

# Spatio-Temporal Event Segmentation for Wildlife Extended Videos

Ramy Mounir<sup>1((\Box)</sup>, Roman Gula<sup>2</sup>, Jörn Theuerkauf<sup>2</sup>, and Sudeep Sarkar<sup>1</sup>

<sup>1</sup> University of South Florida, Tampa, Florida 33647, USA {ramy, sarkar}@usf.edu
<sup>2</sup> Museum and Institute of Zoology, PAS, Warsaw, Poland {rgula,jtheuer}@miiz.eu

Abstract. Using offline training schemes, researchers have tackled the event segmentation problem by providing full or weak-supervision through manually annotated labels or self-supervised epoch-based training. Most works consider videos that are at most 10's of minutes long. We present a self-supervised perceptual prediction framework capable of temporal event segmentation by building stable representations of objects over time and demonstrate it on long videos, spanning several days at 25 FPS. The approach is deceptively simple but quite effective. We rely on predictions of high-level features computed by a standard deep learning backbone. For prediction, we use an LSTM, augmented with an attention mechanism, trained in a self-supervised manner using the prediction error. The self-learned attention maps effectively localize and track the eventrelated objects in each frame. The proposed approach does not require labels. It requires only a single pass through the video, with no separate training set. Given the lack of datasets of very long videos, we demonstrate our method on video from 10 d (254 h) of continuous wildlife monitoring data that we had collected with required permissions. We find that the approach is robust to various environmental conditions such as day/night conditions, rain, sharp shadows, and windy conditions. For the task of temporally locating events at the activity level, we had an 80% activity recall rate for one false activity detection every 50 min. We will make the dataset, which is the first of its kind, and the code available to the research community. Project page is available at https://ramymounir.com/ publications/EventSegmentation/.

Keywords: Self-supervised event segmentation  $\cdot$  Spatial object localization  $\cdot$  Streaming input

## 1 Introduction

One of the tasks involved in wild-life monitoring, or even in the video monitoring of other contexts, is detecting significant events in long videos, spanning several days. The goal is to flag temporal segments and highlight possible events in the video snippets flagged, i.e., spatial and temporal localization of possible events, such as bird {leaving,

**Supplementary Information** The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-11349-9\_5.

<sup>©</sup> The Author(s), under exclusive license to Springer Nature Switzerland AG 2022 B. Raman et al. (Eds.): CVIP 2021, CCIS 1568, pp. 48–59, 2022. https://doi.org/10.1007/978-3-031-11349-9\_5



**Fig. 1.** Plots of the two kinds of errors before, during, and after an activity: (top) feature prediction loss over the frames (bottom) motion weighted feature prediction loss over the frames. Errors for some selected frames are shown for both plots, overlaid with the corresponding attention map.

entering, building} a nest. One cannot rely on low-level features for this task as they may change due to environmental conditions. We need high-level features that are sufficient to capture object-level representations and a model to capture these features' temporal evolution over time. There are very few works in the literature that show performance on video spanning several days.

Event segmentation research has largely focused on offline epoch-based training methods which requires training the model on the entire training dataset prior to testing its performance. This poses a challenge for many real world applications, where the entire dataset is simply non-existent and has to be collected sequentially in a stream over time [24]. Our training scheme completely disregards datapoints after being processed by the network. Training and inference are done simultaneously, alleviating the need for epoch-based training in order to appeal to more practical applications and reduce training time.

Our framework follows key ideas from the perceptual prediction line of work in cognitive psychology [15, 34–36]. Research has shown that "event segmentation is an ongoing process in human perception, which helps form the basis of memory and learning". Humans can identify event boundaries, in a purely bottom up fashion, using a biological perceptual predictive model which predicts future perceptual states based on the current perceived sensory information. Experiments have shown that human perceptual system identifies event boundaries based on the appearance and motion cues in the video [22,26,33]. Our model implements this perceptual predictive framework and introduces a motion weighted loss function to allow for the localization and processing of motion cues.

Our approach uses a feature encoding network to transform low-level perceptual information to higher level feature representation. The model is trained to predict the future perceptual encoded input and signal an event if the prediction is significantly different from the future perceived features. The prediction signal also incorporates a higher level representation of the movement cues within frames. The intuition of our approach is presented in Fig. 1 for a "walk in and out" event. The error signal is used for temporal event segmentation, while an attention map segments a frame spatially.

**Novel contributions:** To the best of our knowledge, we are among the first to (1) introduce the attention-based mechanism to temporal event segmentation models, allowing the model to localize the event in each processed frame, in a purely self-supervised manner, without the need for labels or training data; (2) introduce the idea of motion weighted loss function to stabilize the attention maps that works even when the object of interest does not move; and (3) evaluate and report the performance of temporal segmentation on a remarkably long dataset (over ten days of continuous wildlife monitoring).

## 2 Relevant Work

**Supervised Temporal Event Segmentation** uses direct labelling (of frames) to segment videos into smaller constituent events. Fully supervised models are heavily dependent on vast amount of training data to achieve good segmentation results. Different model variations and approaches have been tested, such as using an encoder-decoder temporal convolutional network (ED-TCN) [13], or a spatiotemporal CNN model [14]. To alleviate the need for expensive direct labelling, weakly supervised approaches [7,9,12,23] have emerged with an attempt to use metadata (such as captions or narrations) to guide the training process without the need for explicit training labels [4, 18]. However, such metadata are not always available as part of the dataset, which makes weakly supervised approaches inapplicable to most practical applications.

**Self-supervised Temporal Event Segmentation** attempts to completely eliminate the need for annotations [20,25]. Many approaches rely heavily on higher level features clustering of frames to sub-activities [6,28]. The performance of the clustering algorithms in unsupervised event segmentation is proportional to the performance of the embedding/encoding model that transforms frames to higher level feature representations. Clustering algorithms can be highly computationally expensive depending on the number of frames to be clustered. Recent work [2] uses a self-supervised perceptual predictive model to detect event boundaries; we improve upon this model to include attention unit, which helps the model focus on event-causing objects. Other work [19] uses a self-supervised perceptual prediction model that is refined over significant amount of reinforcement learning iterations.

**Frame Predictive Models** have attempted to provide accurate predictions of the next frame in a sequence [10, 16, 21, 29, 30]; however, these models are focusing on predicting future frames in raw pixel format. Such models may generate a prediction loss that only captures frame motion difference with limited understanding of higher level features that constitutes event boundaries.

Attention Units have been applied to image captioning [31], and natural language processing [5,8,17,27,32] fully supervised applications. Attention is used to expose different temporal - or spatial - segments of the input to the decoding LSTM at every time step using fully supervised model architectures. We use attention in a slightly different form, where the LSTM is decoded only once (per input frame) to predict future features and uses attention weighted input to do so. Unlike [5,8,17,27,31,32], our attention weights and biases are trained using an unsupervised loss functions.

Recent work [3] has used the prediction loss, with the assistance of region proposal networks (RPNs) and multi-layer LSTM units, to localize actions. We eliminate the need for RPNs and multi-layer LSTM units by extracting Bahdanau [5] attention weights prior to the LSTM prediction layer, which allows our model to localize objects of interest, even when stationary. From our experiments, we found out that prediction loss attention tends to fade away as moving objects become stationary, which makes its attention map more similar to results extracted from background subtraction or optical flow. In contrast, our model proves to be successful in attending to moving and stationary objects despite variations in environmental conditions, such as moving shadows and lighting changes, as presented in the supplementary videos.

## 3 Methodology

The proposed framework is inspired by the works of Zacks *et al.* on perceptual prediction for events segmentation [34]. The proposed architecture, summarised in Fig. 2, can be divided into several individual components. In this section, we explain the role of each component starting by the encoder network and attention unit in Sect. 3.1 and 3.2, followed by a discussion on the recurrent predictive layer in Sect. 3.3. We conclude by introducing the loss functions (Sect. 3.4) used for self-supervised learning as well as the adaptive thresholding function (Sect. 3.5). Full pseudocode is provided in the appendix.

## 3.1 Input Encoding

The raw input images are transformed from pixel space into a higher level feature space by utilizing an encoder (CNN) model. This encoded feature representation allows the network to extract features of higher importance to the task being learned. We denote the output of the CNN layers by  $I'_t = f(I_t, \theta_e)$  where  $\theta_e$  is the learnable weights and biases parameters and  $I_t$  is the input image.

## 3.2 Attention Unit

In this framework, we utilize Bahdanau attention [5] to spatially localize the event in each processed frame. The attention unit receives as an input the encoded features and



**Fig. 2.** The architecture of the self-learning, perceptual prediction algorithm. Input frames from each time instant are encoded into high-level features using a deep-learning stack, followed by an attention overlay that is based on inputs from previous time instant, which is input to an LSTM. The training loss is composed based on the predicted and computed features from current and next frames.

outputs a set of attention weights  $(A_t)$  with dimensions  $N \times N \times 1$ . The hidden feature vectors  $(h_{t-1})$  from the prediction layer of the previous time step is used to calculate the output set of weights using Eq. 1, expressed visually in Fig. 2.

$$A_t = \gamma \left( FC(\varphi(FC(h_{t-1}) + FC(I'_t))) \right)$$
(1)

where  $\varphi$  represents hyperbolic tangent (tanh) function, and  $\gamma$  represents a softmax function. The weights  $(A_t)$  are then multiplied by the encoded input feature vectors  $(I'_t)$  to generate the masked feature vectors  $(I''_t)$ .

## 3.3 Future Prediction Layer

The process of future prediction requires a layer capable of storing a flexible internal state (event model) of the previous frames. For this purpose, we use a recurrent layer, specifically long-short term memory cell (LSTM) [11], which is designed to



Fig. 3. Qualitative Results. (Top): Samples of images from the Kagu bird wildlife monitoring dataset. (Bottom): Corresponding attention maps overlayed on the input image

output a future prediction based on the current input and a feature representation of the internal state. More formally, the LSTM cell can be described using the function  $h_t = g(I''_t, W_{lstm}, h_{t-1})$ , where  $h_t$  and  $h_{t-1}$  are the output hidden state and previous hidden state respectively,  $I''_t$  the attention-masked input features at time step t and  $W_{lstm}$  is a set of weights and biases vectors controlling the internal state of the LSTM. The input to the LSTM can be formulated as:

$$FC(y_{t-1} \oplus I_t'') \tag{2}$$

where  $I''_t$  is the masked encoded input feature vector and  $h_{t-1}$  is the hidden state from the previous time step. The symbol  $\oplus$  represents vectors concatenation.

## 3.4 Loss Function

The perceptual prediction model aims to train a model capable of predicting the feature vectors of the next time step. We define two different loss functions, prediction loss and motion weighted loss.

*Prediction Loss.* This function is defined as the L2 Euclidean distance loss between the output prediction  $y_t$  and the next frame encoded feature vectors  $I'_{t+1}$ .

$$e_t = ||(I'_{t+1} - y_t)||^2 \tag{3}$$

*Motion Weighted Loss.* This function aims to extract the motion related feature vectors from two consecutive frames to generate a motion dependent mask, which is applied to the prediction loss. The motion weighted loss function allows the network to benefit from motion information in higher level feature space rather than pixel space. This function is formally defined as:

$$e_t = ||(I'_{t+1} - y_t)^{\odot 2} \odot (I'_{t+1} - I'_t)^{\odot 2}||^2$$
(4)

where  $\odot$  denotes Hadamard (element-wise) operation.

#### 3.5 Error Gate

The error gating function receives, as an input, the error signal defined in Sect. 3.4, and applies a thresholding function to classify each frame. In this framework, we define two types of error gating functions. A simple threshold function  $f(e_t, \psi)$  and an adaptive threshold function  $f(\{e_{t-m} \dots e_t\}, \psi)$ . Equation 5 formally defines the smoothing function for the adaptive error gating implementation. Both error gating functions use Eq. 6 to threshold the error signal. Equations 5 and 6 apply the smoothing function to the full loss signal for analyses purposes; however, the convolution operation can be reduced to element-wise multiplication to calculate a single smoothed value at time step t.

$$e = \{e_{t-m} \dots e_t\} \in \mathbb{R}^m$$

$$e_s = e - [e \circledast \left[\left\{\frac{1}{n} \dots \frac{1}{n}\right\} \in \mathbb{R}^n\right]\right]\right\} n < m$$
(5)

$$f(e_s(t)) = \begin{cases} 1, & \text{if } e_s(t) \ge \psi \\ 0, & \text{otherwise} \end{cases}$$
(6)

where  $\circledast$  represents a 1D convolution operation.

## **4** Experimental Evaluation

In this section, we present the results of our experiments for our approach defined in Sect. 3. We begin by defining the wildlife extended video dataset used for testing, followed by explaining the evaluation metrics used to quantify performance. We discuss the model variations evaluated and conclude by presenting quantitative and qualitative results in Sects. 4.4 and 4.5.

#### 4.1 Dataset

We analyze the performance of our model on a wildlife monitoring dataset. The dataset consists of 10 d (254 h) continuous monitoring of a nest of the Kagu, a flightless bird of New Caledonia. The labels include four unique bird activities, {feeding the chick, incubation/brooding, nest building while sitting on the nest, nest building around the nest}. Start and end times for each instance of these activities are provided with the annotations. We modified the annotations to include walk in and walk out events representing the transitioning events from an empty nest to incubation and vice versa. Our approach can flag the nest building (on and around the nest), feeding the chick, walk in and out events. Other events based on climate, time of day, lighting conditions are ignored by our segmentation network. Figure 3 (Top) shows a sample of images from the dataset.



**Fig. 4.** Frame-level event segmentation ROCs when activities are detected based on simple thresholding of the prediction and motion weighted loss signals. Plots are shown for different ablation studies.



Fig. 5. Activity-level event segmentation ROCs when activities are detected based on simple thresholding of the prediction and motion weighted loss signals. Plots are shown for different ablation studies.

#### 4.2 Evaluation Metrics

We provide quantitative ROC results for both frame level (Fig. 4) and activity level (Figs. 5 and 6) event segmentation. Frame window size ( $\psi$ ) is defined as the maximum joining window size between events; a high  $\psi$  value can causes separate detected events to merge, which decreases the overall performance.

*Frame Level.* The recall value in frame level ROC is calculated as the ratio of true positive frames (event present) to the number of positive frames in the annotations dataset, while the false positive rate is expressed as ratio of the false positive frames to the total number of negative frames (event not present) in the annotation dataset. Threshold value ( $\phi$ ) is varied to obtain a single ROC line, while varying the frame window size ( $\psi$ ) results in a different ROC line.

*Activity Level.* The Hungarian matching (Munkres assignment) algorithm is utilized to achieve one to one mapping between the ground truth labeled events and the detected events. Recall is defined as ratio of the number of correctly detected events (overlapping frames) to the total number of groundtruth events. For the activity level ROC chart, the recall values are plotted against the false positive rate per minute, defined as the ratio of



**Fig. 6.** Activity-level event segmentation ROCs when activities are detected based on adaptive thresholding of the prediction and motion weighted loss signals. Plots are shown for different ablation studies.

the total number of false positive detected events to the total duration of the dataset in minutes. The false positive rate per minute evaluation metric is also used in the ActEV TRECVID challenge [1]. Frame window size value ( $\psi$ ) is varied to obtain a single ROC line, while varying the threshold value ( $\phi$ ) results in a different ROC line.

#### 4.3 Ablative Studies

Different variations of our framework (Sect. 3) have been evaluated to quantify the effect of individual components on the overall performance. In our experiments, we tested the base model, which trains the perceptual prediction framework - including attention unit - using the prediction loss function for backpropagation of the error signal. We refer to the base model as LSTM+ATTN. We also experimented with the effect of removing the attention unit, from the model architecture, on the overall segmentation performance; results of this variation are reported under the model name LSTM. Further testing includes using the motion weighted loss for backpropagation of the error signal. We refer to the motion weighted model as LSTM+ATTN+MW. Each of the models has been tested extensively; results are reported in Sects. 4.4 and 4.5, as well as visually expressed in Figs. 3, 4, 5 and 6.

Comparing the results shown in Figs. 5 and 6 indicate a significant increase of overall performance when using an adaptive threshold for loss signal gating. The efficacy of adaptive thresholding is evident when applied to activity level event segmentation. Comparing the results (LSTM & LSTM+ATTN) show that the model can effectively generate attention maps (Sect. 4.5) without impacting the segmentation performance.

## 4.4 Quantitative Evaluation

We tested three different models, *LSTM*, *LSTM*+*ATTN*, and *LSTM*+*ATTN*+*MW*, for frame level and activity level event segmentation. Simple and adaptive gating functions (Sect. 3.5), were applied to prediction and motion weighted loss signals (Sect. 3.4) for frame level and activity level experiments. For each model we vary parameters such as the threshold value  $\psi$  and the frame window size  $\phi$  to achieve the ROC charts presented in Figs. 4, 5 and 6.

It is to be noted that thresholding a loss signal does not necessarily imply that the model was trained to minimize this particular signal. In other words, loss functions used for backpropagating the error to the models' learnable parameters are identified only in the model name (Sect. 4.3); however, thresholding experiments have been conducted on different types of loss signals, regardless of the backpropagating loss function used for training.

The best performing model, for frame level segmentation (*LSTM*+*ATTN*, $\psi = 1000$ ) is capable of achieving {40%, 60%, 70%} frame recall value at {5%, 10%, 20%} frame false positive rate respectively. Activity level segmentation can recall {80%, 90%, 95%} of the activities at {0.02, 0.1, 0.2} activity false positive rate per minute, respectively, for the model (*LSTM*+*ATTN*+*MW*, $\phi = 0.0001$ ) as presented in Fig. 6. A 0.02 false positive activity rate per minute can also be interpreted as one false activity detection every 50 min of training (for detecting 80% of the groundtruth activities).

## 4.5 Qualitative Evaluation

A sample of the qualitative attention results is presented in Fig. 3. The attention mask, extracted from the model, has been trained to track the event in all processed frames. Our results show that the events are tracked and localized in various lighting (shadows, day/night) and occlusion conditions. Attention has also learned to indefinitely focus on the bird regardless of its motion state (stationary/Non-stationary), which indicates that the model has acquired a high-level temporal understanding of the events in the scene and learned the underlying structure of the bird. Supplementary results<sup>1</sup> display a timelapse of attention weighted frames during illumination changes and moving shadows. We also provide a supplementary video showing the prediction loss signal, motion weighted loss signal and attention mask during a walk in and out event (summarized in Fig. 1). Additional qualitative results are provided in the appendix.

## 5 Conclusion

We demonstrate a self-supervised approach to temporal event segmentation. Our framework can effectively segment a long sequence of activities (video) into a set of individual events. We introduce a novel approach to extract attention results from unsupervised temporal event segmentation network. Gating the loss signal with different threshold values can result in segmentation at different granularities. Quantitative and qualitative results are presented in the form of ROC charts and attention weighted frames. Our results demonstrate the effectiveness of our approach in understanding the higher level spatio-temporal features required for practical temporal event segmentation.

Acknowledgment. This research was supported in part by the US National Science Foundation grants CNS 1513126 and IIS 1956050. The bird video dataset used in this paper was made possible through funding from the Polish National Science Centre (grant NCN 2011/01/M/NZ8/03344 and 2018/29/B/NZ8/02312). Province Sud (New Caledonia) issued all permits - from 2002 to 2020 - required for data collection.

<sup>&</sup>lt;sup>1</sup> Available at https://ramymounir.com/publications/EventSegmentation/.

## References

- 1. Actev: Activities in extended video. https://actev.nist.gov/
- Aakur, S.N., Sarkar, S.: A perceptual prediction framework for self supervised event segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1197–1206 (2019)
- 3. Aakur, S.N., Sarkar, S.: Action localization through continual predictive learning (2020). arXiv preprint arXiv:2003.12185
- Alayrac, J.B., Bojanowski, P., Agrawal, N., Sivic, J., Laptev, I., Lacoste-Julien, S.: Unsupervised learning from narrated instruction videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4575–4583 (2016)
- Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate (2014). arXiv preprint arXiv:1409.0473
- Bhatnagar, B.L., Singh, S., Arora, C., Jawahar, C.: CVIT K: Unsupervised learning of deep feature representation for clustering egocentric actions. In: IJCAI, pp. 1447–1453 (2017)
- Bojanowski, P., et al.: Weakly supervised action labeling in videos under ordering constraints. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 628–643. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1\_41
- 8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: pre-training of deep bidirectional transformers for language understanding (2018). arXiv preprint arXiv:1810.04805
- Ding, L., Xu, C.: Weakly-supervised action segmentation with iterative soft boundary assignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6508–6516 (2018)
- Finn, C., Goodfellow, I., Levine, S.: Unsupervised learning for physical interaction through video prediction. In: Advances in Neural Information Processing Systems, pp. 64–72 (2016)
- 11. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Comput. **9**(8), 1735–1780 (1997)
- Huang, D.A., Fei-Fei, L., Niebles, J.C.: Connectionist temporal modeling for weakly supervised action labeling. In: European Conference on Computer Vision, pp. 137–153. Springer (2016). https://doi.org/10.48550/arXiv.1607.08584
- Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 156–165 (2017)
- Lea, C., Reiter, A., Vidal, R., Hager, G.D.: Segmental spatiotemporal CNNs for fine-grained action segmentation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9907, pp. 36–52. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9\_3
- Loschky, L.C., Larson, A.M., Smith, T.J., Magliano, J.P.: The scene perception & event comprehension theory (spect) applied to visual narratives. Top. Cogn. Sci. 12(1), 311–351 (2020)
- 16. Lotter, W., Kreiman, G., Cox, D.: Deep predictive coding networks for video prediction and unsupervised learning (2016). arXiv preprint arXiv:1605.08104
- 17. Luong, M.T., Pham, H., Manning, C.D.: Effective approaches to attention-based neural machine translation (2015). arXiv preprint arXiv:1508.04025
- Malmaud, J., Huang, J., Rathod, V., Johnston, N., Rabinovich, A., Murphy, K.: What's cookin'? interpreting cooking videos using text, speech and vision (2015). arXiv preprint arXiv:1503.01558
- 19. Metcalf, K., Leake, D.: Modelling unsupervised event segmentation: learning event boundaries from prediction errors. In: CogSci. (2017)
- Garcia del Molino, A., Lim, J.H., Tan, A.H.: Predicting visual context for unsupervised event segmentation in continuous photo-streams. In: Proceedings of the 26th ACM International Conference on Multimedia, pp. 10–17 (2018)

- 21. Qiu, J., Huang, G., Lee, T.S.: A neurally-inspired hierarchical prediction network for spatiotemporal sequence learning and prediction (2019). arXiv preprint arXiv:1901.09002
- Radvansky, G.A., Krawietz, S.A., Tamplin, A.K.: Walking through doorways causes forgetting: further explorations. Quart. J. Exper. Psychol. 64(8), 1632–1645 (2011)
- Richard, A., Kuehne, H., Gall, J.: Weakly supervised action learning with RNN based fineto-coarse modeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 754–763 (2017)
- 24. Sahoo, D., Pham, Q., Lu, J., Hoi, S.C.: Online deep learning: Learning deep neural networks on the fly (2017). arXiv preprint arXiv:1711.03705
- Sener, F., Yao, A.: Unsupervised learning and segmentation of complex activities from video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8368–8376 (2018)
- Speer, N.K., Swallow, K.M., Zacks, J.M.: Activation of human motion processing areas during event perception. Cogn. Aff. Behav. Neurosci. 3(4), 335–345 (2003)
- Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)
- 28. VidalMata, R.G., Scheirer, W.J., Kuehne, H.: Joint visual-temporal embedding for unsupervised learning of actions in untrimmed sequences (2020). arXiv preprint arXiv:2001.11122
- 29. Wang, Y., Gao, Z., Long, M., Wang, J., Yu, P.S.: Predrnn++: towards a resolution of the deep-in-time dilemma in spatiotemporal predictive learning (2018). arXiv preprint arXiv:1804.06300
- Wichers, N., Villegas, R., Erhan, D., Lee, H.: Hierarchical long-term video prediction without supervision (2018). arXiv preprint arXiv:1806.04768
- Xu, K., et al.: Show, attend and tell: neural image caption generation with visual attention. In: International Conference on Machine Learning, pp. 2048–2057 (2015)
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: Xlnet: generalized autoregressive pretraining for language understanding. In: Advances in Neural Information Processing Systems, pp. 5754–5764 (2019)
- Zacks, J.M.: Using movement and intentions to understand simple events. Cogn. Sci. 28(6), 979–1008 (2004)
- 34. Zacks, J.M., Swallow, K.M.: Event segmentation. Curr. Dir. Psychol. Sci. 16(2), 80-84 (2007)
- Zacks, J.M., Tversky, B.: Event structure in perception and conception. Psychol. Bull. 127(1), 3 (2001)
- 36. Zacks, J.M., Tversky, B., Iyer, G.: Perceiving, remembering, and communicating structure in events. J. Exp. Psychol. Gen. **130**(1), 29 (2001)