

This ICCV Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Weakly-Supervised Completion Moment Detection using Temporal Attention

Farnoosh Heidarivincheh, Majid Mirmehdi and Dima Damen Department of Computer Science University of Bristol, Bristol, UK

{farnoosh.heidarivincheh, M.Mirmehdi, Dima.Damen}@bristol.ac.uk

Abstract

Monitoring the progression of an action towards completion offers fine grained insight into the actor's behaviour. In this work, we target detecting the completion moment of actions, that is the moment when the action's goal has been successfully accomplished. This has potential applications from surveillance to assistive living and human-robot interactions. Previous effort [14] required human annotations of the completion moment for training (i.e. full supervision). In this work, we present an approach for moment detection from weak video-level labels. Given both complete and incomplete sequences, of the same action, we learn temporal attention, along with accumulated completion prediction from all frames in the sequence. We also demonstrate how the approach can be used when completion moment supervision is available. We evaluate and compare our approach on actions from three datasets, namely HMDB, UCF101 and RGBD-AC, and show that temporal attention improves detection in both weakly-supervised and fully-supervised settings.

1. Introduction

In the past few decades, with an increasing ubiquity and accessibility of video records, a significant part of research has been devoted to analysing human behaviour in video, including analysing human actions. This active area of research has vast applications, such as health-care, surveillance, video retrieval, entertainment, robotics and human-computer interaction. At its heart, action recognition has evolved from traditional hand-crafted features [18, 34, 35] to deep learning based approaches [16, 31, 4, 11] and achieved remarkable results. However, recent works have focused on proposing network architectures to deal with the spatio-temporal input, and neglected to explore fake or incomplete action instances.

In this paper, we focus on incomplete actions, whether intentional or accidental, which could be crucial in contexts such as surveillance and health-care applications. These are actions which are attempted but their goals remain incomplete. Such incomplete sequences could be incorrectly recognised, or localised, by current state-of-the-art methods. As an example, consider an *incomplete pick*, where the subject only pretends to pick an object up. Standard action recognition classifiers would identify this as a *pick* action, because of its similar motion to successfully completed picks. Likewise, a patient picking up a medicine tablet, but not ingesting it, would also be incorrectly recognised as "take medicine" action, posing risks to automatic monitoring of their health.

Action completion was first introduced in [13] to assess whether the action's goal is achieved. The approach outputs sequence-level predictions of completion to distinguish *complete* sequences from *incomplete* ones. Subsequent works [14, 2] proposed finer-grained analysis of an action's progression towards completion. For example, [14] looks for visual clues which confirm the goal's completion and detects the *completion moment* from frame-level *pre-* and *post-completion* labels.

In this work, we also investigate *completion moment* detection. However, we differ from [14] in the supervision by which our method learns to detect completion. Frame-level annotations are not only expensive to collect, but importantly, highly subjective and often noisy [24, 30, 23]. We offer the first attempt to *completion moment* detection with weak supervision, i.e. using only sequence-level *complete* and *incomplete* labels. Fig 1 illustrates frame-level and sequence-level labels for a *complete pick* action. Given weak labels, we show that *completion moment* detection could be achieved, by learning temporal attention.

We propose to use convolutional and recurrent cells with learnt temporal attention, and accumulate evidence for completion from all frames along the sequence, where evidence is weighted according to the frame's importance to the *completion moment* prediction. A similar approach was attempted in [14], however with full supervision and without temporal attention, where all frames contributed equally to the *completion moment* detection. We show that our proposed approach outperforms [14] when fully supervised,



Frame-level *pre-* and *post-completion* labels Sequence-level *complete* label Figure 1. While frame-level labels (left) are used in a fully-supervised approach like [13], we detect the *completion moment* only using weak labels, i.e. sequence level *complete* and *incomplete* (right).

but importantly is also able to detect *completion* using weak video-level supervision.

We evaluate our approach on selected actions from HMDB [17], UCF101 [33] and RGBD-AC [13]. We show that learning temporal attention decreases the completion detection error, i.e. the relative distance between the predicted and ground truth *completion moment*, by 15% of the sequence length with weak supervision and by 3% when fully supervised.

The remainder of this paper is organised as follows: related work in Sec. 2, proposed method in Sec. 3, experiments and results in Sec. 4 and conclusion and future work in Sec. 5.

2. Related Work

In this section, we differentiate our work from approaches that attempted moment detection in actions, including for action completion. We also review works that utilised temporal attention learning, including for action localisation.

Moment Detection: Temporal action detection from untrimmed videos [38, 29, 5] involves localising start and end points of actions. These works assume all actions are successfully completed, and do not consider incomplete attempts. A few methods [2, 15, 9, 43, 22], on the other hand, have adopted approaches that model action progression or detect particular key moments within actions. Ma et al. [22] detect an action by learning its progression through time. They devise a loss which maximises the margin between the correct action class and other classes as the action progresses further. Hoai and De la Torre [15] also detect actions in untrimmed sequences, where the action progression is modelled by a score function, learned using a Support Vector Machine classifier that peaks when the action ends. Similarly, Becattini et al. [2] attempt to recognise actions by modelling their evolution through time, where the progress is assumed to be linear, reaching the highest at the end of the sequence. Dwibedi et al. [9] propose a selfsupervised approach to learn temporal alignment between sequences based on the similarity between their frames and then observe an action's progression between key frames

given the learnt alignment. Yeung et al. [43] detect actions by looking at individual frames through the sequence, where the location of the next input frame is predicted relative to the current frame. Although these works present a fine-grained analysis from the action progression, they also consider complete attempts and do not detect or localise the *completion moment*.

Action completion [13] differs from these works, as it focuses on the action's *goal*. In [14], *completion moment* detection was addressed using a classification-regression network which outputs frame-level predictions. These predictions are accumulated, using voting, to detect the *completion moment*. However, the method is fully supervised, requiring the *completion moment* annotations for training. In contrast, we solve the same problem using only sequencelevel *complete* and *incomplete* labels (i.e. *weak labels*), through utilising temporal attention learning.

Attention Learning has proven beneficial for research problems, such as image captioning [40, 41, 6], object detection and tracking [3, 7, 1] and person reidentification [12, 20, 39]. Recently action recognition and localisation have also used attention networks to learn which spatial and/or temporal regions contain the most discriminative information. While some works [28, 10, 21, 37] have only focused on frame-level attention (spatial and motion), many others [8, 32, 19, 27, 42, 36, 26, 25] have also incorporated temporal attention in their models. They learn attention scores on the temporal dimension which are then used to weight the frames according to their importance to the final prediction. Of these, Pei et al. [27] introduce a recurrent unit for sequence classification in which a high attention score at each time step pushes the network to focus on the current observations rather than the past ones. Song et al [32] use LSTM for learning temporal attention from skeleton data in action recognition. Du et al. [8] also propose an approach for action recognition using an LSTM with temporal softmax normalisation. Weighted observations, through learnt attention, from all frames in the sequence are combined to recognise the current frame's ongoing action.

For action localisation with weak supervision, several

approaches have also attempted learning temporal attention, such as [36, 26, 25, 19, 42]. For example, Yeung et al. [42] learn temporal attention for dense labelling in action localisation. Since they use trimmed sequences for training, but apply the learnt attention to localise actions in untrimmed sequences, their detection is considered weakly supervised. Other works, however, use untrimmed sequences in training. Li et al. [19] apply attention for action recognition and action detection in untrimmed sequences, using features from multiple modalities as the input to the temporal attention LSTM before softmax normalisation. Nguyen et al. [25] learn attention for action classification. They normalise the attention scores by a sigmoid function, and then use these to estimate the discriminative class-specific temporal regions for localising actions. Wang et al. [36] predict the action's temporal extents by combining hard and soft selection methods, where the soft selection relies on the attention weights for the clip proposals sampled from the untrimmed sequences. In [26], the attention scores are first predicted as a temporal softmax on the class-wise activations and used during training. They then apply a threshold on class-wise activations for localising actions.

In our method, we also use an LSTM for learning attention and a temporal softmax for its normalisation. However, our method differs not only in the problem of *completion moment* detection, but in how we accumulate evidence from all frames in the sequence based on learnt attention. We localise the *completion moment* within trimmed sequences for both training and evaluation.

3. Temporal Attention for *Completion Moment* Detection

We now present our approach to **weakly-supervised** completion moment detection, when video-level annotations only are present. Assume $X_i = \{x_i^1, \dots, x_i^T\}$ are the frames in a sequence of length T where an action has been attempted, and $y_i \in \{0, 1\}$ is the binary video-level label, indicating whether the attempt has been successfully completed or not. Our method takes as input both complete and incomplete sequences of the same action.

To predict the *completion moment* with weak supervision, we propose a network architecture that contains a convolutional frame-level feature extracting network, followed by two recurrent cells for completion prediction and temporal attention prediction, trained jointly with a cross-entropy loss function. Fig 2 depicts our architecture, showing the per-frame feature extraction and recurrent nodes (left) along with the training loss (top left). The frame-level predictions are then accumulated (right) to infer the *completion moment*.

For feature extraction, we train a convolutional network, by propagating the video-level label y_i to *all* frames in a video, and optimise it using the cross-entropy loss,

$$L = \sum_{i=1}^{M} \sum_{t=1}^{T_i} - \left(y_i log f^c(f(x_i^t)) + (1 - y_i) log (1 - f^c(f(x_i^t))) \right)$$
(1)

where M is the number of sequences. The loss is optimised over all frames in all sequences, comparing the video-level labels y_i against classification outputs $f^c()$, while framelevel features $f(x_i^t)$ are accordingly trained. These learnt features form a good base for *completion moment* detection, to be refined by the recurrent cells. This is based on the realistic assumption that, up to the completion moment, both *complete* and *incomplete* sequences are indistinguishable. However, after completion, there are appearance distinctions between the frames, to signify completion.

We train two recurrent models, namely LSTMs, jointly, one for temporal attention, i.e. to learn the relevance of each frame t to completion moment detection, a_t , and one to predict temporally-evidenced completion scores, s_t . The temporal attention network is a standard LSTM, taking the features $f(x_t)$ as input - note that we simplified the notation x_i^t to x_t as the LSTM is trained and evaluated on one sequence. We compute the attention scores by applying a softmax function to the output nodes of this LSTM, o_t^a , across the temporal dimension, such that

$$a_t = \frac{e^{o_t^a}}{\sum\limits_{j=1}^T e^{o_j^a}} \,. \tag{2}$$

The second LSTM, also takes the same input $f(x_t)$, and its output o_t^s is then combined with the attention a_t to produce completion scores per frame

$$s_t = \frac{1}{1 + e^{-a_t o_t^s}} \,. \tag{3}$$

The scores s_t are the confidence of observing completion at frame t. In other words, a frame with a high s_t has observed distinctive signatures for completion, making it more confident that the sequence has been *completed*, with $1 - s_t$ reflecting the confidence for incompletion. We use these frame-level predictions to compute the *completion moment*, such that

$$\hat{y} = \arg\max_{j} \left(\sum_{t=1}^{j} (1 - s_t) + \sum_{t=j+1}^{T} s_t \right).$$
 (4)

The predicted *completion moment* \hat{y} is one where the score for completion beyond frame j as well as the score for incompletion before frame j are the maximum.

During training, only video-level labels are available, and the ground-truth *completion moment* is unknown. We



Figure 2. Weakly supervised, the learnt attention scores (red distribution) and completion confidence scores are combined to infer the *completion moment* \hat{y} (Eq 4). The frames depict a *complete* sequence of action *blowing candles* and colors green and blue represent the observed evidence for completion and incompletion, respectively.

thus train for sequence level prediction, such that

completion moment is finally predicted as

$$\hat{y}_{i}^{tr} = \frac{1}{1 + e^{-\sum_{t} (a_{t} o_{t}^{s})}}, \qquad (5)$$

$$L = \sum_{i} -(y_i \hat{y}_i^{tr} + (1 - y_i)(1 - \hat{y}_i^{tr})), \quad (6)$$

where y_i^{tr} indicates whether the sequence has been completed, *somewhere* along its frames. These predictions are optimised against the video-level completion labels, for all sequences. Note that, using a_t in the training loss makes the model learn to weight highly the temporal regions which contain discriminative evidence for completion.

While focusing on weakly-supervised *completion moment* detection, we also evaluate our proposed architecture in a **supervised** approach. We similarly combine completion detection with temporal attention, when supervision for the *completion moment* is available. We thus train the output of the confidence scores o_t^s in the same way as the regression-based supervision in [14], using the relative distance $r_t = \frac{t-\tau_i}{\tau_i}$ between the frame t and the ground-truth *completion moment* τ_i , allowing the approaches to be directly comparable. The sequence-level loss L_i would then be:

$$L_i^{supervised} = \sum_{t=1}^{T_i} a_t (o_t^s - r_t)^2.$$
 (7)

Using these scores,

$$s_t^{supervised} = a_t \frac{t}{o_t^s + 1},\tag{8}$$

estimates the *completion moment* from each frame, weighted by the learnt attention scores. The sequence-level

$$\hat{y}^{supervised} = \sum_{t=1}^{T} s_t.$$
(9)

Fig. 3 illustrates the supervised completion detection where the frame-level evidences are accumulated across the sequence during inference.

4. Experimental Results

Dataset and Implementation Details - We evaluate our approach on the 16 actions used in [14] as the only prior work to attempt *completion moment* detection, and using the publicly available annotations provided by [14]. These actions have been collected from three public datasets: HMDB [17], UCF101 [33] and RGBD-AC [13]. As stated in [14], these actions cover sport-based and daily actions, for which completion can be defined, and include both complete and incomplete sequences for training. We report results on all 16 actions when supervised. However, in the weakly supervised setting, we require sufficient *incomplete* sequences per action to be able to train with only videolevel weak labels. Of these 16 actions, we only evaluate on 10 actions which have both complete and incomplete sequences, while the remaining 6 have less than 5% incomplete sequences.

For feature extraction, we used the spatial stream of VGG-16 architecture, pre-trained on UCF101. We then fine-tuned it for 20 epochs to acquire frame-level features. The learning rate was started at 10^{-3} , divided by 10 at epochs 3 and 5. The features were extracted from the output of the fc7 layer. Both LSTM cells (attention and



Figure 3. In the Supervised model, the sequence level *completion moment* is the weighted average of the frame-level predictions, using the attention scores.

completion moment prediction) had a single layer with 128 hidden units. When fully supervised, we first trained the completion prediction LSTM s_t for 10 epochs for stability, then jointly trained both LSTMs for 5 more epochs. When weakly-supervised, we initialised both LSTMs from random and trained them jointly for 10 epochs. The learning rates for the LSTM training in both approaches was 10^{-2} for the first 5 epochs and then was divided by 10 for the rest. For temporal prediction, we normalised the sequences to a fixed length, equal to the minimum length of any sequence in that action. Note that our method is not dependent on the sequence length and thus is robust to any other pre-specified length. Additionally, the attention scores were normalised between zero and one and those less than 0.5 were truncated to 0 during inference .

Evaluation Metrics – As in [14], we report the accuracy as the average percentage of frames that are correctly labeled into *pre*- and *post-completion*, given the ground-truth τ_i and the predicted *completion moment* \hat{y}_i , such that

Accuracy =
$$\frac{1}{M} \sum_{i=1}^{M} \frac{1}{T_i} \sum_{t=1}^{T_i} \left[(t < \hat{y}_i \land t < \tau_i) \right] (10)$$
$$\vee (t \ge \hat{y}_i \land t \ge \tau_i) .$$

We also report RD as the relative distance between the predicted and ground truth *completion moment*, averaged on all sequences.

$$RD = \frac{1}{M} \sum_{i=1}^{M} \frac{||\hat{y}_i - \tau_i||}{T_i} \,. \tag{11}$$

Weakly Supervised Completion Detection – Table 1 shows the results of our proposed method for weakly supervised *completion moment* detection using uniform attention (WS-U) as well as with learnt temporal attention

		Incomplete	Accuracy		RD	
		%	WS-U	WS-Att	WS-U	WS-Att
HMDB	pick	22.4	34.8	48.6	0.65	0.51
UCF101	basketball	23.9	38.2	58.0	0.62	0.42
	soccer penalty	30.7	34.4	55.6	0.66	0.44
	blowing candles	45.9	43.8	70.9	0.56	0.29
RGBD-AC	switch	47.8	86.9	89.9	0.13	0.1
	plug	49.3	62.0	78.6	0.38	0.21
	open	47.1	74.6	77.1	0.25	0.23
	pull	52.1	79.1	83.8	0.21	0.16
	pick	52.2	58.4	83.0	0.42	0.17
	drink	48.5	57.2	69.6	0.43	0.30
total		39.3	53.8	69.4	0.46	0.31

Table 1. Results comparing **weakly supervised** *completion moment* detection with and without temporal attention learning.

		Accuracy			RD		
		[14]	S-U	S-Att	[14]	S-U	S-Att
HMDB	catch	80.5	82.3	83.5	0.20	0.18	0.17
	drink	78.0	80.3	81.1	0.22	0.20	0.19
	pick	79.9	81.8	83.2	0.20	0.18	0.17
	pour	80.0	77.8	78.3	0.20	0.22	0.22
	throw	74.6	76.9	78.3	0.25	0.23	0.22
UCF101	basketball	79.5	82.8	83.6	0.20	0.17	0.16
	blowing candles	84.2	89.9	90.1	0.16	0.10	0.10
	frisbee catch	78.3	86.6	86.9	0.22	0.13	0.13
	pole vault	88.4	88.4	89.7	0.12	0.12	0.10
	soccer penalty	87.1	87.0	87.9	0.13	0.13	0.12
RGBD-AC	switch	98.1	94.6	98.2	0.02	0.05	0.02
	plug	96.1	95.1	96.8	0.04	0.05	0.03
	open	86.7	88.2	91.3	0.13	0.12	0.09
	pull	94.1	92.4	95.1	0.06	0.08	0.05
	pick	93.2	90.5	92.0	0.07	0.09	0.08
	drink	90.9	89.0	91.2	0.09	0.11	0.09
total		84.9	86.1	87.4	0.15	0.14	0.12

Table 2. Results comparing **fully supervised** *completion moment* detection with and without temporal attention learning.

(WS-Att). In WS-U, we do not learn attention, and use uniform weighting in inference. Learning temporal attention improves results for all actions and both metrics. For actions with a smaller percentage of *incomplete* sequences,





Example of failure

Figure 4. Qualitative results for weakly supervised *completion moment* detection. Top to bottom: UCF101-*soccer penalty*, RGBD-AC-*pull*, UCF101-*blowing candles* and HMDB-*pick*, respectively.

i.e. HMDB-*pick*, UCF101-*basketball* and UCF101-*soccer penalty*, the performance is lower for both metrics, though temporal attention consistently improves the results. In to-tal, i.e. on all sequences from the three datasets, RD drops to 0.31 with WS-Att.

We also present some qualitative results for the weakly supervised approach in Fig. 4 where the first bar depicts the completion scores s_t - with green and blue representing the observed evidence for completion and incompletion respectively. The attention is shown in red, and results in orange and purple represent *pre* and *post-completion* labels, respectively. In the first two sequences, the temporal attention significantly improves the results by correctly weighting the frames after completion where discriminative featuers are observed. In the third example from action *blowing candles*, while WS-U has been misled by the completion scores at the end of the sequence, WS-Att correctly detected no completion. The last sequence shows a failure case for action *pick*. We believe this would be improved with more *incomplete* sequences during training.

Supervised Temporal Attention Learning – Table 2 shows the results of the supervised approach, compared to the R-R method in [14], which is comparable to ours as it does not use frame-level *pre/post-completion* classification, but directly predicts the *completion moment*. We also compare uniform weighting (S-U) to the learnt attention (S-Att). Learning temporal attention outperforms uniform weighting on all 16 actions, and outperforms the baseline on 14 out of the 16 actions. In total, RD drops to 0.12 with S-Att.

We present qualitative results for our method, when supervised, in Fig 5. The first bar represents the frame-level regression **error**, i.e. $||o_t^s - r_t||$ (darker is lower error). The





Examples of failalure

Figure 5. Qualitative results, using supervised learning. Top left to bottom right: UCF101-pole vault, RGBD-AC-open, HMDB-throw and UCF101-basketball, respectively.



Figure 6. Frame-level regression errors (in blue) are depicted against the attention scores in both fully (red) and weakly (orange) supervised approaches. Note that '0' on the x-axis indicates the *completion moment* for all *complete* sequences, and the end of the sequence for all *incomplete* sequences.

examples of success (left) show two sequences from actions UCF101-*pole vault* and RGBD-AC-*open*. Temporal attention improves the *completion moment* detection for both complete (top left) and incomplete (bottom left) sequences, as high attention correctly aligns to regions with small prediction error. The examples of failure (right) represent two sequences from actions HMDB-*throw* and UCF101-*basketball*, where attention has not been able to pick the regions with small error. In the *basketball* example, the sequence is detected as *complete* with and without attention, despite being *incomplete*.

Frame-level Analysis – We plot the frame-level errors, as well as the attention scores, averaged for all actions, on the three datasets in Fig. 6. The figure shows lower prediction errors (blue), both *before* and *after* the completion moment, in two clear minimas. Increased confusion around the completion moment comes from the very similar features before *completion moment*. We also show the learnt temporal attention for both supervised (red) and weakly supervised (orange) approaches. Generally, higher attention corresponds to lower prediction error - signifying that these frames will have a higher impact in the overall *completion moment* prediction. When weakly-supervised, the attention scores are comparable to full-supervision though understandably softer attention is learnt.

5. Conclusion and Future Work

In this paper, we proposed a method to detect the com*pletion moment* in a variety of actions, suitable for both weakly-supervised and fully supervised sequences. In weak-supervision, video-level labels of completion or in*completion* are only required, for the same action. When a sufficient number of incomplete sequences is available during training, our approach, (1) learns discriminative features for frames pre- and post- completion, by propagating videolevel labels to individual frames, (2) learns temporal attention, to weight discriminative frame-level features, and then (3) accumulates evidence for completion, weighted by the learnt attention, from all frames to predict the completion moment, or identify the attempt as incomplete. We evaluated our approach on 16 actions (with full supervision) and 10 actions (with weak supervision), from 3 datasets. When weakly-supervised, learning attention significantly improved the results on all tested actions. Under full supervision, we outperform prior work [14] on 14 out of the 16 actions.

For future work, we aim to augment the temporal attention with within-frame spatial attention to learn image regions that are most discriminative for completion. We will also combine our soft attention with hard attention mechanisms, similar to [36]. Further, we will investigate *completion moment* detection from untrimmed sequences, which contain multiple action instances.

Acknowledgements

The 1st author wishes to thank the University of Bristol for partial funding of her studies. Public datasets were used in this work.

References

- J. Ba, V. Mnih, and K. Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014. 2
- [2] F. Becattini, T. Uricchio, L. Ballan, L. Seidenari, and A. Del Bimbo. Am I done? Predicting action progress in videos. arXiv preprint arXiv:1705.01781, 2017. 1, 2
- [3] J. Caicedo and S. Lazebnik. Semantic guidance of visual attention for localizing objects in scenes. In *ICCV*, 2015. 2
- [4] J. Carreira and A. Zisserman. Quo Vadis, Action recognition? A new model and the kinetics dataset. In *CVPR*, 2017.
- [5] Y. Chao, S. Vijayanarasimhan, B. Seybold, D. Ross, J. Deng, and R. Sukthankar. Rethinking the faster R-CNN architecture for temporal action localization. In *CVPR*, 2018. 2
- [6] L. Chen, H. Zhang, J. Xiao, L. Nie, J. Shao, W. Liu, and T. Chua. SCA-CNN: Spatial and channel-wise attention in convolutional networks for image captioning. In *CVPR*, 2017. 2
- [7] M. Denil, L. Bazzani, H. Larochelle, and N. de Freitas. Learning where to attend with deep architectures for image tracking. *Neural Computation*, 2012. 2
- [8] W. Du, Y. Wang, and Y. Qiao. Recurrent spatial-temporal attention network for action recognition in videos. *IEEE Transactions on Image Processing*, 2018. 2
- [9] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman. Temporal cycle-consistency learning. *arXiv preprint* arXiv:1904.07846, 2019. 2
- [10] R. Girdhar and D. Ramanan. Attentional pooling for action recognition. In NIPS, 2017. 2
- [11] R. Girdhar, D. Ramanan, A. Gupta, J. Sivic, and B. Russell. ActionVLAD: Learning spatio-temporal aggregation for action classification. In CVPR, 2017. 1
- [12] A. Haque, A. Alahi, and L. Fei-Fei. Recurrent attention models for depth-based person identification. In CVPR, 2016. 2
- [13] F. Heidarivincheh, M. Mirmehdi, and D. Damen. Beyond action recognition: Action completion in RGB-D data. In *BMVC*, 2016. 1, 2, 4
- [14] F. Heidarivincheh, M. Mirmehdi, and D. Damen. Action completion: a temporal model for moment detection. *BMVC*, 2018. 1, 2, 4, 5, 6, 8
- [15] M. Hoai and F. De la Torre. Max-margin early event detectors. *International Journal of Computer Vision (IJCV)*, 2014.
 2
- [16] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 1
- [17] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: A large video database for human motion recognition. In *ICCV*, 2011. 2, 4

- [18] I. Laptev. On space-time interest points. International Journal of Computer Vision (IJCV), 2005. 1
- [19] D. Li, T. Yao, L. Duan, T. Mei, and Y. Rui. Unified spatiotemporal attention networks for action recognition in videos. *IEEE Transactions on Multimedia*, 2019. 2, 3
- [20] S. Li, S. Bak, P. Carr, and X. Wang. Diversity regularized spatiotemporal attention for video-based person reidentification. In *CVPR*, 2018. 2
- [21] Z. Li, K. Gavrilyuk, E. Gavves, M. Jain, and C. Snoek. VideoLSTM convolves, attends and flows for action recognition. *Computer Vision and Image Understanding (CVIU)*, 2018. 2
- [22] S. Ma, L. Sigal, and S. Sclaroff. Learning activity progression in LSTMs for activity detection and early detection. In *CVPR*, 2016. 2
- [23] D. Moltisanti, S. Fidler, and D. Damen. Action Recognition from Single Timestamp Supervision in Untrimmed Videos. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
 1
- [24] D. Moltisanti, M. Wray, W. Mayol-Cuevas, and D. Damen. Trespassing the boundaries: Labeling temporal bounds for object interactions in egocentric video. In *ICCV*, 2017. 1
- [25] P. Nguyen, T. Liu, G. Prasad, and B. Han. Weakly supervised action localization by sparse temporal pooling network. In *CVPR*, 2018. 2, 3
- [26] S. Paul, S. Roy, and A. Roy-Chowdhury. W-TALC: Weaklysupervised temporal activity localization and classification. In *ECCV*, 2018. 2, 3
- [27] W. Pei, T. Baltrusaitis, D. Tax, and L. Morency. Temporal attention-gated model for robust sequence classification. In *CVPR*, 2017. 2
- [28] S. Sharma, R. Kiros, and R. Salakhutdinov. Action recognition using visual attention. arXiv preprint arXiv:1511.04119, 2015. 2
- [29] Z. Shou, D. Wang, and S. Chang. Temporal action localization in untrimmed videos via multi-stage CNNs. In *CVPR*, 2016. 2
- [30] G. Sigurdsson, O. Russakovsky, A. Farhadi, I. Laptev, and A. Gupta. Much ado about time: Exhaustive annotation of temporal data. *arXiv preprint arXiv:1607.07429*, 2016. 1
- [31] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In NIPS, 2014. 1
- [32] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In AAAI, 2017. 2
- [33] K. Soomro, A. Roshan Zamir, and M. Shah. A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402, 2012. 2, 4
- [34] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013. 1
- [35] L. Wang, Y. Qiao, and X. Tang. Motionlets: Mid-level 3D parts for human motion recognition. In CVPR, 2013. 1
- [36] L. Wang, Y. Xiong, D. Lin, and L. Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *CVPR*, 2017. 2, 3, 8
- [37] Y. Wang, S. Wang, J. Tang, N. O'Hare, Y. Chang, and B. Li. Hierarchical attention network for action recognition in videos. arXiv preprint arXiv:1607.06416, 2016. 2

- [38] H. Xu, A. Das, and K. Saenko. R-C3D: Region convolutional 3D network for temporal activity detection. In *ICCV*, 2017.
- [39] J. Xu, R. Zhao, F. Zhu, H. Wang, and W. Ouyang. Attentionaware compositional network for person re-identification. In *CVPR*, 2018. 2
- [40] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 2
- [41] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville. Describing videos by exploiting temporal structure. In *ICCV*, 2015. 2
- [42] 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)*, 2018. 2, 3
- [43] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. Endto-end learning of action detection from frame glimpses in videos. In CVPR, 2016. 2