

Multi-shot Temporal Event Localization: a Benchmark

Xiaolong Liu^{1*} Yao Hu² Song Bai^{2,3†} Fei Ding² Xiang Bai¹ Philip H.S. Torr³

¹Huazhong University of Science and Technology

²Alibaba Group ³University of Oxford

{liuxl, xbai}@hust.edu.cn, songbai.site@gmail.com

{feifei.df, yaohu}@alibaba-inc.com, philip.torr@eng.ox.ac.uk

Abstract

Current developments in temporal event or action localization usually target actions captured by a single camera. However, extensive events or actions in the wild may be captured as a sequence of shots by multiple cameras at different positions. In this paper, we propose a new and challenging task called multi-shot temporal event localization, and accordingly, collect a large-scale dataset called *MULTI-Shot EventS* (MUSES). MUSES has 31,477 event instances for a total of 716 video hours. The core nature of MUSES is the frequent shot cuts, for an average of 19 shots per instance and 176 shots per video, which induces large intra-instance variations. Our comprehensive evaluations show that the state-of-the-art method in temporal action localization only achieves an mAP of 13.1% at IoU=0.5. As a minor contribution, we present a simple baseline approach for handling the intra-instance variations, which reports an mAP of 18.9% on MUSES and 56.9% on THUMOS14 at IoU=0.5. To facilitate research in this direction, we release the dataset and the project code at <https://songbai.site/muses/>.

1. Introduction

Driven by the increasing number of videos generated, shared and consumed every day, video understanding has attracted greater attention in computer vision especially in recent years. As one of the pillars in video understanding, temporal event (or action) localization [10, 47, 61, 77, 84, 85, 87, 88] is a challenging task that aims to predict the semantic label of an action, and in the meantime, locate its start time and end time in a long video. Automating this process is of great importance for many applications, e.g., security surveillance, home care, human-computer interaction, and sports analysis.

*Work done during an internship at Alibaba Group.

†Corresponding author

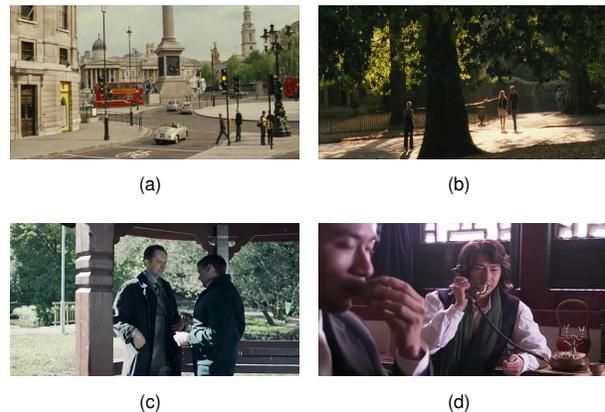


Figure 1. Different types of shot cuts, such as cutting on action (a), dissolve (b), cut-in (c), and cross-cut (d). **Clickable**: best viewed with Adobe Acrobat Reader; click to watch the animation.

While fruitful progress has been made in this field, our work focuses on a real-world scenario that has not been seriously treated so far, that is, localizing events in TV shows and movies. We cast it as a new research topic named multi-shot temporal event localization, which is detailed below:

1) Motivation: Our work is motivated by some commercial applications of multi-shot temporal event localization, where truncated videos are automatically generated from TV shows and movies. For example, a trailer is generated for presenting the highlight to tease the audience, a video summary is generated to help the viewer grasp the important plot points and characters, or a mashup is generated for blending multiple clips of the same theme. To achieve this, a basic but time-consuming step is to find the materials, *i.e.*, the video segments of interest, while the development of multi-shot temporal event localization will enable efficient material extraction and greatly improve the productivity of video content generation.

2) Characteristic: Compared with user-generated videos or surveillance videos, a unique characteristic of TV shows and movies is the highly frequent shot cuts. Herein,



Figure 2. Examples of multi-shot events. In each row, we show three consecutive shots in an instance and select two frames per shot for illustration. The scissor icons indicate the shot boundaries.

a shot means a single sequence of video frames taken by one camera without interruption¹. Because of the use of multi-camera shooting and professional editing techniques, a complete action or event in such videos is usually expressed as a sequence of meaningful short shots connected by cuts of various types, such as cutting on action, dissolve, cut-in, and cross-cut (see Fig. 1 for examples). In other words, the termination of a shot does not mean the end of the corresponding event or action.

3) Challenge: The key challenge of localizing events in TV shows and movies is the large intra-instance variation, induced by the nature of shot cuts. As can be observed in Fig. 2, the view angles and the depth of fields across shots vary dramatically. Meanwhile, due to the existence of shot cuts, some side effects occur, such as scene change, actor change, and heavy occlusions. With such large variations within a single instance, the difficulty of localizing a complete event across shots is significantly increased.

4) Our Contribution: To facilitate the study of multi-shot temporal event localization, the main contribution of this work is a collection of a large-scale dataset named **MULTI-Shot EventS** (MUSES). Unlike existing datasets (*e.g.*, THUMOS14 [30], ActivityNet-1.3 [8], HACS Segments [86]) that are mainly built upon user-generated videos, the data source of MUSES is drama videos processed by professional editing techniques of the entertainment industry. Each instance is defined and annotated as an individual event that may take place across multiple shots. As a consequence, the number of shots in MUSES is as great as 19 per instance and 176 per video. MUSES is also a large-scale dataset that is suitable for training deep learn-

¹Note that the meaning of “shot” in our work is different from those in one-shot/zero-shot/few-shot learning [16].

ing models, which contains 31,477 instances for a total of 716 video hours.

We conduct a comprehensive evaluation of state-of-the-art methods on MUSES, including P-GCN [85], G-TAD [76] and MR [87]. The results show that current methods cannot well handle the large intra-instance variations. The best performance is achieved by P-GCN [85], which is 13.1% mAP at IoU=0.5. It reveals the difficulty of our dataset in temporal event localization, and in the meantime, dramatically exposes the necessity of exploring methods specifically focusing on capturing intra-instance variations. As a minor contribution of this work, we present a simple baseline approach for multi-shot temporal event localization. The proposed baseline achieves an mAP of 18.9% on MUSES and 56.9% on THUMOS14 [30] at IoU=0.5. The dataset and the project code are publicly available for fellow researchers.

2. Related Work

Our work targets temporal event or action localization by contributing a new benchmark dataset. Existing datasets are mainly built upon user-generated videos, where less professional editing is involved. For example, THUMOS14 [30] focuses on sports events. ActivityNet-1.3 [8] extends the taxonomy from sports to human daily activities and significantly increases the number of categories and samples. HACS Segments [86] shares the same lexicon as ActivityNet and further increases the size. In comparison, our dataset is based on drama videos processed by professional editing with frequent shot cuts, so the intra-instance variations are much greater.

Previous methods on temporal event localization can be roughly categorized into two groups, *i.e.*, **two-stage meth-**

Datasets	#Videos	#Categories	#Instances	#Total Hours	#Categories per Video	Multi-shot
THUMOS14 [30]	413	20	6,365	30	1.2	-
ActivityNet-1.3 [8]	19,994	200	30,971	648	1	-
HACS Segments [86]	50k	200	139k	2,048	1	-
MUSES (ours)	3,697	25	31,477	716	3.3	✓

Table 1. Comparing MUSES with existing datasets for temporal event localization.

ods [7, 25, 51, 45, 54, 83, 88] and **one-stage methods** [3, 35, 42, 47, 48, 59, 84], according to whether a standalone proposal generation stage is used. For **two-stage methods**, we first generate a set of proposals, *i.e.*, temporal segments that may contain an event, then leverage a classifier for semantic prediction. Proposal generations can be fulfilled by scoring handcrafted anchors [6, 10, 14, 20, 61, 76], grouping potential event boundaries [43, 59, 87, 88], or a combination of both [19, 46]. For example, Shou *et al.* [61] sample regularly distributed segments as proposals and train a binary classifier to remove background segments. Lin *et al.* [43] locate action boundaries by per-frame classification and group pairs of potential boundaries as proposals. To classify the proposals, Convolutional Neural Networks (CNNs) [20, 61, 88], Recurrent Neural Networks (RNNs) [83] or Graph Convolutional Networks (GCNs) [85] are usually used. For instance, Shou *et al.* [61] employ the C3D [67] network for proposal classification. Zeng *et al.* [85] model proposal-proposal relations with GCNs to extract contextual information. In **one-stage methods**, action localization can be implemented by classifying each pre-defined anchor [5, 35, 42, 47] or each frame [48, 59, 84]. For example, Lin *et al.* [42] predict the classes and boundaries of all anchors simultaneously. Meanwhile, reinforcement learning is also exploited [81]. Some methods [44, 50, 52, 60, 70, 82] also try to learn action localization models in a weakly supervised manner, where only video-level labels are accessible.

This task is also related to action recognition [17, 18, 40, 53, 63, 68, 69, 71, 75, 78], spatio-temporal action localization [27, 31, 39, 64, 73, 79] and temporal action segmentation [13, 35, 36, 37, 38, 55, 72]. Action recognition aims to classify action in a given video, which is usually trimmed to include one action only without background. Representative datasets include KTH [57], Weizemann [22], HMDB51 [34], UCF101 [65], Sports-1M [32], Charades [62], Kinetics [9], Moments in Time [49], Something-Something [23], EventNet [80], Youtube-8M [1] and FineGym [58]. Spatial-temporal action localization, which requires to locate both the temporal segment and spatial locations of actions. UCF101-24 [65], JHMDB [29] and AVA [24] are often used for evaluation. Temporal action segmentation requires to label each frame in a long video by an action class. Some popular datasets include 50Salads [66], GTEA [15] and Breakfast [33], Hollywood Extended [4], and MPII Cooking 2 [56].

MUSES is also related to several movie datasets with action annotations, such as ThreadSafe [26], AVA [24] and MovieNet [28]. However, ThreadSafe and AVA are designed for action recognition and spatio-temporal action localization respectively. MovieNet does not annotate action boundaries. Therefore, none of them is tailored for multi-shot temporal event localization.

3. MUSES Dataset

The goal of this work is to build a large-scale dataset for temporal event localization, especially in the multi-shot scenario. In this section, we present the collection process, the statistic, the characteristic, the evaluation metric, and the baseline approach of the MUSES dataset.

3.1. Data Collection

Category Selection. 25 categories are chosen according to the occurrence frequency in drama, the difficulty of recognizing, and the audience appeal, including *conversation, quarrel, crying, fight, drinking, eating, telephone conversation, horse riding, hugging, stroll, driving, chasing, gunfight, modern meeting, speech, ancient meeting, kissing, war, playing an instrument, dance, (human) flying, cooking, singing, bike riding, and desk work*. Some categories also appear in existing datasets (*e.g.*, ActivityNet-1.3 [8]), most of which are daily-life actions such as *drinking*. In addition, we include some events that are rare in daily life but common in dramas, such as *gunfight*, and *war*.

Data Collection. We begin with the drama database hosted by an online video sharing platform, where each drama is tagged with the genre, such as “romance”, “action”, “crime”, “comedy”, “fantasy”, “war”, “life”, “sports”, and “science fiction”. For each genre, we first exclude those dramas with a small number of views or of a low resolution, then randomly select dramas to ensure diversity. The 500 selected dramas contain 1,003 episodes that cover various years of production, genres, directors and actors, and also convey stories of different periodic eras.

Annotation. The annotation of an instance is given as a set of the start time, the end time, and the category. Our annotation team is composed of 6 experienced annotators. The annotators are trained for a week before the formal annotation. To guarantee the quality, each video is labeled by at least two annotators. Two experts review the annotations provided by the annotators every week and resolve possible

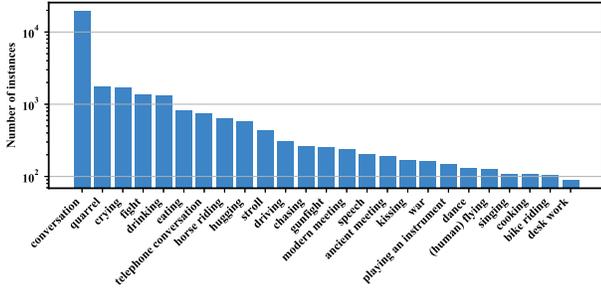


Figure 3. Number of instances per category for MUSES. Note that the Y-axis is in logarithm scale.

inconsistencies between different annotations. The whole annotation process lasts for 3 months.

Two noteworthy comments should be made here. First, as the motivation of this work reveals, an event may go through a series of shots, and sometimes the event is invisible or occluded during the shot cuts. For example, a shot of the audience may be inserted between two shots of a *dance* event. We ask the annotation team to preserve the completeness of the *dance* instance in this case.

Second, different instances may overlap temporally with each other. For example, *telephone conversation* and *driving* can happen simultaneously in a car. In this case, we ask the annotation team to extract two instances belonging to *telephone conversation* and *driving* separately.

3.2. Statistics

The dataset is split into two subsets, with 702 episodes for training and 301 for testing. The average length per episode is 42.8 minutes. Each episode is segmented into videos of around 10 minutes. After removing those that do not contain any events or actions, we obtain 3,697 videos, with 2,587 for training and 1,110 for testing. As a result, MUSES consists of 31,477 instances for a total of 716 video hours.

In Table 1, we give a comparison of dataset statistics between MUSES and some representative datasets, including THUMOS14 [30], ActivityNet-1.3 [8] and HACS Segments [86]. In terms of scale, MUSES is slightly larger than ActivityNet-1.3 but smaller than HACS Segments. The length of videos varies from 300 to 1,151 seconds with an average of 698 seconds, which is about 3 times longer than THUMOS14.

Fig. 3 depicts the category distribution over instances. As it shows, MUSES contains at least 90 instances per category and an average of 1,260 instances. We present the category distribution over videos in Fig. 4(a), which suggests that most videos in MUSES contain instances from 3 categories while those in the other datasets are mainly from 1 semantic class. The instance distribution over videos is given in Fig. 4(b). It can be observed that most videos in

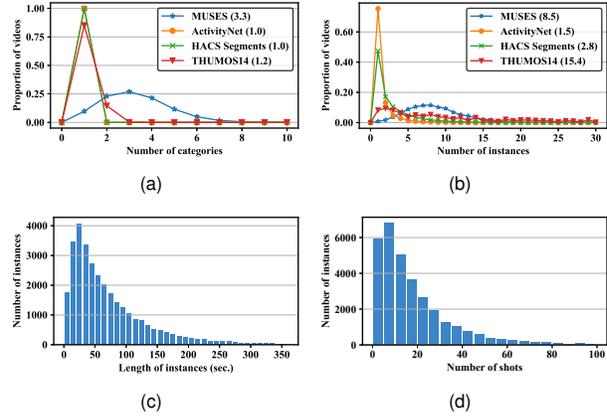


Figure 4. Comparison of datasets: the number of categories (a) and instances per video (b). The numbers in the brackets are the average values of instances and categories per video. The distribution of instance length (c) and the number of shots (d) on MUSES.

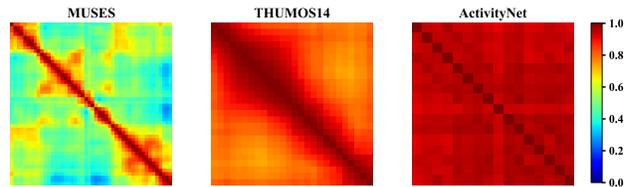


Figure 5. Illustrations of self-similarity samples on MUSES, THUMOS14, and ActivityNet-1.3. The intra-instance variations on MUSES are much greater than that on the other datasets.

MUSES contain 5 to 10 instances with an average of 8.5 instances. Richer yet noisy contextual information is brought by the nature of multi-category and multi-instance in a single video, which encourages the development of advanced techniques for contextual reasoning. Fig. 4(c) presents the distribution of instance length, from which we find that most instances last for 20 to 40 seconds.

Besides, it is found that 15.8% of the instances overlap with another instance with an IoU>0.5. Some frequent pairs include *playing an instrument* ↔ *singing*, *dance* ↔ *singing*, *kissing* ↔ *hugging*, *crying* ↔ *hugging*, *conversation* ↔ *eating*, *war* ↔ *horse riding*, *fight* ↔ *(human) flying* ↔ *driving* and *telephone conversation* ↔ *driving*. The reason behind is that those events or actions usually happen in the same scene.

3.3. Characteristics

The primary property of MUSES is frequent shot cuts inside a single instance. Different from previous datasets [8, 30, 86] for temporal event localization that mainly collect user-generated videos where less editing is involved, MUSES is built upon drama videos made by the industry. In these videos, professional editing techniques are widely used for different purposes, such as removing less informa-

tive segments or guiding the viewer’s attention. As a consequence, an event in such videos is usually expressed by a sequence of short shots connected by various types of shot cuts. Fig. 4(d) presents the distribution of the number of shots over instances. The average number of shots is 19 per instance and 176 per video. Frequent shot cuts induce large intra-instance variations, as we can see from Fig. 2 that different shots differ significantly in the view angles, depth of fields, actors, and backgrounds.

The intra-instance variations can be further evidenced in the feature space. 1) **Self-similarity**: we compute the cosine similarity of each pair of snippets of a single instance using the I3D [9] features, then report the average standard deviation of such self-similarity. The average standard deviation on MUSES is 0.16, while that on THUMOS14 and ActivityNet-1.3 is 0.01 and 0.09, respectively. It suggests that the self-similarity of instances on MUSES is generally smaller than that on the other datasets. Fig. 5 further depicts some examples of self-similarities on different datasets. We find that in MUSES, the snippet pairs with high similarities usually occur in the same shot. 2) **Class activation mapping**: we leverage CAM [89] to visualize the contribution of each temporal snippet to the final prediction. As shown in Fig. 6, the individual contribution is significantly different, which again reveals the intra-instance variations. It also indicates that a considerable number of snippets are less discriminative, caused by the shot cuts, to a precise prediction of the semantic meaning and localization of the corresponding instance.

3.4. Evaluation Metrics

To evaluate the performance of multi-shot temporal event localization, we employ mean Average Precision (mAP). A prediction is considered as a true positive if its IoU with the ground truth is greater than a threshold α and the predicted category is correct, denoted as mAP_α . α is set to 0.3 to 0.7 with a stride of 0.1 following the convention of THUMOS14. Meanwhile, we also report an average of mAPs, which is computed over the 5 IoU thresholds.

Moreover, to quantify the capability of algorithms in handling shot cuts, we further resort to mAPs with different numbers of shots. To be concrete, we divide all the instances into 3 groups according to the number of shots, and report $\text{mAP}_{\text{small}}$ (fewer than 10 shots), $\text{mAP}_{\text{medium}}$ (10 to 20 shots), and $\text{mAP}_{\text{large}}$ (no fewer than 20 shots). In MUSES, the 3 groups make up 39.8%, 27.5%, and 32.7% of all instances, respectively.

3.5. Baseline Approach

Pipeline. We adopt the detection by classifying proposal paradigm following Fast R-CNN [21]. Taking a video and a set of temporal proposals (refer to the supplementary material for proposal generation) as input, our baseline ap-

proach is comprised of three steps: 1) feature extraction. As [85, 10, 87], we use the I3D [9] network to extract a 1D feature for the input video; 2) temporal aggregation. The feature will be forwarded into a temporal aggregation module, which is detailed below, to mitigate the intra-instance variations caused by shot cuts; 3) proposal evaluation. For each proposal, a feature representation is extracted via RoI pooling [21]. Two classifiers are utilized respectively to predict the category and the completeness of the proposal. A boundary regressor is also employed to adjust the boundaries.

Temporal Aggregation. As analyzed in Sec. 3.3, shot cuts will result in large inter-snippet variations. One of the keys is to enhance the feature coherence within a single instance. To this end, we present a simple yet effective module called temporal aggregation to improve the feature discriminative power of each snippet.

Let $\mathbf{X} \in \mathbb{R}^{T \times C}$ be the video feature extracted from the I3D network, where T and C denote its length and the dimension respectively. We first evenly divide \mathbf{X} into H units of length W as $\mathbf{X}' = \text{reshape}(\mathbf{X}) \in \mathbb{R}^{H \times W \times C}$, where $\mathbf{X}'_{ij} = \mathbf{X}_{iW+j}$ and $T = H \times W$. As a result, each row in \mathbf{X}' correspond to consecutive snippets in the same unit, and each column in \mathbf{X}' correspond to non-adjacent snippets with a stride of W (a configurable parameter) in \mathbf{X} .

We then apply the standard 2D convolution to \mathbf{X}' as $\mathbf{Y}' = \mathbf{W} * \mathbf{X}'$, where $\mathbf{W} \in \mathbb{R}^{k_H \times k_W}$ is the convolution kernel. In this way, we can obtain both short-term (within each row) and long-term (across multiple rows) information for each multi-shot instance, so that the feature variation is reduced significantly. It is easy to find that the receptive field is equivalent to $(k_H - 1)W + k_W$ and one can adjust the receptive field by tuning W without changing the kernel size.

The output feature map \mathbf{Y}' is re-casted into 1D view, as $\mathbf{Y} = \text{reshape}(\mathbf{Y}') \in \mathbb{R}^{T \times C}$, which is of the same size as the input feature \mathbf{X} . Therefore, the temporal aggregation module is easy to implement using 2D convolutions and acts as a cheap plug-in. There are a few alternatives for modeling the long-term information, such as dilated 1D convolution, deformable 1D convolution [11], *etc.* Our experiments suggest that they achieve inferior performance to our temporal aggregation module.

Moreover, as the duration of events and shots in MUSES vary significantly, we adopt a split-transform-merge strategy [74] and build a multi-branch block to handle the scale changes. More precisely, we use K temporal aggregation modules, each of which uses a convolution kernel of different sizes. The output of each branch is fused by summation. Thanks to the various kernel sizes, the learned feature integrates temporal information of different scales and is potentially capable of better dealing with events of different durations.

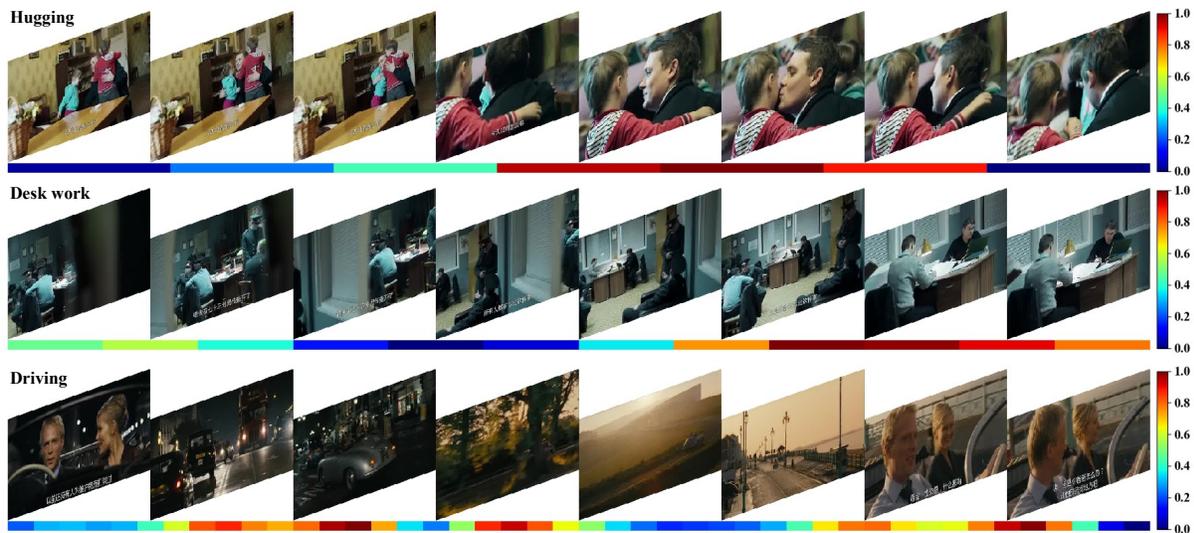


Figure 6. Class activation mapping (CAM) of some ground truth instances in MUSES.

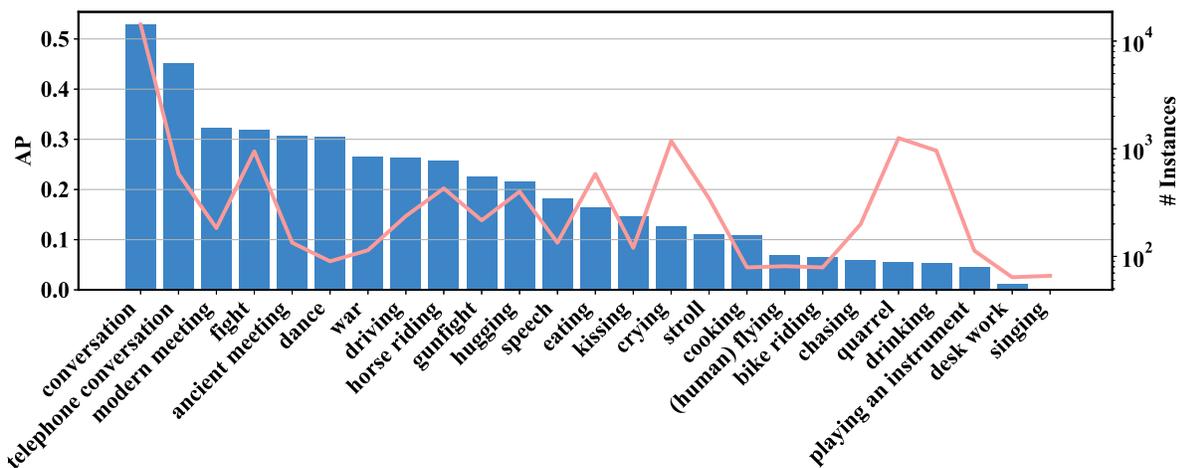


Figure 7. Event localization performance in terms of AP (bars) and the number of training instances per category (lines) on MUSES.

4. Experiments

In this section, we give a thorough study of the newly collected MUSES dataset. In addition, we report the results on THUMOS14 [30] and ActivityNet-1.3 [8] under the conventional experiment setup following [61, 88, 43, 41]. Due to space limitation, we put the experimental details, such as the network architecture, the loss functions, the training strategy, some tricks like boundary extension [12] and the hyper-parameters, in the supplementary material.

4.1. Analysis of MUSES Dataset

Which Categories Are More Challenging? We present the localization performance and the number of training instances per category in Fig. 7. The best performance is witnessed in the *conversation* category, probably owing to the

largest number of training instances. Meanwhile, categories related to objects and scenes, such as *telephone conversation* and *modern meeting*, and categories that have salient motion features, such as *fight* and *dance*, are relatively easier to be coped with. The top-5 challenging categories are *singing*, *desk work*, *playing an instrument*, *drinking*, and *quarrel*. The poor performance might be caused by less training data or larger diversity.

Error Analysis. Using the protocol devised by [2], we present an error analysis on MUSES in Fig. 8(a). According to its IoU with the matched ground truth (denoted as gIoU) and whether the predicted label is correct, a false positive can be divided into five types, including Double Detection Error (a duplicate prediction with $\text{gIoU} \geq \alpha$ and a correct label that is not the highest scoring one), Wrong Label Error ($\text{gIoU} \geq \alpha$, incorrect label), Localization Er-

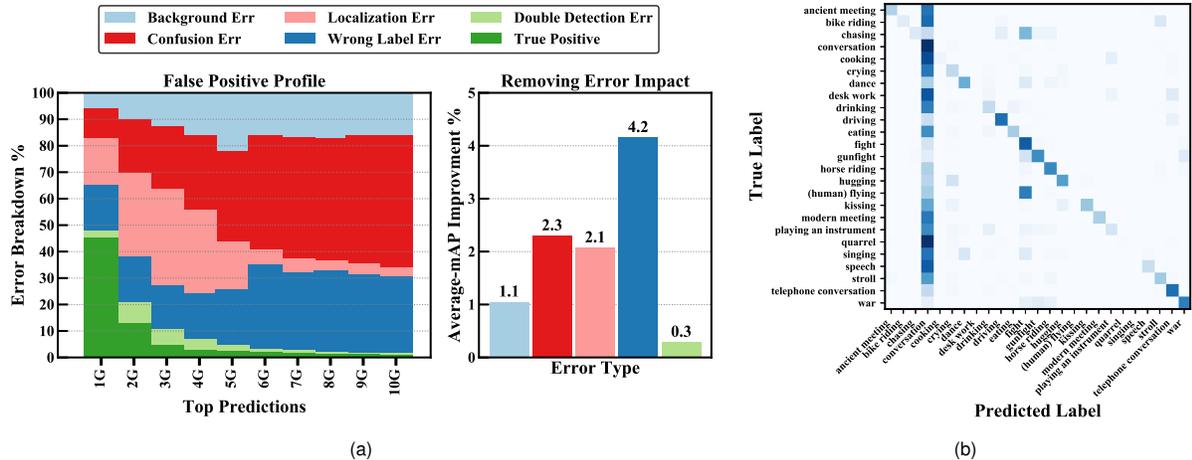


Figure 8. Error analysis on MUSES. (a) **Left**: the error distribution over the number of predictions per video. G means the number of ground truth instances. (a) **Right**: the impact of error types, measured by the improvement gained from resolving a particular type of error. See Sec. 4.1 for the definitions of the errors. (b) Confusion matrix.

ror ($0.1 \leq \text{gIoU} < \alpha$, correct label), Confusion Error ($0.1 \leq \text{gIoU} < \alpha$, incorrect label), and Background Error ($\text{gIoU} < 0.1$).

As Fig. 8(a) left shows, Confusion Error, Localization Error, and Wrong Label Error are the three main sources of false positives. It indicates that incorrect boundaries and incorrect labels are responsible for false positive predictions, which are the side effect of frequent shot cuts. Different from the error distribution on ActivityNet-1.3 where Localization Error occurs most (quoted from [2]), more Wrong Label Errors and Confusion Errors are observed on MUSES. This phenomenon suggests that the classification task is more difficult on MUSES compared with that on ActivityNet-1.3. Meanwhile, Fig. 8(a) right suggests that the performance boost is more remarkable if the classification task is well addressed.

To dissect Wrong Label Error, we present the confusion matrix in Fig. 8(b). It is observed that instances in many categories are easy to be mis-classified into *conversation*. The possible reasons are two folds: 1) imbalanced number of training samples among different classes; 2) as one of the key visual cues in *conversation*, the close-ups of characters also exist in the other categories, such as *modern meeting*, *quarrel*, *speech*, *desk work*, etc. Besides, event pairs with high concurrency, such as *fight* \leftrightarrow (*human*) *flying* and *singing* \leftrightarrow *dance*, are easy to confuse.

Effect of Temporal Aggregation. In Table 2, we compare the performance of different variants of our method. As can be observed, removing temporal aggregation (vanilla) leads to a decrease of 2.8% in terms of $\text{mAP}_{0.5}$ on MUSES. Replacing temporal aggregation with dilated 1D convolution or deformable 1D convolution also leads to a performance drop. Meanwhile, an integration of multiple tem-

Methods	$k_H \times k_W$	W	$\text{mAP}_{0.5}$
Vanilla	-	-	16.1
Dilated Conv.	-	-	17.1
Deformable Conv. [11]	-	-	17.0
Ours (Single-scale)	1×3	3	16.0
	3×3	3	16.8
	3×3	6	17.5
	3×3	9	17.2
Ours (Multi-scale)	-	-	18.9

Table 2. Performance analysis of temporal aggregation in terms of $\text{mAP}_{0.5}$.

poral aggregation modules at different scales outperforms its single-scale counterparts. The performance gain stems from the fact that temporal aggregation is capable of enhancing the feature coherence within each single instance to a certain extent. To verify this, we follow the gist described in Sec. 3.3 to compute the standard deviation of self-similarities, and find that the standard deviation is decreased from 0.16 to 0.09 after temporal aggregation is applied.

State-of-the-art Evaluations on MUSES. We present an evaluation of state-of-the-art methods whose code is publicly available, including P-GCN [85], G-TAD [77], and MR [87]. Note that G-TAD and MR only focus on class-agnostic proposals, so we train an MLP on I3D features to generate category predictions. To ensure a fair comparison, the video features are kept the same for all the methods and the same proposals are used in P-GCN and our approach. As shown in Table 3, our method achieves the highest mAPs among all the methods. However, the achievement is far below the state-of-the-art performance on THUMOS14 (an mAP of around 52% at $\text{IoU}=0.5$, refer to [77, 85]) and

Methods	mAP _{0.3}	mAP _{0.4}	mAP _{0.5}	mAP _{0.6}	mAP _{0.7}	mAP	mAP _{small}	mAP _{medium}	mAP _{large}
Random	1.20	0.64	0.29	0.10	0.03	0.45	0.02	0.06	0.41
MR [87]	12.9	11.3	9.2	7.6	5.9	9.4	4.1	6.9	7.5
G-TAD [77]	19.1	14.8	11.1	7.4	4.7	11.4	7.3	9.6	4.6
P-GCN [85]	19.9	17.1	13.1	9.7	5.4	13.0	6.3	8.7	7.9
Ours	25.9	22.6	18.9	15.0	10.6	18.6	7.4	11.9	17.3

Table 3. Performance evaluation of state-of-the-art methods on the newly collected MUSES dataset.

Methods	mAP _{0.3}	mAP _{0.4}	mAP _{0.5}	mAP _{0.6}	mAP _{0.7}
SCNN [61]	36.3	28.7	19.0	10.3	5.3
SMS [84]	36.5	27.8	17.8	-	-
CDC [59]	41.3	30.7	24.7	14.3	8.8
Dai <i>et al.</i> [12]	-	33.3	25.6	15.9	9.0
R-C3D [76]	44.8	35.6	28.9	19.1	9.3
SS-TAD [5]	45.7	-	29.2	-	9.6
SSN [88]	50.6	40.8	29.1	-	-
TAL-Net [10]	53.2	48.5	42.8	33.8	20.8
GTAN [47]	57.8	47.2	38.8	-	-
P-GCN [85]	63.6	57.8	49.1	-	-
G-TAD [77]	66.4	60.4	51.6	37.6	22.9
MR [87]	53.9	50.7	45.4	38.0	28.5
Ours	68.9	64.0	56.9	46.3	31.0

Table 4. Performance comparison on THUMOS14 in terms of mAP at IoU thresholds from 0.3 to 0.7.

ActivityNet-1.3 (an mAP of around 50% at IoU=0.5, refer to [41, 77]), which reveals the difficulty of our dataset. Note that mAP_{large} can be higher than mAP_{medium} and mAP_{small}, as events with more shots are often longer thus easier to localize.

4.2. Evaluations on THUMOS14 and ActivityNet

Our baseline approach leverages temporal aggregation which is originally designed for handling intra-instance variations caused by frequent shot cuts. Here, we report the performance of temporal aggregation on THUMOS14 and ActivityNet-1.3. Table 4 presents the comparison with the state-of-the-art methods on THUMOS14. Our approach achieves the highest mAPs at different IoU thresholds. In particular, we outperform the second-best entry, that is G-TAD [77], by an absolute improvement of 5.3% in terms of mAP at IoU=0.5. Table 5 presents the performance comparison on ActivityNet-1.3. We build our model upon BMN [41], one of the state-of-the-art algorithms with publicly available code. As it shows, integrating temporal aggregation with BMN is also useful and we report an average mAP of 33.99%.

5. Conclusion

Truncating TV shows and movies into concise and attractive short videos has been a popular way of increasing click-through rates in video sharing platforms, where localizing temporal segments of interest is the kick-off step.

Methods	mAP _{0.5}	mAP _{0.75}	mAP _{0.95}	mAP
R-C3D [76]	26.80	-	-	12.70
TAL-Net [10]	38.23	18.30	1.30	20.22
GTAN [47]	52.61	34.14	8.91	34.31
BMN [41]	50.07	34.78	8.29	33.85
BMN* [41]	49.66	33.85	7.86	33.45
Ours	50.02	34.97	6.57	33.99

Table 5. Performance comparison on the validation set of ActivityNet-1.3 at different IoU thresholds. * indicates the result is our re-implementation with the publicly available code.

However, temporal event or action localization in such video sources is more or less ignored by our research community. And existing methods cannot well handle the intra-instance variations caused by the frequent shot cuts as TV shows and movies are usually post-processed by professional editing techniques.

To enable automatic video content generation with efficient and scalable material extraction in TV shows and movies, we define a new task called multi-shot temporal event localization that aims to localize events or actions captured in multiple shots. A large-scale dataset, called MUSES, is collected for this study. Built upon drama videos, MUSES provide rich multi-shot instances with frequent shot cuts, which induces great intra-instance variations and brings new challenges to current approaches. A comprehensive evaluation shows that the state-of-the-art methods in this field fail to cope with frequent shot cuts and reveal the difficulty of the MUSES dataset. MUSES could serve as a benchmark dataset and facilitate research in temporal event localization.

MUSES will be updated in the future by adding multi-modal annotations, such as subtitles, audio, facial expression, *etc.* The setting of weakly-supervised multi-shot temporal event localization will be also defined and explored on top of MUSES. We hope that MUSES could trigger and advance the application of temporal event localization techniques to commercial products in video content generation.

Acknowledgments. This work was supported by National Key R&D Program of China (No. 2018YFB1004600), to Dr. Xiang Bai by the National Program for Support of Top-notch Young Professionals, Alibaba Group through Alibaba Research Intern Program, Royal Academy of Engineering under the Research Chair and Senior Research Fellowships scheme, EPSRC/MURI grant EP/N019474/1 and FiveAI.

References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. 3
- [2] Humam Alwassel, Fabian Caba Heilbron, Victor Escorcia, and Bernard Ghanem. Diagnosing error in temporal action detectors. In *ECCV*, pages 256–272, 2018. 6, 7
- [3] Humam Alwassel, Fabian Caba Heilbron, and Bernard Ghanem. Action search: Spotting actions in videos and its application to temporal action localization. In *ECCV*, pages 251–266, 2018. 3
- [4] Piotr Bojanowski, Rémi Lajugie, Francis Bach, Ivan Laptev, Jean Ponce, Cordelia Schmid, and Josef Sivic. Weakly supervised action labeling in videos under ordering constraints. In *ECCV*, pages 628–643, 2014. 3
- [5] S Buch, V Escorcia, B Ghanem, L Fei-Fei, and JC Niebles. End-to-end, single-stream temporal action detection in untrimmed videos. In *BMVC*, 2017. 3, 8
- [6] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Niebles. Sst: Single-stream temporal action proposals. In *CVPR*, pages 6373–6382, 2017. 3
- [7] Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *CVPR*, pages 1914–1923, 2016. 3
- [8] 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, 3, 4, 6
- [9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 4724–4733, 2017. 3, 5
- [10] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Re-thinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018. 1, 3, 5, 8
- [11] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, pages 764–773, 2017. 5, 7
- [12] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S Davis, and Yan Qiu Chen. Temporal context network for activity localization in videos. In *ICCV*, pages 5727–5736, 2017. 6, 8
- [13] Li Ding and Chenliang Xu. Weakly-supervised action segmentation with iterative soft boundary assignment. In *CVPR*, pages 6508–6516, 2018. 3
- [14] Victor Escorcia, Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Daps: Deep action proposals for action understanding. In *ECCV*, pages 768–784, 2016. 3
- [15] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, pages 3281–3288, 2011. 3
- [16] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *TPAMI*, 28(4):594–611, 2006. 2
- [17] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019. 3
- [18] Christoph Feichtenhofer, Axel Pinz, and Richard P. Wildes. Spatiotemporal residual networks for video action recognition. In *NIPS*, pages 3468–3476, 2016. 3
- [19] Jiyang Gao, Kan Chen, and Ram Nevatia. Ctap: Complementary temporal action proposal generation. In *ECCV*, September 2018. 3
- [20] Jiyang Gao, Zhenheng Yang, Chen Sun, Kan Chen, and Ram Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. In *ICCV*, pages 3648–3656, 2017. 3
- [21] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. 5
- [22] Lena Gorelick, Moshe Blank, Eli Shechtman, Michal Irani, and Ronen Basri. Actions as space-time shapes. *TPAMI*, 29(12):2247–2253, 2007. 3
- [23] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, number 4, page 5, 2017. 3
- [24] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, pages 6047–6056, 2018. 3
- [25] Fabian Caba Heilbron, Wayner Barrios, Victor Escorcia, and Bernard Ghanem. Scc: Semantic context cascade for efficient action detection. In *CVPR*, pages 3175–3184, 2017. 3
- [26] Minh Hoai and Andrew Zisserman. Thread-safe: Towards recognizing human actions across shot boundaries. In *ACCV*, pages 222–237. Springer, 2014. 3
- [27] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (t-cnn) for action detection in videos. In *ICCV*, pages 5822–5831, 2017. 3
- [28] Qingqiu Huang, Yu Xiong, Anyi Rao, Jiase Wang, and Dahua Lin. Movienet: A holistic dataset for movie understanding. In *ECCV*, 2020. 3
- [29] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *ICCV*, pages 3192–3199, 2013. 3
- [30] YG Jiang, Jingen Liu, A Roshan Zamir, G Toderici, I Laptev, Mubarak Shah, and Rahul Sukthankar. Thumos challenge: Action recognition with a large number of classes, 2014. 2, 3, 4, 6
- [31] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *ICCV*, pages 4405–4413, 2017. 3
- [32] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, pages 1725–1732, 2014. 3

- [33] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, pages 780–787, 2014. [3](#)
- [34] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso A. Poggio, and Thomas Serre. HMDB: A large video database for human motion recognition. In Dimitris N. Metaxas, Long Quan, Alberto Sanfeliu, and Luc Van Gool, editors, *ICCV*, pages 2556–2563, 2011. [3](#)
- [35] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *CVPR*, pages 156–165, 2017. [3](#)
- [36] Colin Lea, Austin Reiter, René Vidal, and Gregory D Hager. Segmental spatiotemporal cnns for fine-grained action segmentation. In *ECCV*, pages 36–52, 2016. [3](#)
- [37] Peng Lei and Sinisa Todorovic. Temporal deformable residual networks for action segmentation in videos. In *CVPR*, pages 6742–6751, 2018. [3](#)
- [38] Shi-Jie Li, Yazan AbuFarha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. Ms-tcn++: Multi-stage temporal convolutional network for action segmentation. *TPAMI*, 2020. [3](#)
- [39] Yixuan Li, Zixu Wang, Limin Wang, and Gangshan Wu. Actions as moving points. *ECCV*, 2020. [3](#)
- [40] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. [3](#)
- [41] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, pages 3889–3898, 2019. [6](#), [8](#)
- [42] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *ACM MM*, pages 988–996, 2017. [3](#)
- [43] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, September 2018. [3](#), [6](#)
- [44] Daochang Liu, Tingting Jiang, and Yizhou Wang. Completeness modeling and context separation for weakly supervised temporal action localization. In *CVPR*, pages 1298–1307, 2019. [3](#)
- [45] Xiaolong Liu, Yuchao Sun, Jianghu Lu, Cong Yao, and Yu Zhou. Self-similarity action proposal. *IEEE Signal Processing Letters*, 27:2064–2068, 2020. [3](#)
- [46] Yuan Liu, Lin Ma, Yifeng Zhang, Wei Liu, and Shih-Fu Chang. Multi-granularity generator for temporal action proposal. In *CVPR*, pages 3604–3613, 2019. [3](#)
- [47] 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. [1](#), [3](#), [8](#)
- [48] Shugao Ma, Leonid Sigal, and Stan Sclaroff. Learning activity progression in lstms for activity detection and early detection. In *CVPR*, June 2016. [3](#)
- [49] Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. *TPAMI*, 42(2):502–508, 2019. [3](#)
- [50] Phuc Nguyen, Ting Liu, Gautam Prasad, and Bohyung Han. Weakly supervised action localization by sparse temporal pooling network. In *CVPR*, pages 6752–6761, 2018. [3](#)
- [51] Bingbing Ni, Xiaokang Yang, and Shenghua Gao. Progressively parsing interactional objects for fine grained action detection. In *CVPR*, pages 1020–1028, 2016. [3](#)
- [52] Sujoy Paul, Sourya Roy, and Amit K. Roy-Chowdhury. Wtalc: Weakly-supervised temporal activity localization and classification. In *ECCV*, September 2018. [3](#)
- [53] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, pages 5533–5541, 2017. [3](#)
- [54] Alexander Richard and Juergen Gall. Temporal action detection using a statistical language model. In *CVPR*, pages 3131–3140, 2016. [3](#)
- [55] Alexander Richard, Hilde Kuehne, and Juergen Gall. Weakly supervised action learning with rnn based fine-to-coarse modeling. In *CVPR*, pages 754–763, 2017. [3](#)
- [56] Marcus Rohrbach, Anna Rohrbach, Michaela Regneri, Sikandar Amin, Mykhaylo Andriluka, Manfred Pinkal, and Bernt Schiele. Recognizing fine-grained and composite activities using hand-centric features and script data. *IJCV*, 119(3):346–373, 2016. [3](#)
- [57] Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: a local svm approach. In *ICPR*, volume 3, pages 32–36, 2004. [3](#)
- [58] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Finegym: A hierarchical video dataset for fine-grained action understanding. In *CVPR*, pages 2616–2625, 2020. [3](#)
- [59] 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 *ICCV*, pages 1417–1426, 2017. [3](#), [8](#)
- [60] Zheng Shou, Hang Gao, Lei Zhang, Kazuyuki Miyazawa, and Shih-Fu Chang. Autoloc: Weakly-supervised temporal action localization in untrimmed videos. In *ECCV*, pages 154–171, 2018. [3](#)
- [61] 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](#), [3](#), [6](#), [8](#)
- [62] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, pages 510–526, 2016. [3](#)
- [63] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, pages 568–576, 2014. [3](#)
- [64] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip HS Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *ICCV*, pages 3637–3646, 2017. [3](#)
- [65] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012. [3](#)
- [66] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *ACM International Joint*

- Conference on Pervasive and Ubiquitous Computing*, pages 729–738, 2013. 3
- [67] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015. 3
- [68] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018. 3
- [69] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, pages 3551–3558, 2013. 3
- [70] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *CVPR*, pages 4325–4334, 2017. 3
- [71] 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. 3
- [72] Zhenzhi Wang, Ziteng Gao, Limin Wang, Zhifeng Li, and Gangshan Wu. Boundary-aware cascade networks for temporal action segmentation. In *ECCV*, 2020. 3
- [73] Jianchao Wu, Zhanghui Kuang, Limin Wang, Wayne Zhang, and Gangshan Wu. Context-aware rcnn: A baseline for action detection in videos. *ECCV*, 2020. 3
- [74] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, pages 1492–1500, 2017. 5
- [75] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, pages 305–321, 2018. 3
- [76] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: region convolutional 3d network for temporal activity detection. In *ICCV*, pages 5794–5803, 2017. 2, 3, 8
- [77] 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, 7, 8
- [78] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *CVPR*, pages 591–600, 2020. 3
- [79] Xitong Yang, Xiaodong Yang, Ming-Yu Liu, Fanyi Xiao, Larry S Davis, and Jan Kautz. Step: Spatio-temporal progressive learning for video action detection. In *CVPR*, pages 264–272, 2019. 3
- [80] Guangnan Ye, Yitong Li, Hongliang Xu, Dong Liu, and Shih-Fu Chang. Eventnet: A large scale structured concept library for complex event detection in video. In *ACM MM*, pages 471–480, 2015. 3
- [81] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, pages 2678–2687, 2016. 3
- [82] Tan Yu, Zhou Ren, Yuncheng Li, Enxu Yan, Ning Xu, and Junsong Yuan. Temporal structure mining for weakly supervised action detection. In *ICCV*, pages 5522–5531, 2019. 3
- [83] 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. 3
- [84] Ze-Huan Yuan, Jonathan C Stroud, Tong Lu, and Jia Deng. Temporal action localization by structured maximal sums. In *CVPR*, volume 2, page 7, 2017. 1, 3, 8
- [85] Runhao Zeng, Wenbing Huang, Mingkui 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, 3, 5, 7, 8
- [86] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. HACS: human action clips and segments dataset for recognition and temporal localization. In *ICCV*, pages 8667–8677. IEEE, 2019. 2, 3, 4
- [87] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *ECCV*, 2020. 1, 2, 3, 5, 7, 8
- [88] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. *ICCV*, 2, 2017. 1, 3, 6, 8
- [89] B. Zhou, A. Khosla, Lapedriza. A., A. Oliva, and A. Torralba. Learning Deep Features for Discriminative Localization. *CVPR*, 2016. 5