evaluation on large-scale benchmark datasets - ActivityNet and THUMOS14. Our results (Section 6) show that the proposed schemes consistently boost performance for the state-of-the-art TAL approaches, resulting in an improved mAP of 52.73 for ActivityNet-1.3 and 57.18 mAP for THUMOS14.
- Our experiments include ablations involving multiple fusion schemes, modality combinations, and TAL architectures.

Our code, models, and associated data are available at <https://github.com/skelemoa/tal-hmo>.

2 RELATED WORK

Temporal Action Localization: A popular technique for Temporal Action Localization is inspired from the proposal-based approach for object detection (Girshick, 2015). In this approach, a set of so-called proposals are generated and subsequently refined to produce the final class and boundary predictions. Many recent approaches employ this proposal-based formulation (Shou et al., 2016; Zhao et al., 2017; Xu et al., 2017). Specifically, this is the case for state-of-the-art approaches we consider in this paper – G-TAD (Xu et al., 2020b), PGCN (Zeng et al., 2019b) and MUSES baseline (Liu et al., 2021). Both G-TAD (Xu et al., 2020b) and PGCN (Zeng et al., 2019b) use graph convolutional networks and the concept of edges to share context and background information between proposals. MUSES baseline (Liu et al., 2021) on the other hand, achieves the state-of-the-art results on the benchmark datasets by employing a temporal aggregation module, originally in-

tended to account for the frequent camera changes in their new multi-shot dataset.

The proposal generation schemes in the literature are either anchor-based (Gao et al., 2017; Gao et al., 2018; Liu et al., 2018) or generate a boundary probability sequence (Lin et al., 2018; Lin et al., 2019; Su et al., 2020). Past work in this domain also includes end to end techniques (Buch et al., 2019; Yeung et al., 2015; Lin et al., 2017) which combine the two stages. Frame-level techniques which require merging steps to generate boundary predictions also exist (Shou et al., 2017; Montes et al., 2016; Zeng et al., 2019a). We augment the proposal-based state of the art approaches designed solely for visual modality by incorporating audio into their architectures.

Audio-only based Localization: Speaker diarization (Wang et al., 2018; Zhang et al., 2019) involves localization of speaker boundaries and grouping segments that belong to the same speaker. The DCASE Challenge (Ono et al., 2020) examines sound event detection in domestic environments as one of the challenge tasks (Miyazaki et al., 2020; Hao et al., 2020; Ebberts and Haeb-Umbach, 2020; Lin and Wang, 2019; Delphin-Poulat and Plapous, 2019; Shi, 2019). In our action localization setting, note that the audio modality is unrestricted. It is not confined to speech or labeled sound events which is the case for audio-only localization.

Audio-visual Localization: This task is essentially localization of specific events of interest across modalities. Given a temporal segment of one modality (auditory or visual), we would like to localize the temporal segment of the associated content in the other modality. This is different from our task of temporal action localization (TAL) which focuses on predicting the class labels and temporal segment boundaries of all actions present in a given video.

Fusion Approaches for TAL: Fusion of multiple modalities is an effective technique for video understanding tasks due to its ability to incorporate all the information available in videos. The fusion schemes

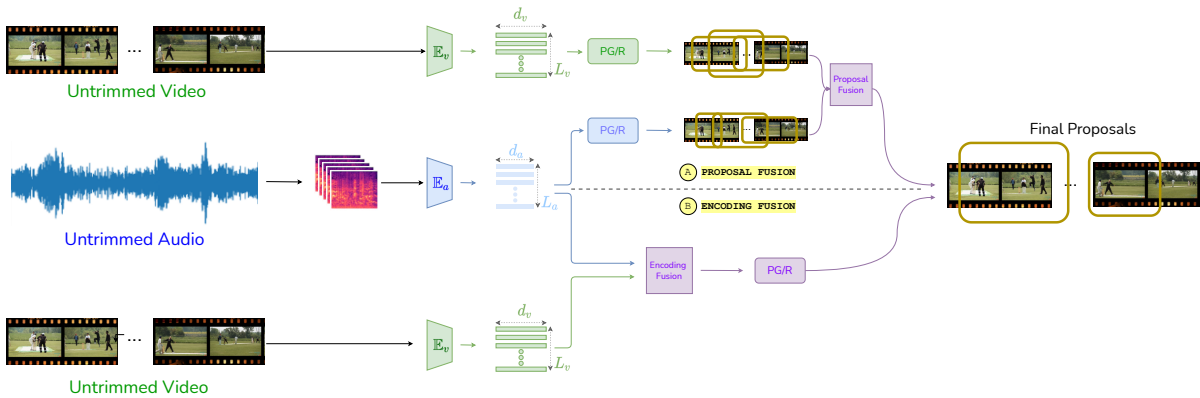


Figure 2: An illustrative overview of our fusion schemes (Section 3).

present in the literature can be divided into 3 broad categories – *early fusion*, *mid fusion* and *late fusion*.

Late fusion combines the representations closer to the output end from each individual modality stream. When per-modality predictions are fused, this technique is also referred to as *decision level fusion*. *Decision level fusion* is used in I3D (Carreira and Zisserman, 2017) which serves as a feature extractor for the current state-of-the-art in TAL. However, unlike the popular classification setting, *decision level fusion* is challenging for TAL since predictions often differ in relative temporal extents. PGCN (Zeng et al., 2019b), introduced earlier, solves this problem by performing Non-Maximal Suppression on the combined pool of proposals from the two modalities (RGB, Optical Flow). MUSES baseline (Liu et al., 2021) fuses the RGB and Flow predictions. *Mid fusion* combines mid-level feature representations from each individual modality stream. Feichtenhofer et. al. (Feichtenhofer et al., 2016) found that fusing RGB and Optical Flow streams at the last convolutional layer yields good visual modality features. The resulting mid-level features have been successfully employed by well-performing TAL approaches (Lin et al., 2020; Lin et al., 2019; Lin et al., 2018; Li et al., 2019). In particular, they are utilized by G-TAD (Zeng et al., 2019b) to obtain feature representations for each temporal proposal. *Early fusion* involves fusing the modalities at the input level. In the few papers that compare different fusion schemes (Jiang et al., 2018; Tian et al., 2018), *early fusion* is generally an inferior choice.

Apart from within (visual) modality fusion mentioned above, audio-visual fusion specifically has been shown to benefit (trimmed) clip-level action recognition (Wu et al., 2016; Long et al., 2018b; Long et al., 2018a; Kazakos et al., 2019) and audio-visual event localization (Zhou et al., 2021; He et al., 2021; Xu et al., 2020a). The audio modality has also been

shown to be beneficial for the weakly supervised version of TAL (Lee et al., 2021) wherein the boundary labels for activity instances are absent. However, the lack of labels is a fundamental performance bottleneck compared to the supervised approach.

In our work, we introduce two *mid-level fusion* schemes along with *decision level fusion* to combine Audio, RGB, Flow modalities for state-of-the-art supervised TAL.

3 PROPOSED FUSION SCHEMES

Processing Stages in Video-only TAL: Temporal Action Localization can be formulated as the task of predicting start and end times (t_s, t_e) and action label a for each action in an untrimmed RGB video $V \in \mathbb{R}^{F \times 3 \times H \times W}$, where F is the number of frames, H and W represent the frame height and width. Despite the architectural differences, state of the art TAL approaches typically consist of three stages: *feature extraction*, *proposal generation* and *proposal refinement*.

The *feature extraction* stage transforms a video into a sequence of feature vectors corresponding to each visual modality (RGB and Flow). Specifically, the feature extractor operates on fixed-size snippets $S \in \mathbb{R}^{L \times C \times H \times W}$ and produces a feature vector $f \in \mathbb{R}^{d_v}$. Here, C is the number of channels and L is the number of frames in the snippet. This results in the feature vector sequence $F_v \in \mathbb{R}^{L_v \times d_v}$ mentioned above where L_v is the number of snippets. This stage is shown as the component shaded green (containing E_v) in Figure 2.

The *proposal generation* stage processes the feature sequence mentioned above to generate action proposals. Each candidate action proposal is associated with temporal extent (start and end time) relative to the input video, and a confidence score. In some

approaches, each proposal is also associated with an activity class label.

The *proposal refinement* stage takes the feature sequence corresponding to each candidate proposal as input and refines the boundary predictions and confidence scores. Some approaches modify the class label predictions as well. The final proposals are generally obtained by applying nonmaximal suppression to weed out the redundancies arising from highly overlapping proposals. Also, note that some approaches do not treat proposal generation and refinement as two different stages. To accommodate this variety, we depict the processing associated with proposal generation and refinement as a single module titled ‘PG/R’ in Figure 2.

3.1 Incorporating Audio

As with video-only TAL approaches, the first stage of audio modality processing consists of extracting a sequence of features from audio snippets (refer to the blue shaded module termed \mathbb{E}_a in Figure 2). This results in the ‘audio’ feature vector sequence $F_a \in \mathbb{R}^{L_a \times d_a}$ mentioned above where L_a is the number of audio snippets and d_a is the length of feature vector for an audio snippet.

Our objective is to incorporate audio as seamlessly as possible into existing video-only TAL architectures. To enable flexible integration, we present two schemes - *proposal fusion* and *encoding fusion*. We describe these schemes next.

3.2 (A) Proposal Fusion

This is a *decision fusion* approach and as suggested by its name, the basic idea is to merge proposals from the audio and video modalities (see ‘Proposal Fusion’ in Figure 2). To begin with, audio proposals are obtained like the procedure used to obtain video proposals. As mentioned earlier, it is straightforward to fuse action class predictions from each modality by simply averaging the probabilities. However, TAL proposals consist of regression scores for action boundaries, where averaging across modality does not make much sense since it is likely to introduce error into the predictions. This makes the fusion task challenging in the TAL setting.

To solve this problem while adhering to our objective of leaving the existing video-only pipeline untouched, we repurpose the corresponding module from the pipelines. Specifically, we use Non-Maximal Suppression (NMS) for iteratively choosing the best proposals which minimally overlap with other proposals. In some architectures (e.g.




Figure 3: Residual Multimodal Attention mechanism with video-only and audio features as a form of encoding fusion (Section 3.3). \oplus indicates tensor addition.

PGCN (Zeng et al., 2019b)), NMS is applied to separate proposals from RGB and Flow components which contribute together as part of the visual modality. We extend this, by initially pooling visual modality proposals with audio proposals, and then applying NMS.

3.3 (B) Encoding Fusion

Instead of the late fusion of per-modality proposals described above, an alternative is to utilize the combined set of audio F_a and video feature sequences F_v to generate a single, unified set of proposals. However, since the encoded representation dimensions d_a, d_v and the number of sequence elements L_a, L_v can be unequal, standard dimension-wise concatenation techniques are not applicable. To tackle this issue, we explore four approaches to make the sequence lengths equal for feature fusion (depicted as purple block titled ‘Encoding Fusion’ in Figure 2).

For the first two approaches, we revisit the feature extraction phase and extract audio features at the frame rate used for videos. As a result, we obtain a paired sequence of audio and video snippets (i.e. $L_a = L_v$).

- **Concatenation (Concat):** The paired sequences are concatenated along the feature dimension.
- **Residual Multimodal Attention (RMAttn):** To refine each modality’s representation using features from other modalities, we employ a residual multimodal attention mechanism (Tian et al., 2018) as shown in Figure 3.

$$a^{t'} = \sigma(a^t + f(a^t, v^t)) \quad (1)$$

$$v^{t'} = \sigma(v^t + f(a^t, v^t)) \quad (2)$$

where $f(\cdot)$ is an additive fusion function, $\sigma(\cdot)$ is the hyperbolic tangent function, and \cdot , and the sum of $a^{t'}$ and $v^{t'}$ is used as the joint representation for the video feature at the timestamp t .

The other two encoding fusion approaches:

- **Duplicate and Trim (DupTrim):** Suppose $L_v < L_a$ and $k = \frac{L_a}{L_v}$. We first duplicate each visual feature to obtain a sequence of length kL_v . We then trim both the audio and visual feature sequences to a common length $L_m = \min(L_a, kL_v)$. A similar procedure is followed for the other case ($L_a < L_v$).
- **Average and Trim (AvgTrim):** Suppose $L_v < L_a$ and $k = \frac{L_a}{L_v}$. We group audio features into subsequences of length $k' = \lceil k \rceil$. We then form a new feature sequence of length L'_a by averaging each k' -sized group. Following a procedure similar to ‘DupTrim’ above, we trim the modality sequences to a common length, i.e. $L_m = \min(L'_a, L_v)$.

For the above approaches involving trimming, the resulting audio and video sequences are concatenated along the feature dimension to obtain the fused multi-modal feature sequence.

For all the approaches, the resulting fused representation is processed by a ‘PG/R’ (proposal generation and refinement) module to obtain the final predictions, similar to its usage mentioned earlier (see Figure 2). Apart from the fusion schemes, we also note that each scheme involves additional choices. In our experiments, we perform a comparative evaluation for all the resulting combinations.

4 IMPLEMENTATION DETAILS

TAL Video-only Architectures: To determine the utility of audio modality and to ensure our method is easily applicable to any video-only TAL approach, we do not change the architectures and hyperparameters (e.g. snippet length, frame rate, optimizers) for the baselines. The feature extraction components of the baselines are summarized in Table 1.

Audio Extraction: For audio, we use VGGish (Hershey et al., 2017), a state of the art approach for audio feature extraction. We use a sampling rate of 16kHz to extract the audio signal and extract 128-D features from 1.2s long snippets. For experiments involving attentional fusion and simple concatenation, we extract features by centering the 1.2s window about the snippets used for video feature extraction to maintain the same feature sequence length for audio and video modalities. Windows shorter than 1.2s (a few starting and ending ones) are padded with zeros at the end to specify that no more information is present and to match the 1.2s window requirement. Although the opposite (i.e. changing sampling rate for video, keeping the audio setup unchanged) is possible, we prefer the former since the video-only architecture and

(video) data processing can be used as specified originally, without worrying about the consequences of such a change on the existing video-only architectural setup and hyperparameter choices.

Proposal Generation and Refinement (PG/R): For proposal generation, we consider state of the art architectures GTAD(Xu et al., 2020b), BMN(Lin et al., 2019) and BSN(Lin et al., 2018). Similarly, for proposal refinement, we consider proposals generated from BMN and BSN, refined in PGCN(Zeng et al., 2019b) and MUSES(Liu et al., 2021) in our experiments with audio.

Optimization: We train all the architectures except PGCN and GTAD with their original setting. We use a batch size of 256 for PGCN and 16 for GTAD. For training, we use 4 GeForce 2080Ti 11GB GPUs. The entire codebase is based on the Pytorch library except for VGGish (Hershey et al., 2017) which is based on Keras.

5 EXPERIMENTS

5.1 Datasets

To compare our results with the SOTA architectures in which we incorporate audio, we evaluate our models on two benchmark datasets for temporal action localization.

Thumos14: (Jiang et al., 2014) contains 1010 untrimmed videos for validation and 1574 for testing. Of these, 200 validation and 213 testing videos contain temporal annotations spanning 20 activity categories. Following the standard setup (Xu et al., 2020b; Zeng et al., 2019b), we use the 200 validation videos for training and the 213 testing videos for evaluation.

ActivityNet-1.3: (Caba Heilbron et al., 2015) contains 20k untrimmed videos with 200 action classes between its training, validation and testing sets. Once again, following the standard setup (Xu et al., 2020b; Zeng et al., 2019b), we train on 10024 videos and test on the 4926 videos from the validation set.

5.2 Evaluation Protocol

We label a temporal (proposal) prediction (with associated start and end time) as correct if (i) its Intersection Over Union (IOU) with ground-truth exceeds a pre-determined threshold (ii) the proposal’s label matches the ground truth counterpart. Following standard protocols, we evaluate the mean Average Precision (mAP) scores at IOU thresholds from 0.3 to 0.7 with a step of 0.1 for THUMOS14 (Jiang et al., 2014)

Table 1: Architectural pipeline components for top-performing TAL approaches. To reduce clutter, only mAP@0.5 is reported. The ‘+Audio’ group refers to the fusion configuration corresponding to the best results (Section 3).

Dataset	Setup Id	Architecture	Visual Features	+Audio			
				mAP	mAP	fusion scheme	fusion type
THUMOS14(Jiang et al., 2014)	1	GTAD(Xu et al., 2020b)	TSN(Wang et al., 2019)	40.20	42.01	encoding	Concat
	2	PGCN(Zeng et al., 2019b)	TSP (Alwassel et al., 2020)	53.50	53.96	encoding	AvgTrim
	3	MUSES(Liu et al., 2021)	I3D(Carreira and Zisserman, 2017)	56.16 ¹	57.18	encoding	Concat
ActivityNet-1.3(Caba Heilbron et al., 2015)	1	GTAD(Xu et al., 2020b)	GES(Xiong et al., 2016)	41.50 ²	42.17	encoding	Concat
	2	GTAD(Xu et al., 2020b)	TSP (Alwassel et al., 2020)	51.26	54.34	encoding	RMAttn

¹ From our best run using the official code from (Liu et al., 2021). In the paper, reported mAP is 56.90

² The result is obtained by repeating the evaluation on the set of videos currently available.

Table 2: [THUMOS14] mAP for MUSES(Liu et al., 2021) + I3D(Carreira and Zisserman, 2017) architecture of video-only, audio only and audio-visual fusional approaches.

Fusion Type(scheme)	Visual	Audio	mAP@IoU				
			0.3	0.4	0.5	0.6	0.7
Video-only	I3D	–	67.91	63.15	56.16	46.07	29.89
Audio-only	–	VGGish	8.23	6.61	4.73	3.05	1.30
Concat (Encoding)	I3D	VGGish	70.18	64.98	57.18	45.42	28.86
Dup-Trim (Encoding)	I3D	VGGish	69.25	64.22	56.53	46.08	30.73
Avg-Trim (Encoding)	I3D	VGGish	65.10	60.47	53.92	42.87	28.09
RM-Attention (Encoding)	I3D	VGGish	67.88	63.05	56.19	45.27	29.98
Proposal (Decision)	I3D	VGGish	52.22	47.42	39.37	30.54	17.36

and {0.5, 0.75, 0.95} for ActivityNet-1.3 (Caba Heilbron et al., 2015).

6 RESULTS

The mAP@0.5 results of the best fusion approach for each video-only baseline can be seen in Table 1. It can be seen that incorporating audio consistently improves performance across all approaches. In particular, this incorporation results in a new state-of-the-art result on both the benchmark TAL datasets. In terms of fusion approaches, the clear dominance of the encoding fusion scheme (Section 3.3) can be seen. The Residual Multimodal Attention mechanism from this scheme enables the best performance for the relatively larger ActivityNet-1.3 dataset. Similarly, our mechanism of resampling the audio modality followed by a concatenation of per-modality features enables the best performance for the THUMOS14 dataset. A fuller comparison of the existing best video-only and best audio-visual results obtained via our work can be seen in Tables 2,3. The results once again reinforce the utility of audio for TAL.

6.1 Ablations

To analyze the utility of the fusion schemes proposed in Section 3, we compared their performances with audio-only and video-only methods for the best performing approach in each dataset. Looking at Tables 2 and 3, it is readily evident that fusing audio and video gives the best results. Specifically, RM-Attention fu-

Table 3: [ActivityNet-1.3] mAP for GTAD(Xu et al., 2020b) + TSP(Alwassel et al., 2020) of video-only, audio-only and audio-visual fusional approaches.

Fusion Type (scheme)	Visual	Audio	mAP@IoU			
			0.5	0.75	0.95	avg.
Video-only	TSP	–	51.26	37.12	9.29	35.01
Audio-only	–	VGGish	43.07	28.27	5.82	28.19
Concat (Encoding)	TSP	VGGish	52.6	37.55	9.19	36.37
Dup-Trim (Encoding)	TSP	VGGish	52.31	37.66	9.49	36.47
Avg-Trim (Encoding)	TSP	VGGish	51.91	37.53	9.55	36.33
RM-Attention (Encoding)	TSP	VGGish	54.34	37.66	9.29	36.82
Proposal (Decision)	TSP	VGGish	51.6	37.34	9.41	35.95

sion enables the best result for ActivityNet-1.3 while simple concat works best for Thumos14. The reason for simple concat’s superior performance for Thumos14 can be explained by the fact that audio content is less informative regarding the action boundaries in Thumos14 compared to ActivityNet-1.3. This is also evident from the audio-only baselines – compare the second rows of Tables 2,3. We hypothesize that RM-Attention is more effective at fusing the modalities than filtering out noise when the audio modality is uninformative. In contrast, for simple concat, the separate, non-modulated contribution of audio and visual features makes the fusion scheme less susceptible to noise in any one modality.

DupTrim seems to perform better than AvgTrim, while both are inferior to simple concatenation. This indicates that preserving the ideal frame rate for each modality may not be that crucial to performance and it is probably better to extract features at the same rate for each modality rather than artificially making them equal after extraction. Among the fusion schemes, proposal fusion performs the worst for both ActivityNet-1.3 and Thumos14. This is to be expected because it just selects the best proposal out from the audio and visual streams.

The performances of the audio-only baselines for each dataset suggest that audio information present in ActivityNet-1.3 is much more indicative of activity boundaries compared to that in Thumos14. This is also consistent with the degree of improvement due to fusion for both datasets.

6.2 Class-wise Analysis

To examine the effect of audio at the action class level, we plot the change in Average Precision (AP) relative to the video-only score for the best performing setup. Figure 4 depicts the plot for ActivityNet-1.3, sorted in decreasing order of AP improvement. The majority of action classes show positive AP improvement. This correlates with observations made in the context of aggregate performance (Table 1, Table 3). The action classes which benefit the most from audio (e.g. ‘Playing Ten Pins’, ‘Curling’, ‘Blow-drying hair’) tend to have signature audio transitions marking the beginning and end of the action. The classes at the opposite end (e.g. ‘Painting’, ‘Doing nails’, ‘Putting in contact lenses’) are characterized by very little association between visual and audio modalities. For these classes, we empirically observed that ambient background sounds are present which induce noisy features. However, the gating mechanism enabled by Residual Multimodal Attention ensures that the effect of such noise from the modalities is appropriately mitigated. This can be seen from the smaller magnitude of the drop in AP.

Figure 5 depicts the sorted AP improvement plot for the relatively smaller Thumos14 dataset. Similar qualitative trends as for ActivityNet-1.3 mentioned earlier can be seen, i.e. signature audio transitions characterizing largest AP improvement classes and weak inter-modality associations characterizing least AP improvement classes. However, as mentioned in the previous section, the relatively weak association between audio and video modalities in Thumos14 causes the % of categories which are negatively impacted by audio inclusion to be greater compared to ActivityNet-1.3.

6.3 Instance-wise Analysis

Modifying the approach used by Alwassel et al. (Alwassel et al., 2018) for their approach (DETAD), we analyze two salient attributes of data to analyze the effect of adding audio. These attributes are (i) coverage - the proportion of untrimmed video that the ground truth instance spans (ii) length (temporal duration)

To measure coverage, we normalize the duration of an action instance relative to the duration of the video. Thus, the larger the coverage, the larger the extent the instance occupies in the video. Note that coverage lies between 0 and 1. We group the resulting coverage values into five buckets: Extra Small (XS: (0, 0.2]), Small (S: (0.2, 0.4]), Medium (M: (0.4, 0.6]), Large (L:(0.6, 0.8]), and Extra Large (XL: (0.8, 1.0]). The length is measured as the instance

duration in seconds. We create five different length groups: Extra Small (XS: (0, 30]), Small (S: (30, 60]), Medium (M: (60, 120]), Long (L: (120, 180]), and Extra Long (XL: > 180).

From the numbers below the bucket labels on the x-axis in Figures 6 and 7, we see that most action instances fall in Extra Small buckets. Also, the distributions of coverage and length of the ground truth instances are skewed towards the left (shorter extents).

The change in mAP due to the inclusion of audio can be viewed in Figures 6 and 7 on a per-bucket basis. The overall gain in performance for both datasets is well explained by the overwhelmingly large proportion of the total instances showing improvement due to audio : 63.3% by coverage and 79.81% by length for ActivityNet-1.3 and 95.5% by coverage and length for Thumos14.

From the figures, we see that audio fusion enables consistent improvements for XS and M instances for both datasets while for XL instances, the mAP decreases or remained unchanged. This can be attributed in part to the fact that the shorter instances have an audio ‘signature’ for the action that spans the majority of the instance which assists detection. For the longer action instances, the action-characteristic audio spans a small section of the instance which might not aid detection as much.

6.4 False Positive Analysis

Following (Alwassel et al., 2018), we consider the following error sub-categories within false positive predictions:

- Double Detection error: $IoU > \alpha$ and correct label but not the highest IoU.
- localization error: $0.1 < IoU < \alpha$, correct label
- Confusion error: $0.1 < IoU < \alpha$, wrong label
- Wrong Label error: $IoU > \alpha$, wrong label
- Background error: $IoU < 0.1$

The change in the distribution of possible prediction outcomes with the inclusion of audio can be viewed in Figure 8. The false-positive errors except Background errors have increased. However, their relative frequency is smaller. The large decrease in the number of Background errors more than mitigates the combined increase in other error sub-categories, explaining the overall improvement in performance. The trends in false-positive errors also suggest that audio information is most useful in discriminating between activity instances and the background in untrimmed videos. In addition, we observe that the number of true positive predictions (prediction with




Figure 4: [ActivityNet-1.3] Relative change in per-class AP of the best multimodal setup (Table 1) with inclusion of audio.

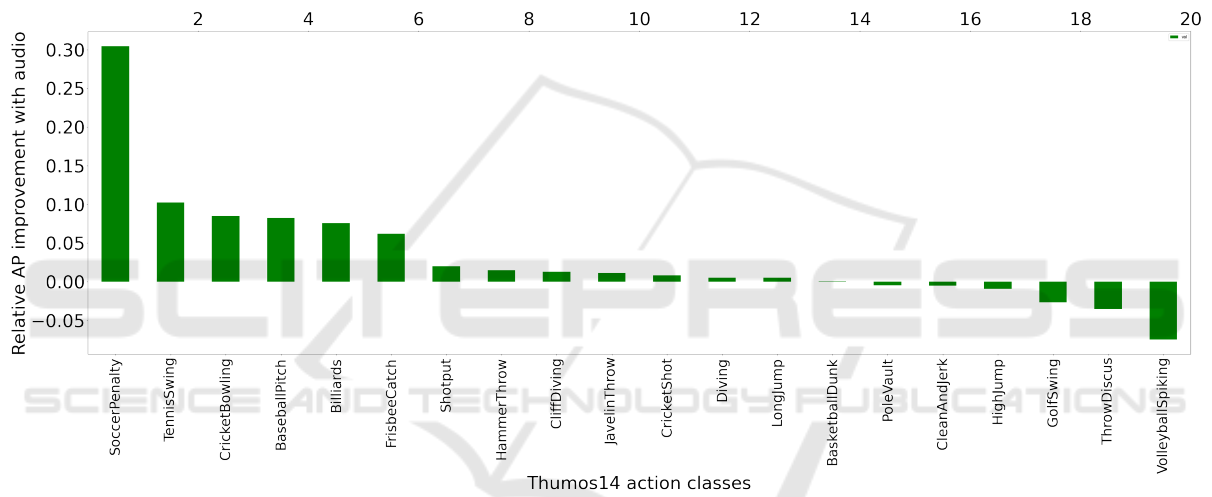


Figure 5: [THUMOS14] Relative change in per-class AP of the best multimodal setup (Table 1) with inclusion of audio.




Figure 6: [ActivityNet-1.3] Relative change in average mAP of the best multimodal setup (Table 1) classified by instance length and coverage, with inclusion of audio. The numbers below X-labels represent the percentage of each type of instance class in the dataset.




Figure 7: [THUMOS14] Relative change in average mAP of the best multimodal setup (Table 1) classified by instance length and coverage, with inclusion of audio. The numbers below X-labels represent the percentage of each type of instance class in the dataset.

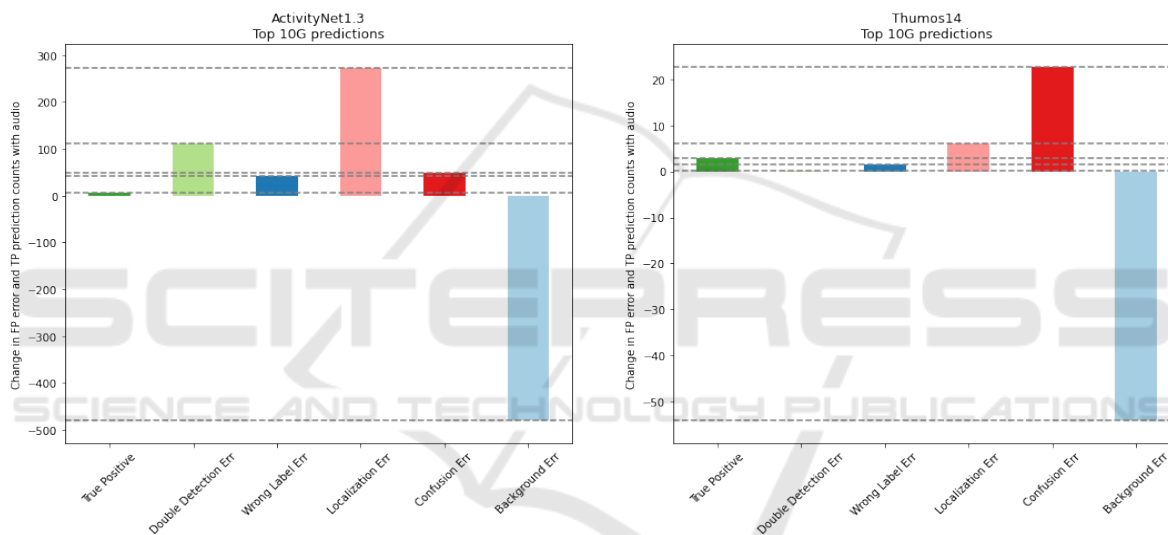


Figure 8: Change in number of True Positive (TP) predictions and False Positive (FP) errors of each type of the best multimodal setup (Table 1) for each dataset with the inclusion of audio. The dashed lines are added to distinguish very close values.

highest IoU , with $IoU > \alpha$ and correctly predicted label, where α is the IoU threshold) increase for both THUMOS14 and ActivityNet-1.3, with the inclusion of audio.

7 CONCLUSION

In this paper, we have presented multiple simple but effective fusion schemes for incorporating audio into existing video-only TAL approaches. To the best of our knowledge, our multimodal effort is the first of its kind for fully supervised TAL. An advantage of our schemes is that they can be readily incorporated into a variety of video-only TAL architectures – a capability we expect to be available for future approaches as

well. Experimental results on two large-scale benchmark datasets demonstrate consistent gains due to our fusion approach over video-only methods and state-of-the-art performance. Our analysis also sheds light on the impact of audio availability on overall as well as per-class performance. Going ahead, we plan to expand and improve the proposed families of fusion schemes.

ACKNOWLEDGEMENTS

We wish to acknowledge the support from MeITY towards this work via the IMSA project (Ref. No.4(16)/2019-ITEA).

REFERENCES

- Alwassel, H., Caba Heilbron, F., Escorcia, V., and Ghanem, B. (2018). Diagnosing error in temporal action detectors. In *ECCV*.
- Alwassel, H., Giancola, S., and Ghanem, B. (2020). TSP: temporally-sensitive pretraining of video encoders for localization tasks. *CoRR*, abs/2011.11479.
- Arandjelovic, R. and Zisserman, A. (2017). Look, listen and learn. In *ICCV*.
- Buch, S., Escorcia, V., Ghanem, B., Fei-Fei, L., and Niebles, J. (2019). End-to-end, single-stream temporal action detection in untrimmed videos. In *BMVC*.
- Caba Heilbron, F., Escorcia, V., Ghanem, B., and Carlos Niebles, J. (2015). Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*.
- Carreira, J. and Zisserman, A. (2017). Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*.
- Delphin-Poulat, L. and Plapous, C. (2019). Mean teacher with data augmentation for dcase 2019 task 4. Technical report, Orange Labs Lannion, France.
- Ebbers, J. and Haeb-Umbach, R. (2020). Convolutional recurrent neural networks for weakly labeled semi-supervised sound event detection in domestic environments. Technical report, DCASE2020 Challenge.
- Feichtenhofer, C., Pinz, A., and Zisserman, A. (2016). Convolutional two-stream network fusion for video action recognition.
- Gao, J., Chen, K., and Nevatia, R. (2018). CTAP: complementary temporal action proposal generation. *CoRR*, abs/1807.04821.
- Gao, J., Yang, Z., Sun, C., Chen, K., and Nevatia, R. (2017). TURN TAP: temporal unit regression network for temporal action proposals. *CoRR*, abs/1703.06189.
- Girshick, R. (2015). Fast r-cnn. In *ICCV*, pages 1440–1448.
- Hao, J., Hou, Z., and Peng, W. (2020). Cross-domain sound event detection: from synthesized audio to real audio. Technical report, DCASE2020 Challenge.
- He, Y., Xu, X., Liu, X., Ou, W., and Lu, H. (2021). Multi-modal transformer networks with latent interaction for audio-visual event localization. In *ICME*.
- Hershey, S., Chaudhuri, S., Ellis, D. P., Gemmeke, J. F., Jansen, A., Moore, R. C., Plakal, M., Platt, D., Saurous, R. A., Seybold, B., et al. (2017). Cnn architectures for large-scale audio classification. In *ICASSP*.
- Jiang, H., Li, Y., Song, S., and Liu, J. (2018). *Rethinking Fusion Baselines for Multi-modal Human Action Recognition: 19th Pacific-Rim Conference on Multimedia, Hefei, China, September 21-22, 2018, Proceedings, Part III*.
- Jiang, Y.-G., Liu, J., Roshan Zamir, A., Toderici, G., Laptev, I., Shah, M., and Sukthankar, R. (2014). THUMOS challenge: Action recognition with a large number of classes.
- Kazakos, E., Nagrani, A., Zisserman, A., and Damen, D. (2019). Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *ICCV*.
- Lee, J.-T., Jain, M., Park, H., and Yun, S. (2021). Cross-attentional audio-visual fusion for weakly-supervised action localization. In *ICLR*.
- Li, X., Lin, T., Liu, X., Gan, C., Zuo, W., Li, C., Long, X., He, D., Li, F., and Wen, S. (2019). Deep concept-wise temporal convolutional networks for action localization.
- Lin, C., Li, J., Wang, Y., Tai, Y., Luo, D., Cui, Z., Wang, C., Li, J., Huang, F., and Ji, R. (2020). Fast learning of temporal action proposal via dense boundary generator. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Lin, C., Xu, C., Luo, D., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., and Fu, Y. (2021). Learning salient boundary feature for anchor-free temporal action localization. In *CVPR*, pages 3320–3329.
- Lin, L. and Wang, X. (2019). Guided learning convolution system for dcase 2019 task 4. Technical report, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China.
- Lin, T., Liu, X., Li, X., Ding, E., and Wen, S. (2019). BMN: boundary-matching network for temporal action proposal generation. *CoRR*, abs/1907.09702.
- Lin, T., Zhao, X., and Shou, Z. (2017). Single shot temporal action detection. *CoRR*, abs/1710.06236.
- Lin, T., Zhao, X., Su, H., Wang, C., and Yang, M. (2018). BSN: boundary sensitive network for temporal action proposal generation. *CoRR*, abs/1806.02964.
- Liu, X., Hu, Y., Bai, S., Ding, F., Bai, X., and Torr, P. H. S. (2021). Multi-shot temporal event localization: A benchmark. In *CVPR*, pages 12596–12606.
- Liu, Y., Ma, L., Zhang, Y., Liu, W., and Chang, S. (2018). Multi-granularity generator for temporal action proposal. *CoRR*, abs/1811.11524.
- Long, X., Gan, C., De Melo, G., Liu, X., Li, Y., Li, F., and Wen, S. (2018a). Multimodal keyless attention fusion for video classification. In *AAAI 2018*.
- Long, X., Gan, C., de Melo, G., Wu, J., Liu, X., and Wen, S. (2018b). Attention clusters: Purely attention based local feature integration for video classification. In *CVPR*.
- Miyazaki, K., Komatsu, T., Hayashi, T., Watanabe, S., Toda, T., and Takeda, K. (2020). Convolution-augmented transformer for semi-supervised sound event detection. Technical report, DCASE2020 Challenge.
- Montes, A., Salvador, A., and Giró-i-Nieto, X. (2016). Temporal activity detection in untrimmed videos with recurrent neural networks. *CoRR*, abs/1608.08128.
- Ono, N., Harada, N., Kawaguchi, Y., Mesaros, A., Imoto, K., Koizumi, Y., , and Komatsu, T. (2020). *Proceedings of the Fifth Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE 2020)*.
- Owens, A. and Efros, A. A. (2018). Audio-visual scene analysis with self-supervised multisensory features. In *ECCV*.
- Owens, A., Wu, J., McDermott, J. H., Freeman, W. T., and Torralba, A. Learning sight from sound: Ambient sound provides supervision for visual learning. *Int. J. Comput. Vis.*

- Shi, Z. (2019). Hodgepodge: Sound event detection based on ensemble of semi-supervised learning methods. Technical report, Fujitsu Research and Development Center, Beijing, China.
- Shou, Z., Chan, J., Zareian, A., Miyazawa, K., and Chang, S. (2017). CDC: convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. *CoRR*, abs/1703.01515.
- Shou, Z., Wang, D., and Chang, S. (2016). Action temporal localization in untrimmed videos via multi-stage cnns. *CoRR*, abs/1601.02129.
- Su, H., Gan, W., Wu, W., Yan, J., and Qiao, Y. (2020). BSN++: complementary boundary regressor with scale-balanced relation modeling for temporal action proposal generation. *CoRR*, abs/2009.07641.
- Tian, Y., Shi, J., Li, B., Duan, Z., and Xu, C. (2018). Audio-visual event localization in unconstrained videos.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., and Van Gool, L. (2019). Temporal segment networks for action recognition in videos. *TPAMI*.
- Wang, Q., Downey, C., Wan, L., Mansfield, P. A., and Moreno, I. L. (2018). Speaker diarization with lstm. In *ICASSP*.
- Wu, Z., Jiang, Y.-G., Wang, X., Ye, H., and Xue, X. (2016). Multi-stream multi-class fusion of deep networks for video classification. In *ACM Multimedia*.
- Xiong, Y., Wang, L., Wang, Z., Zhang, B., Song, H., Li, W., Lin, D., Qiao, Y., Gool, L. V., and Tang, X. (2016). Cuhk & ethz & siat submission to activitynet challenge 2016.
- Xu, H., Das, A., and Saenko, K. (2017). R-C3D: region convolutional 3d network for temporal activity detection. *CoRR*, abs/1703.07814.
- Xu, H., Zeng, R., Wu, Q., Tan, M., and Gan, C. (2020a). Cross-modal relation-aware networks for audio-visual event localization. In *ACM Multimedia*, pages 3893–3901.
- Xu, M., Zhao, C., Rojas, D. S., Thabet, A., and Ghanem, B. (2020b). G-tad: Sub-graph localization for temporal action detection. In *CVPR*.
- Yeung, S., Russakovsky, O., Mori, G., and Fei-Fei, L. (2015). End-to-end learning of action detection from frame glimpses in videos. *CoRR*.
- Zeng, R., Gan, C., Chen, P., Huang, W., Wu, Q., and Tan, M. (2019a). Breaking winner-takes-all: Iterative-winners-out networks for weakly supervised temporal action localization. *IEEE Transactions on Image Processing*.
- Zeng, R., Huang, W., Tan, M., Rong, Y., Zhao, P., Huang, J., and Gan, C. (2019b). Graph convolutional networks for temporal action localization. In *ICCV*.
- Zhang, A., Wang, Q., Zhu, Z., Paisley, J., and Wang, C. (2019). Fully supervised speaker diarization. In *ICASSP*.
- Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Lin, D., and Tang, X. (2017). Temporal action detection with structured segment networks. *CoRR*, abs/1704.06228.
- Zhou, J., Zheng, L., Zhong, Y., Hao, S., and Wang, M. (2021). Positive sample propagation along the audio-visual event line. In *CVPR*.