

Learning Latent Super-Events to Detect Multiple Activities in Videos

AJ Piergiovanni and Michael S. Ryoo

Department of Computer Science, Indiana University, Bloomington, IN 47408

{ajpiergi, mryoo}@indiana.edu

Abstract

In this paper, we introduce the concept of learning latent super-events from activity videos, and present how it benefits activity detection in continuous videos. We define a super-event as a set of multiple events occurring together in videos with a particular temporal organization; it is the opposite concept of sub-events. Real-world videos contain multiple activities and are rarely segmented (e.g., surveillance videos), and learning latent super-events allows the model to capture how the events are temporally related in videos. We design temporal structure filters that enable the model to focus on particular sub-intervals of the videos, and use them together with a soft attention mechanism to learn representations of latent super-events. Super-event representations are combined with per-frame or per-segment CNNs to provide frame-level annotations. Our approach is designed to be fully differentiable, enabling end-to-end learning of latent super-event representations jointly with the activity detector using them. Our experiments with multiple public video datasets confirm that the proposed concept of latent super-event learning significantly benefits activity detection, advancing the state-of-the-arts.

1. Introduction

Activity detection is an important computer vision problem with many societal applications, including smart surveillance, monitoring of patients or elderly (e.g., for quality-of-life systems), online video retrieval, and robot perception. Given a continuous video, the task is to find the frames corresponding to every event occurring in the video. This is more challenging compared to the activity classification problem of categorizing a pre-segmented video or localizing a single activity in a trimmed video. Although activity detection is an important area to study as almost all real-world videos contain multiple activities and are rarely segmented (e.g., surveillance systems), it has been investigated much less, particularly for multi-event videos.

In the past years, end-to-end learning methods using convolutional neural networks (CNNs) obtained a great amount




Figure 1. **Top:** In a basketball video, the block event cannot occur without a shot event. Similarly, the dribble event provides temporal context as to where a shot event can occur. **Bottom:** In a video of a person eating, ‘eating a sandwich’ must occur near ‘holding sandwich’. The drinking action is also related to eating.

of success in video analysis. These approaches successfully modeled per-frame (or per-local-segment) information in activity videos, such as a single RGB frame or optical flows over a small number of frames [7]. Recently, models such as I3D [3] have been developed to capture longer-term dynamics (e.g., 64 frames). However, because such end-to-end models are optimized for capturing per-segment information, the primary focus of existing works has been mainly on activity classification and not detection. There are recent works on activity detection using end-to-end models (e.g., the detection task in [25]), but many of these works also focus on making better per-segment decisions (and their post-processing) rather than learning details of temporal structure/context over the entire (variable length) video.

A video contains multiple activities (in a sequence or in parallel) and they are correlated. This means that detecting the frames of one activity in the video should benefit from information in the frames corresponding to another activity, which are often temporally very separated. Existing approaches of representing/classifying video segments without regard to contextual information is thus limited; continuous videos contain rich temporal structure which can be exploited to improve activity detection. For example, in a

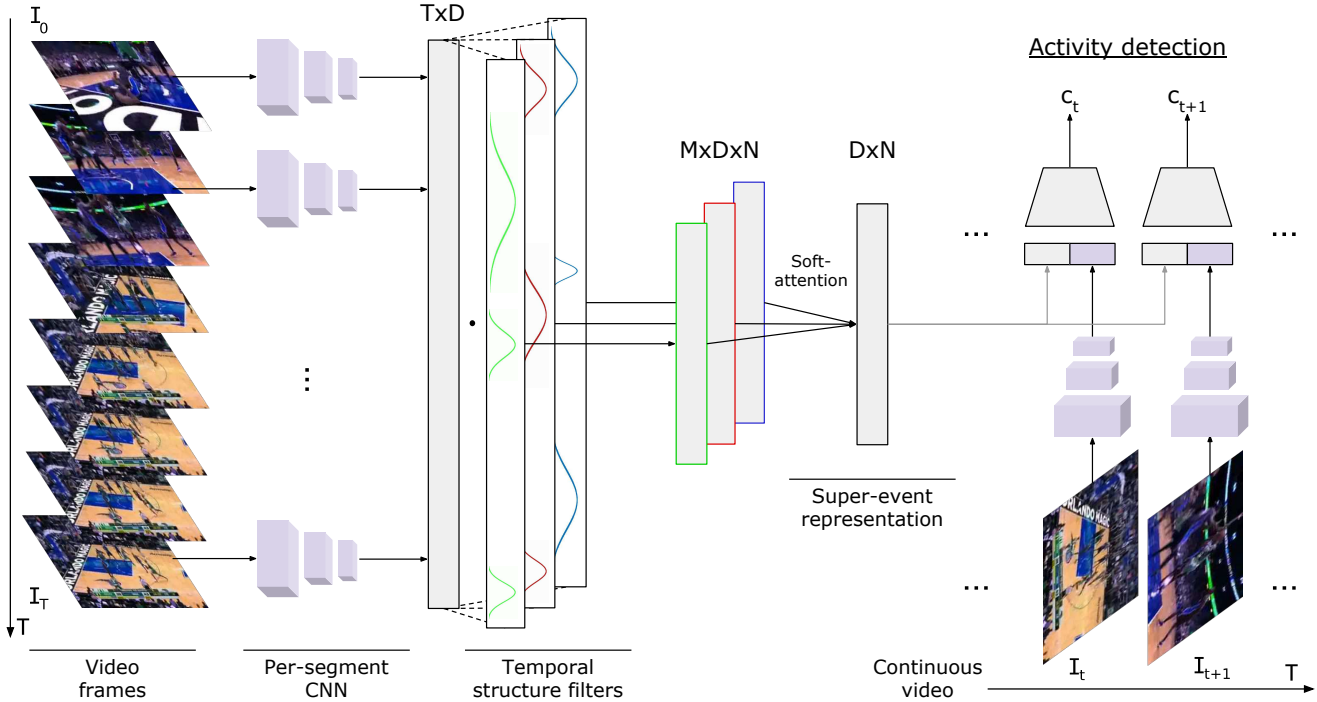


Figure 2. Overview of our approach. The temporal structure filters are applied to the entire video and soft-attention is applied to form a *super-event* representation for each video. We concatenate per-frame (or per-segment) video features with the super-event representation to make per-frame classifications, annotating each frame with its activity class (or no-activity).

video of a basketball game, shooting and blocking events must occur near-by, as shown in Figure 1. A block or rebound event cannot occur without a shot event.

In this paper, we introduce the concept of *super-events* and present how its representation learning can benefit activity detection. We define a super-event as a set of multiple events occurring together in videos with a particular temporal pattern (i.e., structure). More specifically, if the event we are interested in is a part of a longer-term event, we call the longer-term event its super-event. This is the opposite concept of ‘sub-events’. For example, the events of shooting and blocking mentioned above forms a super-event, which may be named as a blocked-shoot. Learning such latent super-events allows the model to capture how the events are temporally related in their videos. Once learned, when making a prediction (i.e., testing), the super-events can serve as temporal context to better detect the events. This enables the detection decision at each frame to be made while considering longer-term temporal structure. Note that such super-events are ‘latent’, meaning that no super-event annotations are provided.

We newly design *temporal structure filters*, and convolve it with the video representation to obtain a super-event representation. Temporal structure filters allow the model to focus on particular sub-intervals. Such temporal structure filters are learned for each event, optimized based on the

training data for the best super-event representation construction. For each frame, we combine the super-event representation with the per-frame or per-segment CNN representation for its binary classification per event. Our method is fully-differentiable and can be learned end-to-end using back propagation, making it suitable for any length video. Our experimental results with three different datasets confirm the benefits of our latent super-event learning, obtaining the state-of-the-art results on MultiTHUMOS and Charades.

2. Related works

Activity recognition has been a popular research topic in computer vision [1]. Hand-crafted features, such as improved dense trajectories [29] gave excellent results on many benchmark datasets in the past. Recently, there have been more works on learning features for action recognition using convolutional models. Two-stream CNN approaches take both RGB frames and optical flow as input [26] to capture motion and image features. 3D spatio-temporal (XYT) convolutional filters also have been learned and applied to many activity recognition tasks [28, 3]. Large scale datasets and challenges such as THUMOS [12], ActivityNet [6] and Charades [25] provided these approaches training videos to learn the model.

Improving activity recognition by temporally aggregat-

ing such per-frame CNN features has also been studied. Ng et al. [16] compared several different forms of temporal pooling over per-frame CNN features. Karpathy et al. [13] compared various methods to combine temporal information within a CNN. Piergiovanni et al. [18] studied learning multiple sub-intervals to improve activity recognition. These works focused mostly on pooling short-term temporal information to classify a single interval as an activity. They did not detect multiple activity instances or explore long-term structure between multiple activities.

Recurrent neural networks (RNNs) such as long short-term memory (LSTM) have been popularly used to model temporal event transitions between frames [32, 5, 31]. These approaches all relied on a single RNN to capture temporal information in different frames, only implicitly capturing relationships between multiple activities. Our approach allows for more explicit modeling of activity relationships as their super-events, which leads to better performance and more insights into what the network learns.

Recently, segment-based 3D CNNs have been used to capture spatio-temporal information simultaneously for the activity detection [22, 30]. Shou et al. [21] uses convolutional upsampling to make dense predictions to better localize start and end times. Zhao et al. [33] use a binary actionness classifier to propose many action segments per video then classify each segment individually. However, these approaches all treated the segments as individual instances and do not exploit the longer-term relationships between actions.

Sigurdsson et al. [24] found that using intent, which they defined as the clustering of similar activities, is helpful to activity detection. Fully-connected CRFs were applied as a post-processing of per-frame CNN features as well as object features. While such approach learns some global context, they do not learn the explicit temporal structure, nor are they learned in an end-to-end fashion.

To our knowledge, this is the first work exploring an end-to-end model for super-event representation learning, capturing temporal structure and relationships between activities. Learning hierarchical structures of activities has been studied in many traditional works [17, 8, 15, 11, 2, 19], but they were not learned end-to-end or required additional labels for intervals. Our model is fully differentiable, enabling joint end-to-end learning of latent super-events and the activity detector.

3. Activity detection with latent super-events

The objective of our model is to annotate each frame as its corresponding activity class (including no-activity) given a continuous video. We present an end-to-end learning model that does such labeling not only by looking at each frame or each local segment but also by considering the overall temporal structure. Our idea is to allow the model

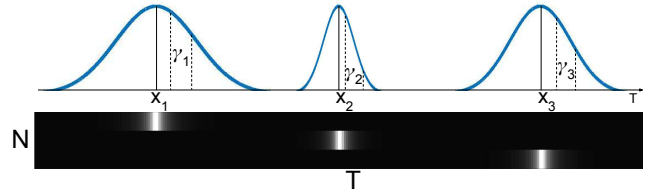


Figure 3. An illustration of a temporal structure filter. Each filter learns independent centers and widths.

to learn the representations of *super-events* summarizing much longer-term temporal intervals, and take advantage of them for the activity detection. A super-event is defined as a longer-term event containing the event of interest, which is the opposite concept of ‘sub-events’. Our approach, shown in Fig. 2, effectively represents super-events using temporal structure filters. We learn per-class soft-attention weights over the filters to create a super-event representation, and take advantage of it for the frame classification.

3.1. Per-frame representation

The base component of our detection is a CNN providing per-frame (or per-local-segment) representation. This is obtained by learning standard video CNN models (e.g., [26, 3]). We train the model to learn binary per-frame classifiers by optimizing:

$$L(v) = \sum_{t,c} z_{t,c} \log(p(c|v_t)) + (1 - z_{t,c}) \log(1 - p(c|v_t)) \quad (1)$$

where v_t is the per-frame or per-segment CNN feature at frame t and $z_{t,c}$ is the ground truth label for class c and time t . This gives a sequence of probabilities for each class which can be used to find activity intervals. Using a fully-connected network to model $p(c|v_t)$ captures minimal temporal information, using just a single frame or segment. RNN models have been used to compute $p(c|v_t)$ which captures some implicit temporal information. The learned CNN producing v_t serves as our base component.

3.2. Temporal structure filter

The *temporal structure filter* we introduce in this paper is a filter designed to capture temporal context formed by multiple activities. It is an extension/generalization of the spatial attention model proposed in [9]. The idea is to represent a variable length video with a fixed dimensional vector, by only focusing on the ‘learned’ frame locations. The previous attention model repeats a single Gaussian distribution several times with a fixed stride, while a temporal structure filter learns several independent distributions.

In particular, each temporal structure filter is modeled as a set of N Cauchy distributions: we found that the Cauchy distribution was easier to train than the Gaussian distribution commonly used (i.e., it converges faster). Each distri-

bution learns a center, x_n and γ_n which controls the width. Given T , the length of the video, each filter is constructed by:

$$\begin{aligned} \hat{x}_n &= \frac{(T-1) \cdot (\tanh(x_n) + 1)}{2} \\ \hat{\gamma}_n &= \exp(1 - 2 \cdot |\tanh(\gamma_n)|) \\ F[t, n] &= \frac{1}{Z_n \pi \hat{\gamma}_n \left(\frac{(t - \hat{x}_n)^2}{\hat{\gamma}_n} \right)} \end{aligned} \quad (2)$$

where Z_n is a normalization constant, $t \in \{1, 2, \dots, T\}$ and $n \in \{1, 2, \dots, N\}$. Figure 3 shows an example.

When trained for the super-event representation learning, our temporal structure filter allows the model to explicitly learn which temporal intervals in the entire video are relevant for the frame-level event detection.

3.3. Super-event representation learning

In its elementary form, we represent the super-event of each event c by applying one temporal structure filter F_c over the entire video representation v . This essentially is a matrix product between each of the N distributions in F_c and v . The super-event representation S_c is obtained as:

$$S_c[n] = \sum_t^T F_c[t, n] \cdot v_t. \quad (3)$$

This operation applies F_c (an $T \times N$ filter) to v (the $T \times D$ video features) and returns an $N \cdot D$ -dimension vector, which we call the super-event representation. This has an effect of summarizing the entire video representation v into a much smaller representation S_c by focusing on the frames specified with the learned parameters of F_c .

Super-event representation with soft attention: As several activities can share the same super-event, it makes sense to learn a set of M different temporal structure filters and share these filters across the classes. Here, M is less than the number of classes C . In order to represent the super-event of each activity class c using such M filters, we learn a set of per-class soft-attention weights allowing each activity class to select some of the M structure filters to use. For a set of C classes, we learn weights $W_{c,m}$ and compute the soft-attention as:

$$A_{c,m} = \frac{\exp(W_{c,m})}{\sum_k^M \exp(W_{c,k})}. \quad (4)$$

We then create a super-event representation by applying these weights to the M temporal structure filters:

$$S_c = \sum_m^M A_{c,m} \cdot \sum_t^T F_m[t] \cdot v_t \quad (5)$$

3.4. Detection with super-events

We perform per-frame binary classification for each class by concatenating the super-event representation with the CNN frame representation:

$$p(c|[v_t, S_c]) = \sigma(W[v_t, S_c]) \quad (6)$$

where W is a learnable parameter and σ is the sigmoid function. Unlike standard per-frame classification approaches, each frame now depends on the abstracts from the entire video features (i.e., our super-event representation) instead of a single frame.

To learn parameters, we optimize the multi-label binary classification loss:

$$\begin{aligned} L(v) &= \sum_{t,c} z_{t,c} \log(\sigma(W[v_t, S_c])) \\ &\quad + (1 - z_{t,c}) \log(1 - \sigma(W[v_t, S_c])) \end{aligned} \quad (7)$$

where $z_{t,c}$ is the ground truth label for class c at time t . We minimize this loss using stochastic gradient descent.

3.5. Relative super-events

We also propose a relative super-event model, where instead of computing the super-event representation for the entire video, we compute a super-event representation relative to the current frame location. This allows us to capture relative relationships like in basketball, shooting and blocking must occur together, regardless of the other content in the video. This works well even when a video contains multiple, unrelated activity classes, for example a highlight video of various sports. This approach is identical to the one above, except the length L of the temporal structure filter is fixed. We use the temporal structure filters F as a convolutional kernel and convolve with the video features, applying F centered at each feature t with a length of L frames. The per-class attention weighting is applied to the structure filters to form a per-frame super-event representation used for the classification of the current frame.

4. Experiments

In order to evaluate the effectiveness of our proposed super-event representation learning, we conducted a set of experiments comparing our method to conventional approaches across various datasets and feature types. We particularly focused on activity detection datasets with videos containing multiple actions/activities, including MultiTHUMOS [31], Charades [25], and AVA [10].

4.1. Settings

Features: We extracted I3D [3] features from the videos at 25fps with a stride of 8 frames. We are using I3D as the base per-segment CNN mentioned in Section 3.1. I3D

is a two-stream 3D CNN that achieved state-of-the-art performance on several action recognition tasks. Using both RGB and optical flow as input, it is able to capture relatively longer-term temporal relationships using a temporal resolution of up to 99 frames. We obtained pre-trained weights for I3D as well as code to fine-tune I3D from the authors, which is identical to what they used for the Charades Challenge 2017. We also tested VGG RGB [27], VGG optical flow, and standard two-stream CNNs [7] at 8fps as our base per-frame CNNs. The optical flow and two-stream features have a temporal resolution of 10 frames and RGB features have a temporal resolution of 1 frame.

Implementation Details: We implemented the temporal structure filters in PyTorch. The learning rate was set to 0.1 and reduced by a factor of 10 every 1000 iterations. We trained the network using a batch size of 32 videos for 5000 iterations using the Adam [14] optimizer with default parameters. We applied dropout with a probability of 0.5 to the input features. We set $N = 3$ (i.e., 3 Cauchy distributions per temporal structure filter), and $M = 5$ (i.e., 5 temporal structure filters). Per-frame CNN representation dimensionality is $D = 1024$ for I3D features and $D = 4096$ for VGG features. In both cases, we used the output of the layer before the final fully-connected layer as the features. In order to make the final per-frame decisions based on per-segment CNN features and our super-event representations, we trained a single fully-connected layer with input size $D + D \cdot N$ and output size of the number of classes. Our code and pretrained models are available at <https://github.com/piergiaj/super-events-cvpr18>.

4.2. MultiTHUMOS

Dataset: MultiTHUMOS [31] is an extension of the THUMOS [12] dataset with the untrimmed videos densely annotated for 65 different action classes. Unlike ActivityNet and THUMOS, MultiTHUMOS has on average 10.5 activity classes per video, 1.5 labels per frame, and up to 25 different activity instances in each video. This allows us to confirm the value of learning super-event representations in complex videos. MultiTHUMOS contains YouTube videos of various sport activities, such as basketball games (shoot, guard, dribble, and block), volleyball games (serve, set, block, and spike), types of weightlifting, throwing, etc.

We followed the standard activity detection evaluation setting of MultiTHUMOS, which is measuring the mean average precision (mAP) by annotating each frame in test videos. We actually used fewer training samples than provided: We only used the 1010 continuous videos for our training, and test on the full 1574 test videos. We did not use the segmented training videos from THUMOS, since super-event learning is not meaningful with trimmed videos. Even without using the full dataset, we were able to outperform

Table 1. Performances of the state-of-the-art methods and our approach on MultiTHUMOS. Our approach meaningfully outperforms all previous results.

	mAP
Two-stream [31]	27.6
Two-stream + LSTM [31]	28.1
Multi-LSTM [31]	29.6
Predictive-corrective [4]	29.7
I3D baseline	29.7
I3D + LSTM	29.9
I3D + Temporal Pyramid	31.2
I3D + our super-event	36.4

Table 2. Performances of different super-event representations, evaluated using MultiTHUMOS. I3D was used as the base CNN.

	mAP
Baseline	29.7
Global max pooling	30.0
Global mean pooling	30.8
Temporal pyramid pooling	31.2
Single super-event (per-class)	31.2
With soft attention	36.4
With soft attention + relative	36.2

the previous approaches.

Results: Table 1 shows the performance (mAP) of our approach compared against previously reported results of the state-of-the-art methods. We are able to observe that our method outperforms previous approaches by a significant margin, achieving a mAP of 36.4%. This is meaningfully higher than our baseline (i.e., I3D) and the previous best approach [4]: 29.7% vs. 36.4%. Using the pre-trained I3D [3] model provided the accuracy identical to the previous best reported performance, and our approach of learning and using latent super-events on top of I3D outperformed it by the margin of 6.7%. Note that this is also higher than using a sequential recurrent model like LSTM on top of the same feature, also by a margin of more than 6%. In addition, we also tested the method of using a temporal pyramid commonly used in the previous works (e.g., [20]) as a super-event representation. Using the temporal pyramid (of level 3) gave us the performance of 31.2 on MultiTHUMOS (while ours was 36.4).

We conducted more experiments to compare different super-event representations in our approach. Table 2 compares (1) the baseline per-frame classifier, (2) the method of using global mean/max pooling as a super-event representation, (3) temporal pyramid pooling, (4) the elementary single super-event representation mentioned in Section 3.3, (5) our soft-attention-based super-event representation, and (6) the relative super-event approach. We found that while the global max/mean pooled representation and the single

super-event improves performance (compared to the baseline), using the soft-attention and shared temporal structure filters gave the best performance. Relative super-events performed similarly to non-relative super-events.

Figure 4 illustrates a qualitative analysis of our results. We can observe that super-events improve the detection of related basketball events, especially the co-occurrence of the shooting and blocking actions. The learned temporal structure filters and super-event representations are shown in Figure 6. The filters are visualized by combining our learned temporal structure filters using their soft-attention weights. We obtain 5 TxN filters, and added them based on the learned attention weights.

4.3. Charades dataset

Dataset: Charades [25] is a large scale dataset consisting of 9848 videos across 157 activities. The videos were recorded by people in their own homes based on a provided script. Each video contains on an average of 6.8 activity instances, often with complex co-occurring activities, making it a suitable dataset to test our super-event learning. The activities were mainly performed at home, and their classes include ‘preparing a meal’, ‘eating’, ‘sitting’, ‘cleaning’, etc.

In our experiments, we follow the original Charades test setting (i.e., Charades_v1_localize evaluation). This is a bit different from the newer test set released for the Charades Challenge 2017. The ground truth labels of the challenge test videos are not publicly available and its evaluation server is not approving any new account access. Instead, we followed the original Charades localization test setting (v1) from the dataset website. The original setting is more challenging than the competition setting in the aspect that the competition allowed more training data to be used than the original setting: the participants were able to use the videos of the original test set as a part of the training in the competition. Note that the exact same I3D code (with fine-tuning) that provided 20.72 mAP in the competition setting is only providing 17.22 mAP in the original test setting.

Similar to MultiTHUMOS, the performances are measured in terms of mAP by evaluating per-frame annotations.

Results: We compared the effectiveness of our approach of using super-event representations with multiple different base per-frame/per-segment CNNs. More specifically, we used the I3D features across RGB frames, flow frames, and two-stream with and without fine-tuning on the Charades dataset. We also tested our approach with the standard two-stream CNN using VGG models. These features capture various amounts of temporal information, from a single RGB frames (VGG RGB), to 10 optical flow frames (VGG Flow and VGG two-stream) and up to 99 frames for I3D.

Table 3 shows the results describing how much our latent

Table 3. Results comparing our approach using super-events with the baselines and LSTMs on the Charades dataset v1. These numbers are raw results without the post-processing method of [23].

	Baseline	LSTM	Ours
Two-stream CNN			
VGG RGB	6.11	6.23	7.64
VGG Flow	5.13	6.08	6.83
VGG Two-Stream	6.56	7.85	8.53
Pre-trained I3D			
I3D RGB	9.78	9.91	11.0
I3D Flow	9.58	9.81	11.0
I3D Two-Stream	10.32	11.50	12.8
Fine-tuned I3D			
I3D RGB	15.63	17.03	18.64
I3D Flow	16.21	17.54	18.52
I3D Two-Stream	17.22	18.12	19.41

Table 4. Results on Charades original dataset (i.e., Charades_v1_localize setting). Note that this setting is a bit different from the Charades Challenge 2017 competition setting, whose evaluation server is not approving any new account access. This setting uses less training data.

	mAP
Random [23]	2.42
RGB [23]	7.89
Predictive-corrective [4]	8.9
Two-Stream [23]	8.94
Two-stream+LSTM [23]	9.6
R-C3D [30]	12.7
Sigurdsson et al. [23]	12.8
I3D [3]	17.22
I3D + LSTM	18.1
I3D + Temporal Pyramid	18.2
I3D + super-events	19.41

super-event learning approach improves the activity detection performance for each per-frame/per-segment baseline. We are reporting the performances of our method using the super-event representations with soft attention. Furthermore, we implemented the standard LSTM method over the same baseline CNNs. The idea was to directly compare the abilities of our approach and the LSTM in capturing long-term temporal dynamics/relations. We found that super-events yield meaningful performance increases regardless of the feature type. We also outperform the LSTM models, confirming that super-events are better able to capture and use temporal structure than LSTMs.

We compare our results with the state-of-the-art in Table 4. Our method is obtaining the best known performance in the localization setting of the Charades dataset. Notably, it is performing better than I3D which obtained the best competition performance in 2017, while using the same feature. Figures 5 and 7 show example detections.

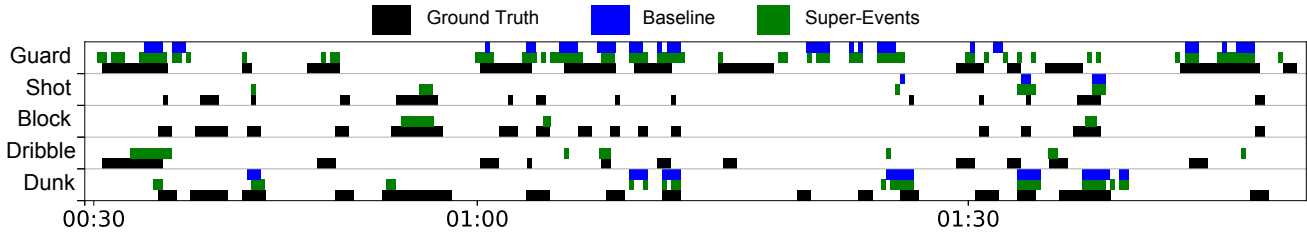


Figure 4. Results from a video in MultiTHUMOS. Super-events especially help the detection of the shooting and blocking events.

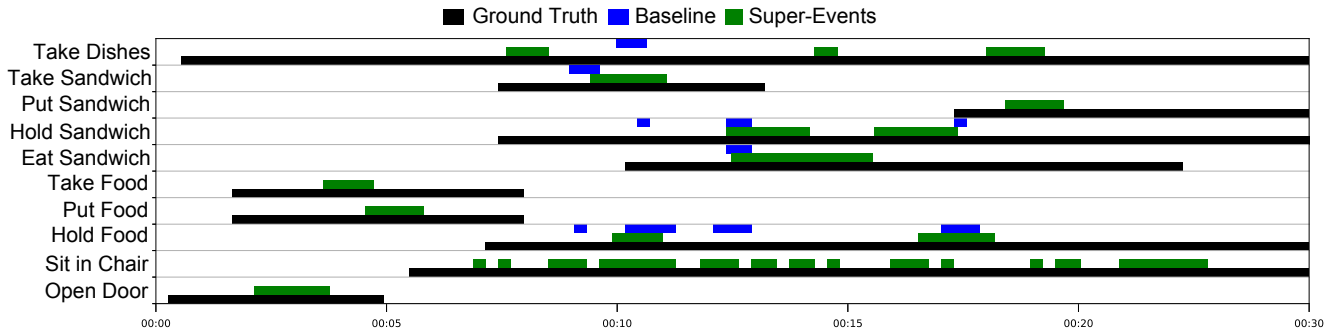


Figure 5. Results from a video in Charades. Super-events improve the detection of related events such as taking, holding and eating a sandwich.

4.4. AVA dataset:

Dataset: AVA [10] is a new large-scale video dataset containing 80 action classes in 57,600 video clips drawn from 192 movies. Unlike Charades, which has individual activity classes per objects, the activities in AVA are very generic, such as sit, stand, walk, carry, etc. A 15 minute segment is selected from each movie and annotated in 3 second intervals for spatial and temporal activity detection. Since we are interested in temporal activity detection, we designed a new setting using AVA in which we label each frame with the activities that are occurring in it regardless of spatial location. Identical to Charades (and MultiTHUMOS), we evaluate our models using per-frame mean average precision (mAP). Because each 3 second interval contains the same labels, we average our predictions and produce one probability vector per 3 second interval. We only use temporal annotations in this experiment.

Results: Table 5 compares our approach with random, I3D baseline, and I3D + LSTM. We used RGB, Flow, and Two-stream versions of I3D. Again, we are reporting the performances of our method using super-event representations with soft attention. The results are very consistent with the results we obtained from MultiTHUMOS and Charades. Learning (latent) super-event representations with our method and using them for the activity detection always achieved the higher performance. Using LSTM was also able to improve the vanilla baseline, but our approach always outperformed the LSTM by a meaningful margin.

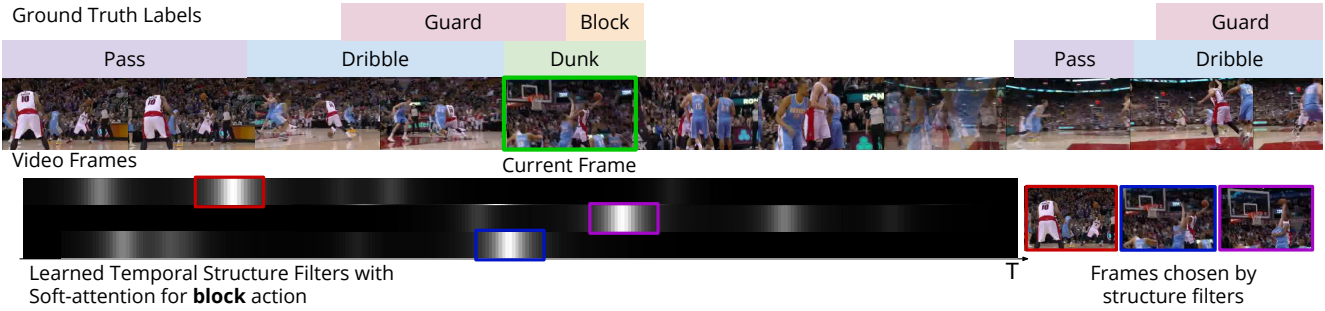
Table 5. Results on Ava Test Set (mAP), with the temporal annotation setting. That is, we evaluated the accuracy solely based on frame-level annotations.

	Baseline	LSTM	Ours
Random	2.65	2.65	2.65
I3D RGB	6.8	7.0	8.3
I3D Flow	7.1	7.2	9.1
I3D Two-Stream	7.5	7.8	9.8

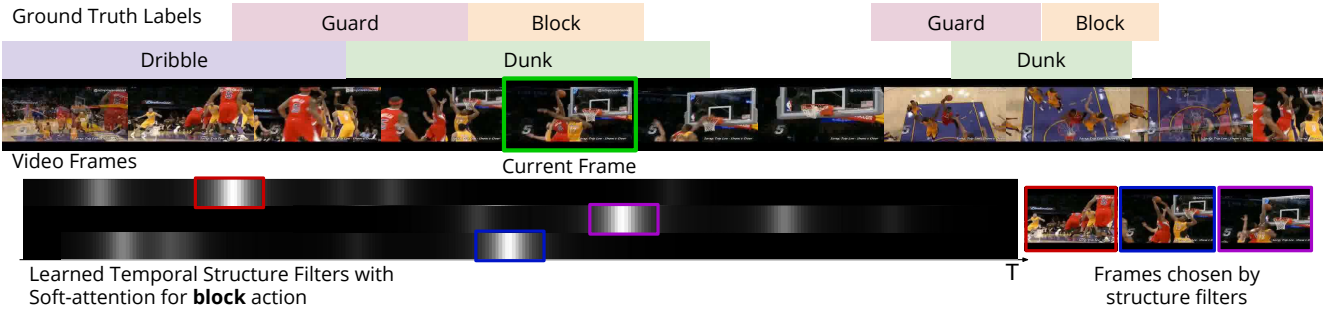
5. Conclusion

We introduced the concept of learning latent super-events in activity videos. We defined a super-event as a set of multiple events occurring together in videos with a particular temporal structure (i.e., the opposite concept of sub-events). We newly designed the temporal structure filters, and presented how they can be used to capture temporal dynamics in multi-activity videos for the super-event representation learning. We provided a fully differentiable end-to-end architecture to jointly learn the latent super-events and the activity detector using them. We were able to confirm that our method performs superior to the state-of-the-art methods in multiple different activity detection datasets, including MultiTHUMOS and Charades.

Ack: This work was supported by the U.S. Army Research Laboratory under Cooperative Agreement DAAD 19-01-2-0012, and by the ICT R&D program of MSIP/IITP, Republic of Korea (17ZF1200, Development of XDMedia Solution for Invigoration of Realistic Media Industry).

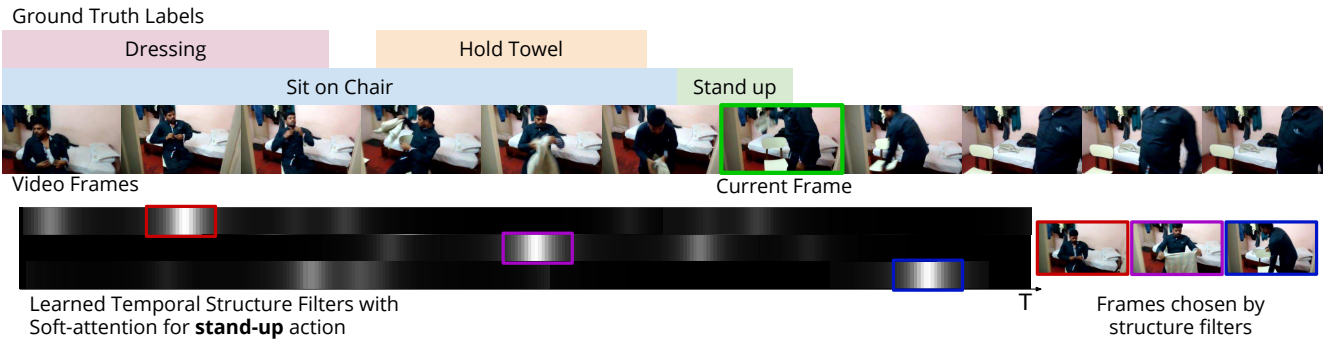


(a) Example video of the block action with passing and dribbling.



(b) Example video of the block action with dribbling.

Figure 6. Illustration of the learned temporal structure filters with soft-attention for the block action in MultiTHUMOS. When applied to two different videos, we can observe that the temporal structure filters capture the temporal relationships between the frames corresponding to shooting and blocking action, as well as the relationship between those corresponding to dribble/pass and blocking action.



(a) Example video of the stand-up action with the dressing and holding towel actions.



(b) Example video of the stand-up action with holding shoes and reading book actions.

Figure 7. Illustration of the learned temporal structure filters with soft-attention for the stand-up action in Charades. The filters are applied to two different videos. It shows that our temporal structure filters are able to capture the relations that the frames of a sitting action come before the frames of standing up, regardless of the other actions that may occur in the video.

References

- [1] J. K. Aggarwal and M. S. Ryoo. Human activity analysis: A review. *ACM Computing Surveys*, 43:16:1–16:43, April 2011.
- [2] W. Brendel and S. Todorovic. Learning spatiotemporal graphs of human activities. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 778–785. IEEE, 2011.
- [3] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [4] A. Dave, O. Russakovsky, and D. Ramanan. Predictive-corrective networks for action detection. *arXiv preprint arXiv:1704.03615*, 2017.
- [5] V. Escorcia, F. C. Heilbron, J. C. Niebles, and B. Ghanem. Daps: Deep action proposals for action understanding. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 768–784. Springer, 2016.
- [6] B. G. Fabian Caba Heilbron, Victor Escorcia and J. C. Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 961–970, 2015.
- [7] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1933–1941, 2016.
- [8] A. Gaidon, Z. Harchaoui, and C. Schmid. Actom sequence models for efficient action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3201–3208. IEEE, 2011.
- [9] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.
- [10] C. Gu, C. Sun, S. Vijayanarasimhan, C. Pantofaru, D. A. Ross, G. Toderici, Y. Li, S. Ricco, R. Sukthankar, C. Schmid, and J. Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. *arXiv preprint arXiv:1705.08421*, 2017.
- [11] A. Gupta, P. Srinivasan, J. Shi, and L. S. Davis. Understanding videos, constructing plots learning a visually grounded storyline model from annotated videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2012–2019. IEEE, 2009.
- [12] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://crcv.ucf.edu/THUMOS14/>, 2014.
- [13] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1725–1732, 2014.
- [14] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] B. Laxton, J. Lim, and D. Kriegman. Leveraging temporal, contextual and ordering constraints for recognizing complex activities in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8. IEEE, 2007.
- [16] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4694–4702. IEEE, 2015.
- [17] J. C. Niebles, C.-W. Chen, and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 392–405. Springer, 2010.
- [18] A. Piergiovanni, C. Fan, and M. S. Ryoo. Learning latent sub-events in activity videos using temporal attention filters. In *Proceedings of the American Association for Artificial Intelligence (AAAI)*, 2017.
- [19] M. S. Ryoo and L. Matthies. First-person activity recognition: What are they doing to me? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [20] M. S. Ryoo, B. Rothrock, and L. Matthies. Pooled motion features for first-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 896–904, 2015.
- [21] Z. Shou, J. Chan, A. Zareian, K. Miyazawa, and S.-F. Chang. Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. *arXiv preprint arXiv:1703.01515*, 2017.
- [22] Z. Shou, D. Wang, and S.-F. Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1049–1058, 2016.
- [23] G. A. Sigurdsson, S. Divvala, A. Farhadi, and A. Gupta. Asynchronous temporal fields for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [24] G. A. Sigurdsson, O. Russakovsky, and A. Gupta. What actions are needed for understanding human actions in videos? *arXiv preprint arXiv:1708.02696*, 2017.
- [25] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016.
- [26] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NIPS)*, pages 568–576, 2014.
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [28] D. Tran, L. D. Bourdev, R. Fergus, L. Torresani, and M. Paluri. C3d: generic features for video analysis. *CoRR*, abs/1412.0767, 2(7):8, 2014.
- [29] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *Proceedings of the*

IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3169–3176. IEEE, 2011.

- [30] H. Xu, A. Das, and K. Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. *arXiv preprint arXiv:1703.07814*, 2017.
- [31] S. Yeung, O. Russakovsky, N. Jin, M. Andriluka, G. Mori, and L. Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision (IJCV)*, pages 1–15, 2015.
- [32] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2678–2687, 2016.
- [33] Y. Zhao, Y. Xiong, L. Wang, Z. Wu, X. Tang, and D. Lin. Temporal action detection with structured segment networks. *arXiv preprint arXiv:1704.06228*, 2017.