

What Would You Expect? Anticipating Egocentric Actions With Rolling-Unrolling LSTMs and Modality Attention

Antonino Furnari Giovanni Maria Farinella

University of Catania - Department of Mathematics and Computer Science

<http://iplab.dmi.unict.it/fpv/> - {furnari,gfarinella}@dmi.unict.it

Abstract

Egocentric action anticipation consists in understanding which objects the camera wearer will interact with in the near future and which actions they will perform. We tackle the problem proposing an architecture able to anticipate actions at multiple temporal scales using two LSTMs to 1) summarize the past, and 2) formulate predictions about the future. The input video is processed considering three complimentary modalities: appearance (RGB), motion (optical flow) and objects (object-based features). Modality-specific predictions are fused using a novel Modality Attention (MATT) mechanism which learns to weigh modalities in an adaptive fashion. Extensive evaluations on two large-scale benchmark datasets show that our method outperforms prior art by up to +7% on the challenging EPIC-Kitchens dataset including more than 2500 actions, and generalizes to EGTEA Gaze+. Our approach is also shown to generalize to the tasks of early action recognition and action recognition. Our method is ranked first in the public leaderboard of the EPIC-Kitchens egocentric action anticipation challenge 2019. Please see the project web page for code and additional details: <http://iplab.dmi.unict.it/rulstm>.

1. Introduction

Anticipating the near future is a natural task for humans and a fundamental one for intelligent systems when it is necessary to react before an action is completed (e.g., to anticipate a pedestrian crossing the street from an autonomous vehicle [9]) or even before it starts (e.g., to notify a user who is performing the wrong action in a known workflow [49]). Additionally, tasks such as action anticipation [18, 29, 54] and early action recognition [2, 9, 36] pose a series of key challenges from a computational perspective. Indeed, methods addressing these tasks need to model the relationships between past, future events and incomplete observations. First Person (Egocentric) Vision [26], in particular, offers an interesting scenario to study anticipation problems. On one hand, whilst being a natural task for humans, anticipating

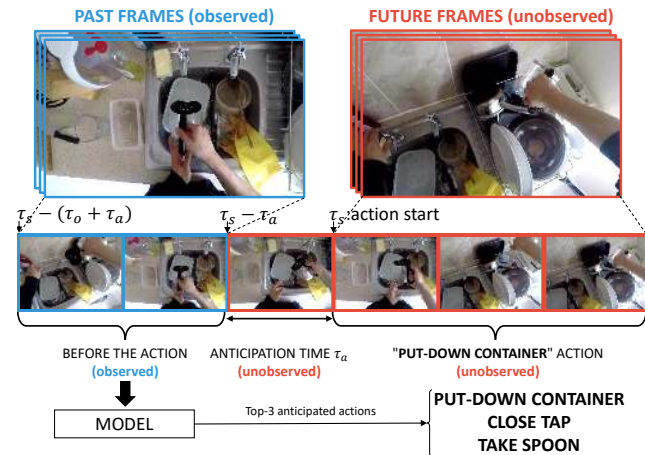


Figure 1. Egocentric Action Anticipation. See text for notation.

the future from egocentric video is computationally challenging due to the ability of wearable cameras to acquire long videos of complex activities involving many objects and actions performed by a user from their unique point of view. On the other hand, investigating these tasks is fundamental for the construction of intelligent wearable systems able to anticipate the user’s goal and assist them [26].

In this paper, we address the problem of egocentric action anticipation. As defined in [8] and illustrated in Figure 1, the task consists in recognizing an action starting at time τ_s by observing a video segment *preceding the action* starting at time $\tau_s - (\tau_o + \tau_a)$ and ending at time $\tau_s - \tau_a$, where the “observation time” τ_o indicates the length of the observed segment, whereas the “anticipation time” τ_a denotes how many seconds in advance actions are to be anticipated. While action anticipation has been investigated in classic third person vision [1, 18, 22, 29, 54, 25], less attention has been devoted to the egocentric scenario [8, 16, 43].

We observe that egocentric action anticipation methods need to address two sub-tasks: 1) summarizing what has been observed in the past (e.g., “a container has been washed” in the observed segment in Figure 1), and 2) making hypotheses about what will happen in the future (e.g.,

“put-down container”, “close tap”, “take spoon” in Figure 1). While previous approaches attempted to address these two sub-tasks jointly [2, 8, 36, 54], our method disentangles them using two separate LSTMs. The “Rolling” LSTM (R-LSTM) is responsible for continuously encoding streaming observations to summarize the past. When the method is required to anticipate actions, the “Unrolling” LSTM (U-LSTM) takes over the current hidden and cell states of the R-LSTM and makes predictions about the future. Differently from previous works which considered a fixed anticipation time [8, 16, 54], the proposed method can anticipate an action at multiple anticipation times, e.g., 2s, 1.5s, 1s and 0.5s before it occurs. The network is pre-trained with a novel “Sequence Completion Pre-training” (SCP) technique, which encourages the disentanglement of the two sub-tasks. To take advantage of the complementary nature of different input modalities, the proposed Rolling-Unrolling LSTM (RU) processes spatial observations (RGB frames), motion (optical flow), as well as object-based features. Multimodal predictions are fused with a novel “Modality ATtention” mechanism (MATT), which adaptively estimates optimal fusion weights for each modality by considering the outputs of the modality-specific R-LSTM components. Experiments on two large-scale datasets of egocentric videos, EPIC-KTICHENS [8] and EGTEA Gaze+ [33], show that the proposed method outperforms several state-of-the-art approaches and baselines in the task of egocentric action anticipation and generalizes to the tasks of early action recognition and action recognition.

The contributions of this work are as follows: 1) we are the first to systematically investigate the problem of egocentric action anticipation within the framework of the challenge proposed in [8]; 2) our investigation benchmarks popular ideas and approaches to action anticipation and leads to the definition of RU, an architecture able to anticipate egocentric actions at multiple temporal scales; 3) the proposed model is shown to benefit from two techniques specific to the investigated problem, i.e., i) “Sequence Completion Pre-training” and ii) adaptive fusion of multi-modal predictions through Modality ATtention; 4) extensive evaluations highlight the limits of previous approaches and show significant improvements of the proposed method over the state of the art. To support future research, the code implementing the proposed method will be released upon publication.

2. Related Work

Action Recognition Our work is related to previous research on action recognition from third person vision [7, 13, 14, 27, 31, 32, 45, 53, 55, 56, 57, 62] and first person vision [11, 12, 33, 34, 35, 40, 44, 48, 50, 51, 52]. Specifically, we build on previous ideas investigated in the context of action recognition such as the use of multiple modalities

for video analysis [45], the use of Temporal Segment Networks [57] as a principled way to train CNNs for action recognition, as well as the explicit encoding of object-based features [11, 35, 40, 47, 52] to analyze egocentric video. However, in contrast with the aforementioned works, we address the problem of egocentric action *anticipation* and show that approaches designed for action recognition, such as TSN [57] and early/late fusion to merge multi-modal predictions [45] are not directly applicable to the problem of egocentric action anticipation.

Early Action Recognition in Third Person Vision Early action recognition consists in recognizing an ongoing action as early as possible from partial observations [9]. The problem has been widely investigated in the domain of third person vision [2, 4, 6, 9, 19, 20, 21, 36, 42]. Differently from these works, we address the task of *anticipating* actions from egocentric video, i.e., the action should be recognized before it starts, hence it cannot even be partially observed at the time of prediction. Given the similarity between early recognition and anticipation, we consider and evaluate some ideas investigated in the context of early action recognition, such as the use of LSTMs to process streaming video [2, 19, 36], and the use of dedicated loss functions [36]. Moreover, we show that the proposed architecture also generalizes to the problem of early egocentric action recognition, achieving state-of-the-art performances.

Action Anticipation in Third Person Vision Action anticipation is the task of predicting an action *before* it occurs [18]. Previous works investigated different forms of action and activity anticipation from third person video [1, 15, 18, 22, 24, 25, 28, 29, 30, 37, 54, 59]. While we consider the problem of action anticipation from egocentric visual data, our work builds on some ideas explored in past works such as the use of LSTMs to anticipate actions [1, 18, 25], the use of the encoder-decoder framework to encode past observations and produce hypotheses of future actions [18], and the use of object specific features [37] to determine which objects are present in the scene. Additionally, we show that other approaches, such as the direct regression of future representations [18, 54], do not achieve satisfactory performance in the considered scenario.

Anticipation in First Person Vision Past works on anticipation from egocentric visual data have investigated different problems and considered different evaluation frameworks [5, 10, 17, 39, 41, 43, 46, 49, 60, 63]. Instead, we tackle the egocentric action anticipation challenge recently proposed in [8], which has been little investigated so far [16]. While a direct comparison of the proposed approach with the aforementioned works is unfeasible due to the lack of a common framework, our method incorporates some ideas from past approaches, such as the analysis of past actions [43] and the detection of the objects present in




Figure 2. Video processing scheme of the proposed method with $S_{enc} = 6$ and $S_{ant} = 8$.

the scene to infer future actions [17].

3. Proposed Approach

Processing Strategy Figure 2 illustrates the processing strategy adopted by the proposed method. The video is processed in an on-line fashion, with a short video snippet V_t consumed every α seconds, where t indexes the current time-step. Specifically, an action occurring at time τ_s is anticipated by processing a video segment of length l starting at time $\tau_s - l - \alpha$ and ending at time $\tau_s - \alpha$. The input video ends at time $\tau_s - \alpha$ as our method aims at anticipating actions *at least* α seconds before they occur. The processing is performed in two stages: an “encoding” stage, which is carried out for S_{enc} time-steps, and an “anticipation” stage, which is carried out for S_{ant} time-steps. In the encoding stage, the model summarizes the semantic content of the S_{enc} input video snippets without producing any prediction, whereas in the anticipation stage the model continues to encode the semantics of the S_{ant} input video snippets and outputs S_{ant} action scores s_t which can be used to perform action anticipation. This scheme effectively allows to formulate S_{ant} predictions for a single action at multiple anticipation times. In our experiments, we set $\alpha = 0.25s$, $S_{enc} = 6$ and $S_{ant} = 8$. In these settings, the model analyzes video segments of length $l = \alpha(S_{enc} + S_{ant}) = 3.5s$ and outputs 8 predictions at the following anticipation times: $\tau_a \in \{2s, 1.75s, 1.5s, 1.25s, 1s, 0.75s, 0.5s, 0.25s\}$. It should be noted that, since the predictions are produced *while* processing the video, at time step t the effective observation time will be equal to $\alpha \cdot t$. Hence, the 8 predictions are performed at the following effective observation times: $\tau_o \in \{1.75s, 2s, 2.25s, 2.5s, 2.75s, 3s, 3.25s, 3.5s\}$. Our formulation generalizes the one proposed in [8], which is illustrated in Figure 1. For instance, at time-step $t = 11$, our model will anticipate actions with an effective observation time equal to $\tau_o = \alpha \cdot t = 2.75s$ and an anticipation time equal to $\tau_a = \alpha(S_{ant} + S_{enc} + 1 - t) = 1s$, as in [8].

Rolling-Unrolling LSTM The proposed method uses two separate LSTMs to encode past observations and formulate predictions about the future. Following previous literature [45], we include multiple identical branches which

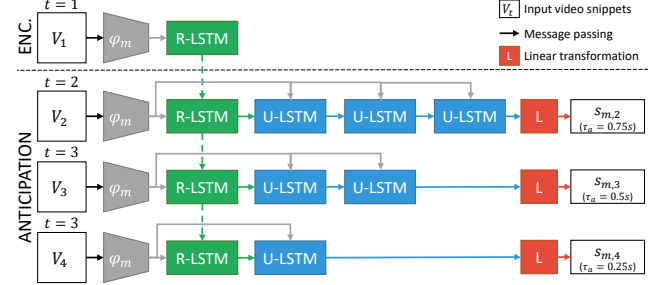


Figure 3. Example of RU modality-specific branch with $S_{enc} = 1$ and $S_{ant} = 3$.

analyze the video according to different modalities. Specifically, at each time-step t , the input video snippet V_t is represented using different modality-specific representation functions $\varphi_1, \dots, \varphi_M$ depending on learnable parameters $\theta^{\varphi_1}, \dots, \theta^{\varphi_M}$. This process allows to obtain the modality-specific feature vectors $f_{1,t} = \varphi_1(V_t), \dots, f_{M,t} = \varphi_M(V_t)$, where M is the total number of modalities (i.e., the total number of branches in our architecture), and $f_{m,t}$ is the feature vector computed at time-step t for the modality m . The feature vector $f_{m,t}$ is fed to the m^{th} branch of the architecture. While our model can easily incorporate different modalities, in this work we consider $M = 3$ modalities, i.e., RGB frames (spatial branch), optical flow (motion branch) and object-based features (object branch).

Figure 3 illustrates the processing taking place in a single branch m of the proposed RU model. For illustration purposes only, the figure shows an example in which $S_{enc} = 1$ and $S_{ant} = 3$. At time step t , the feature vector $f_{m,t}$ is fed to the Rolling LSTM (R-LSTM), which encodes its semantic content recursively, as follows:

$$(h_{m,t}^R, c_{m,t}^R) = LSTM_{\theta_m^R}(f_{m,t}, h_{m,t-1}^R, c_{m,t-1}^R) \quad (1)$$

where $LSTM_{\theta_m^R}$ denotes the R-LSTM of branch m , depending on the learnable parameters θ_m^R , whereas $h_{m,t}^R$ and $c_{m,t}^R$ are the hidden and cell states computed at time t in the modality m . The initial hidden and cell states of the R-LSTM are initialized with zeros: $h_{m,0}^R = \mathbf{0}$, $c_{m,0}^R = \mathbf{0}$.

In the anticipation stage, at time step t , the Unrolling LSTM (U-LSTM) is used to make predictions about the future. The U-LSTM takes over the hidden and cell vectors of the R-LSTM at the current time-step (i.e., $h_{m,t}^R$ and $c_{m,t}^R$) and iterates over the representation of the current video snippet $f_{m,t}$ for a number of times n_t equal to the number of time-steps required to reach the beginning of the action, i.e., $n_t = S_{ant} + S_{enc} - t + 1$. Hidden and cell states of the U-LSTM are computed as follows at the j^{th} iteration:

$$(h_{m,j}^U, c_{m,j}^U) = LSTM_{\theta_m^U}(f_{m,t}, h_{m,j-1}^U, c_{m,j-1}^U) \quad (2)$$

where $LSTM_{\theta_m^U}$ is the U-LSTM of branch m , depending on the learnable parameters θ_m^U , and $h_{m,t}^U, c_{m,t}^U$ are the hidden and cell states computed at iteration j for the modality

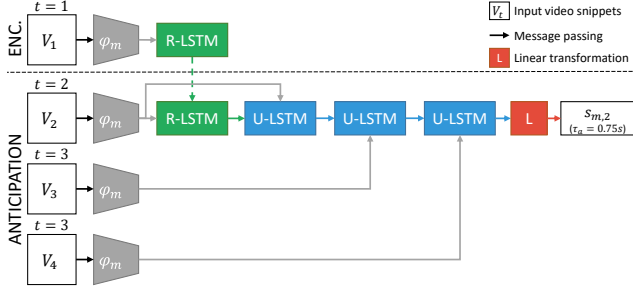


Figure 4. Example of connection scheme used during SCP for time-step $t = 2$.

m . The initial hidden and cell states of the U-LSTM are initialized from the current hidden and cell states computed by the R-LSTM: $h_{m,0}^U = h_{m,t}^R$, $c_{m,0}^U = c_{m,t}^R$. Note that the input $f_{m,t}$ of the U-LSTM does not depend on j (see eq. (2)) because it is fixed during the “unrolling” procedure. The main rationale of “unrolling” the U-LSTM for a different number of times at each time-step is to encourage it to differentiate predictions at different anticipation times.

Modality-specific action scores $s_{m,t}$ are computed at time-step t by processing the last hidden vector of the U-LSTM with a linear transformation with learnable parameters θ_m^W and θ_m^b : $s_{m,t} = \theta_m^W h_{m,n_t}^U + \theta_m^b$.

Sequence Completion Pre-Training (SCP) The two LSTMs composing the RU architecture are designed to address two specific sub-tasks: the R-LSTM is responsible for encoding past observations and summarizing what has happened up to a given time-step, whereas the U-LSTM focuses on anticipating future actions conditioned on the hidden and cell vectors of the R-LSTM. To encourage the two LSTMs to specialize on the two different sub-tasks, we propose to train the architecture using a novel Sequence Completion Pre-training (SCP) procedure. During SCP, the connections of the U-LSTM are modified to allow it to process future representations, rather than iterating on the current one. In practice, the U-LSTM hidden and cell states are computed as follows during SCP:

$$(h_{m,j}^U, c_{m,j}^U) = LSTM_{\theta_m^U}(f_{m,t+j-1}, h_{m,j-1}^U, c_{m,j-1}^U) \quad (3)$$

where the input representations $f_{m,t+j-1}$ are sampled from future time-steps $t+j-1$. Figure 4 illustrates an example of the connection scheme used during SCP for time-step $t = 2$. The main goal of pre-training the RU with SCP is to allow the R-LSTM to focus on summarizing past representations without trying to anticipate the future.

Modality ATTention (MATT) Coherently with past work on egocentric action anticipation [8], we found it sub-optimal to fuse multi-modal predictions with classic approaches such as early and late fusion. This is probably due to the fact that, when anticipating egocentric actions,

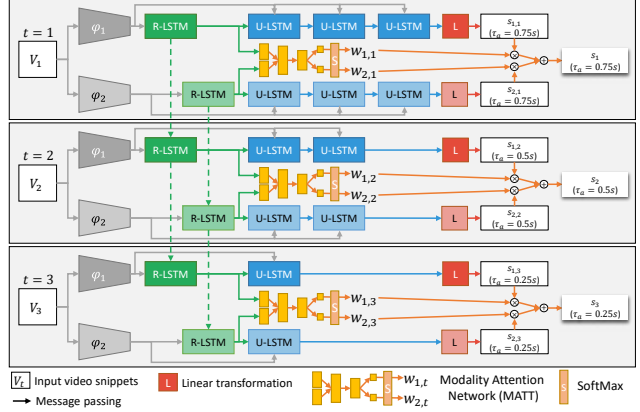


Figure 5. Example of the complete RU architecture with two modalities and the Modality ATTention mechanism (MATT).

one modality might be more useful than another (e.g., appearance over motion), depending on the processed sample. Inspired by previous work on attention [3, 58] and multi-modal fusion [38], we introduce a Modality ATTention (MATT) module which computes a set of attention scores indicating the relative importance of each modality for the final prediction. At a given time-step t , such scores are obtained by processing the concatenation of the hidden and cell vectors of the R-LSTM networks belonging to all branches $m = 1, \dots, M$ with a deep neural network D depending on the learnable parameters θ^{MATT} :

$$\lambda_t = D_{\theta^{MATT}}(\oplus_{m=1}^M (h_{m,t}^R \oplus c_{m,t}^R)) \quad (4)$$

where \oplus denotes the concatenation operator and $\oplus_{m=1}^M (h_{m,t}^R \oplus c_{m,t}^R)$ is the concatenation of the hidden and cell vectors produced by the R-LSTM at time-step t across all modalities. Late fusion weights can be obtained normalizing the score vector λ_t with the softmax function in order to make sure that fusion weights sum to one: $w_{m,t} = \frac{\exp(\lambda_{t,m})}{\sum_k \exp(\lambda_{t,k})}$, where $\lambda_{t,m}$ is the m^{th} component of the score vector λ_t . The final set of fusion weights is obtained at time-step t by merging the modality-specific predictions produced by the different branches with a linear combination as follows: $s_t = \sum_m w_{m,t} \cdot s_{m,t}$. Figure 5 illustrates an example of a complete RU with two modalities and the MATT fusion mechanism. For illustration purposes, the figure shows only three anticipation steps.

Branches and Representation Functions We instantiate the proposed architecture with 3 branches: a spatial branch which processes RGB frames, a motion branch which processes optical flow, and an object branch which processes object-based features. Our architecture analyzes video snippets of 5 frames $V_t = \{I_{t,1}, I_{t,2}, \dots, I_{t,5}\}$, where $I_{t,i}$ is the i^{th} frame of the video snippet. The representation function φ_1 of the spatial branch computes the feature vector $f_{1,t}$ by extracting features from the last frame $I_{t,5}$ of the video snippet using a Batch Normalized Inception CNN [23] trained

for action recognition. The representation function φ_2 of the motion branch extracts optical flow from the 5 frames of the current video snippet as proposed in [57]. The computed x and y optical flow form a tensor with 10 channels, which is fed to a Batch Normalized Inception CNN trained for action recognition to obtain the feature vector $f_{2,t}$. Note that φ_1 and φ_2 allow to obtain “action-centric” representations of the input frame which can be used by the R-LSTM to summarize what has happened in the past. The representation function φ_3 of the object branch extracts objects from the last frame $I_{t,5}$ of the input snippet V_t using an object detector. A fixed-length representation $f_{3,t}$ is obtained by accumulating the confidence scores of all bounding boxes predicted for each object class. Specifically, let $b_{t,i}$ be the i^{th} bounding box detected in image $I_{t,5}$, let $b_{t,i}^c$ be its class and let $b_{t,i}^s$ be its detection confidence score. The j^{th} component of the output representation vector $f_{3,t}$ is obtained by summing the confidence scores of all detected objects of class j , i.e., $f_{3,t,j} = \sum_i [b_{t,i}^c = j] b_{t,i}^s$, where $[\cdot]$ denotes the Iverson bracket. This representation only encodes the presence of an object in the scene, discarding its position in the frame, similarly to the representation proposed in [40] for egocentric activity recognition. We found this holistic representation to be sufficient in the case of egocentric action anticipation. Differently from φ_1 and φ_2 , φ_3 produces object-centric features, which carry information on what objects are present in the scene and hence could be interacted next.¹

4. Experimental Settings

Datasets We perform experiments on two large-scale datasets of egocentric videos: EPIC-Kitchens [8] and EGTEA Gaze+ [33]. EPIC-Kitchens contains 39,596 action annotations, 125 verbs, and 352 nouns. We split the public training set of EPIC-Kitchens (28,472 action segments) into training (23,493 segments) and validation (4,979 segments) sets by randomly choosing 232 videos for training and 40 videos for validation. We considered all unique (*verb, noun*) pairs in the public training set, thus obtaining 2,513 unique actions. EGTEA Gaze+ contains 10,325 action annotations, 19 verbs, 51 nouns and 106 unique actions. Methods are evaluated on EGTEA Gaze+ reporting the average performance across the three splits provided by the authors of the dataset [33].

Evaluation Measures Methods are evaluated using Top-k evaluation measures, i.e., we deem a prediction correct if the ground truth action falls in the top-k predictions. As observed in previous works [16, 29], this evaluation scheme is appropriate given the uncertainty of future predictions (i.e., many plausible actions can follow an observation). Specifically, we use the Top-5 accuracy as a class-agnostic mea-

¹See supp. for details on implementation and training of our method.

sure and Mean Top-5 Recall as a class aware metric. Top-5 recall [16] for a given class c is defined as the fraction of samples of ground truth class c for which the class c is in the list of the top-5 anticipated actions. Mean Top-5 Recall averages Top-5 recall values over classes. When evaluating on EPIC-Kitchens, Top-5 Recall is averaged over the provided list of many-shot verbs, nouns and actions. Results on the EPIC-Kitchens official test set are reported using the suggested evaluation measures, i.e., Top-1 accuracy, Top-5 accuracy, Precision and Recall.

To assess the timeliness of anticipations, we design a new evaluation measure inspired by the AMOC curve [20]. Let s_t be the score anticipated at time-step t for an action of ground truth class c , let τ_t be the anticipation time at time-step t , and let $tk(s_t)$ be the set of top- k actions as ranked by the scores s_t . We define as “time to action” at rank k the largest anticipation time (i.e., the time of earliest anticipation) in which a correct prediction has been made according to the top- k criterion: $TtA(k) = \max\{\tau(s_t) | c \in tk(s_t), \forall t\}$. If an action is not correctly anticipated in any of the time-steps, we set $TtA(k) = 0$. The mean time to action over the whole test set $mTtA(k)$ indicates how early, in average, a method can anticipate actions.

Performances are evaluated for verb, noun and action predictions on the EPIC-Kitchens dataset. We obtain verb and noun scores by marginalization over action scores for all methods except the one proposed in [8], which predicts verb and noun scores directly.

Compared Methods We compare the proposed method with respect to 7 state-of-the approaches and baselines. Specifically, we consider the Deep Multimodal Regressor (DMR) proposed in [54], the Anticipation Temporal Segment Network (ATSN) of [8], the anticipation Temporal Segment Network trained with verb-noun Marginal Cross Entropy Loss (MCE) described in [16], and the Encoder-Decoder LSTM (ED) introduced in [18]. We also consider baselines obtained adapting approaches proposed for early action recognition to the problem of egocentric action anticipation: the Feedback Network LSTM (FN) proposed in [19], a single LSTM architecture (we use the same parameters of our R-LSTM) trained using the Ranking Loss on Detection Score proposed in [36] (RL), and an LSTM trained using the Exponential Anticipation Loss proposed in [25] (EL). These baselines adopt the video processing scheme illustrated in Figure 2 and are implemented as two stream networks with a spatial and a temporal branch whose predictions are fused by late fusion.²

5. Results

Anticipation Results on EPIC-KITCHES Table 1 compares RU with respect to the state-of-the-art on our EPIC-

²See supp. for implementation details of the considered methods.

	Top-5 ACTION Accuracy% @ different τ_a (s)								Top-5 Acc.% @ 1s			M. Top-5 Rec.% @ 1s			Mean $TtA(5)$		
	2	1.75	1.5	1.25	1.0	0.75	0.5	0.25	VERB	NOUN	ACT.	VERB	NOUN	ACT.	VERB	NOUN	ACT.
DMR [54]	/	/	/	/	16.86	/	/	/	73.66	29.99	16.86	24.50	20.89	03.23	/	/	/
ATSN [8]	/	/	/	/	16.29	/	/	/	<u>77.30</u>	39.93	16.29	33.08	32.77	07.60	/	/	/
MCE [16]	/	/	/	/	26.11	/	/	/	73.35	38.86	26.11	34.62	32.59	06.50	/	/	/
ED [18]	21.53	22.22	23.20	24.78	25.75	26.69	27.66	29.74	75.46	42.96	25.75	<u>41.77</u>	<u>42.59</u>	<u>10.97</u>	<u>01.60</u>	<u>01.02</u>	<u>00.63</u>
FN [19]	23.47	24.07	24.68	25.66	26.27	26.87	27.88	28.96	74.84	40.87	26.27	35.30	37.77	06.64	01.52	00.86	00.56
RL [36]	<u>25.95</u>	<u>26.49</u>	<u>27.15</u>	<u>28.48</u>	<u>29.61</u>	<u>30.81</u>	<u>31.86</u>	32.84	76.79	<u>44.53</u>	<u>29.61</u>	40.80	40.87	10.64	01.57	00.94	00.62
EL [25]	24.68	25.68	26.41	27.35	28.56	30.27	31.50	<u>33.55</u>	75.66	43.72	28.56	38.70	40.32	08.62	01.55	00.94	00.62
RU	29.44	30.73	32.24	33.41	35.32	36.34	37.37	38.98	79.55	51.79	35.32	43.72	49.90	15.10	01.62	01.11	00.76
Improv.	+3.49	+4.24	+5.09	+4.93	+5.71	+5.53	+5.51	+5.43	+2.25	+7.26	+5.71	+1.95	+7.31	+4.13	+0.02	+0.09	+0.13

Table 1. Egocentric action anticipation results on the EPIC-KITCHENS dataset

Kitchens validation set. The left part of the table reports Top-5 action accuracy for the 8 anticipation times. Note that some methods [8, 16, 54] can anticipate actions only at a fixed anticipation time. The right part of the table breaks down Top-5 accuracy and Mean Top-5 Recall for verbs, nouns and actions, for anticipation time $\tau_a = 1s$, and reports mean $TtA(5)$ scores. Best results are in bold, whereas second-best results are underlined. The last row reports the improvement of RU with respect to second best results.

The proposed approach outperforms all competitors by consistent margins according to all evaluation measures, reporting an average improvement over prior art of about 5% on Top-5 action accuracy with respect to all anticipation times. Note that this margin is significant given the large number of 2, 513 action classes present in the dataset. Methods based on TSN (ATSN and MCE) generally achieve low performance, which suggests that simply adapting action recognition methods to the problem of anticipation is insufficient. Interestingly, DMR and ED, which are explicitly trained to anticipate future representations, achieve sub-optimal Top-5 action accuracy as compared to methods trained to predict future actions directly from input images (e.g., compare DMR with MCE, and ED with FN/RL/EL). This might be due to the fact that anticipating future representations is particularly challenging in the case of egocentric video, where the visual content changes continuously because of the mobility of the camera. RL consistently achieves second best results with respect to all anticipation times, except $\tau_a = 0.25$, where it is outperformed by EL. RU is particularly strong on nouns, achieving a Top-5 noun accuracy of 51.79% and a mean Top-5 noun recall of 49.90%, which improves over prior art by +7.26% and +7.31%. The small drop in performance when passing from class-agnostic measures to class-aware measures (i.e., 51.79% to 49.90%) suggests that our method does not over-fit to the distribution of nouns of the training set. It is worth noting that mean Top-5 Recall values are averaged over fairly large sets of 26 many-shot verbs, 71 many-shot nouns, and 819 many-shot actions, as specified in [8]. Differently, all compared methods exhibit large drops in verb and action performance when passing from class-agnostic to class-aware measures. Our insight into this different pattern is that anticipating the next object which will be used

	Top-1 Accuracy%			Top-5 Accuracy%			
	VERB	NOUN	ACT.	VERB	NOUN	ACT.	
S1	2SCNN [8]	29.76	15.15	04.32	76.03	38.56	15.21
	ATSN [8]	31.81	<u>16.22</u>	06.00	<u>76.56</u>	<u>42.15</u>	<u>28.21</u>
	MCE [16]	27.92	16.09	<u>10.76</u>	73.59	39.32	25.28
	RU	33.04	22.78	14.39	79.55	50.95	33.73
	Improvement	+1.23	+6.56	+3.63	+2.99	+8.80	+5.52
S2	2SCNN [8]	25.23	09.97	02.29	<u>68.66</u>	27.38	09.35
	ATSN [8]	<u>25.30</u>	<u>10.41</u>	02.39	68.32	<u>29.50</u>	06.63
	MCE [16]	21.27	09.90	<u>05.57</u>	63.33	25.50	<u>15.71</u>
	RU	27.01	15.19	08.16	69.55	34.38	21.10
	Improvement	+1.71	+4.78	+2.59	+0.89	+4.88	+5.39

Table 2. Anticipation results on the EPIC-KITCHENS test sets.

(i.e., anticipating nouns) is much less ambiguous than anticipating the way in which the object will be used (i.e., anticipating verbs and actions). It is worth noting that second best Top-5 verb and noun accuracy are achieved by different methods (i.e., ATSN in the case of verbs and RL in the case of nouns), while both are outperformed by the proposed RU. Despite its low performance with respect to class-agnostic measures, ED systematically achieves second best results with respect to mean Top-5 recall and mean $TtA(5)$. This highlights that there is no clear second-best performing method. Finally, mean $TtA(k)$ highlights that the proposed method can anticipate verbs, nouns and actions 1.62, 1.11 and 0.76 seconds in advance respectively.

Table 2 assesses the performance of the proposed method on the official test sets of EPIC-Kitchens dataset.³ RU outperforms all competitors by consistent margins on both the “seen” test, which includes scenes appearing in the training set (**S1**) and the “unseen” test set, with scenes not appearing in the training set (**S2**). Also in this case, RU is strong on nouns, obtaining +6.56% and +8.8% in **S1**, as well as +4.78% and +4.88 in **S2**. Improvements in terms of actions are also significant: +3.63% and +5.52% in **S1**, as well as +2.59% and +5.39% on **S2**.

Anticipation Results on EGTEA Gaze+ Table 3 reports Top-5 action accuracy scores achieved by the compared methods on EGTEA Gaze+ with respect to the 8 considered anticipation times. The table also reports mean $TtA(5)$ action scores, denoted as TtA . The proposed method outperforms all competitors at all anticipation times. Note that the margins of improvement are smaller on EGTEA Gaze+

³See supp. for precision and recall results.

	Top-5 ACTION Accuracy% @ different τ_a (s)								TIA
	2	1.75	1.5	1.25	1.0	0.75	0.5	0.25	
DMR [54]	/	/	/	/	55.70	/	/	/	/
ATSN [8]	/	/	/	/	40.53	/	/	/	/
MCE [16]	/	/	/	/	56.29	/	/	/	/
ED [18]	45.03	46.22	46.86	48.36	50.22	51.86	49.99	49.17	1.24
FN [19]	54.06	54.94	56.75	58.34	60.12	62.03	63.96	66.45	1.26
RL [36]	55.18	56.31	58.22	60.35	62.56	64.65	67.35	70.42	1.29
EL [25]	<u>55.62</u>	<u>57.56</u>	<u>59.77</u>	<u>61.58</u>	<u>64.62</u>	<u>66.89</u>	<u>69.60</u>	<u>72.38</u>	<u>1.32</u>
RU	56.82	59.13	61.42	63.53	66.40	68.41	71.84	74.28	1.41
Imp	+1.20	+1.57	+1.65	+1.95	+1.78	+1.52	+2.24	+1.89	+0.09

Table 3. Anticipation results on EGTEA Gaze+.

due to its smaller-scale (106 actions in vs 2, 513 actions in EPIC-KITCHENS). Differently from Table 1, EL systematically achieves second best performance on EGTEA Gaze+, which highlights again that there is no clear second best competitor to RU in our evaluations.

Ablation Study on EPIC-Kitchens To assess the role of rolling-unrolling, we consider a strong baseline composed of a single LSTM (same configuration as R-LSTM) and three branches (RGB, Flow, OBJ) with late fusion (BL). To tear apart the influence of rolling-unrolling and MATT, Table 4(a) compares this baseline with the proposed RU architecture, where MATT has been replaced with late fusion. The proposed RU approach brings systematic improvements over the baseline for all anticipation times, with larger improvements in the case of the object branch.

Table 4(b) reports the performances of the single branches of RU and compares MATT with respect to late fusion (i.e., averaging predictions) and early fusion (i.e., feeding the model with the concatenation of the modality-specific representations). MATT always outperforms late fusion, which consistently achieves second best results. Early fusion always leads to sub-optimal results. All fusion schemes always improve over the single branches. Figure 6 shows regression plots of modality attention weights computed on all samples of the validation set. RGB and OBJ weights are characterized by a strong and steep correlation. A similar pattern applies to Flow and OBJ weights, whereas Flow and RGB weights are characterized by a small positive correlation. This suggests that MATT gives more credit to OBJ when RGB and Flow are less informative, whereas it relies on RGB and Flow when objects are not necessary.

Table 4(c) assesses the role of pre-training the proposed architecture with sequence completion. As can be noted, the proposed pre-training procedure brings small but consistent improvements for most anticipation times. Table 4(d) compares RU with the strong baseline of Table 4(a). The comparison highlights the cumulative effect of all the proposed procedures/component with respect to a strong baseline using three modalities. It is worth noting that the proposed architecture brings improvements for all anticipation times, ranging from +1.53% to +4.08%.

Qualitative Results Figure 7 reports two examples of anticipations made by the proposed method at four antici-

	Top-5 ACTION Accuracy% @ different τ_a (s)								TIA
	2	1.75	1.5	1.25	1.0	0.75	0.5	0.25	
BL (Late)	27.96	28.76	29.99	31.09	32.02	33.09	34.13	34.92	0.66
RU (Late)	29.10	29.77	31.72	33.09	34.23	35.28	36.10	37.61	0.73
Imp.	+1.14	+1.01	+1.73	+2.00	+2.21	+2.19	+1.97	+2.69	+0.07

(a) Rolling-Unrolling Mechanism.

RU (RGB)	25.44	26.89	28.32	29.42	30.83	32.00	33.31	34.47	0.69
RU (Flow)	17.38	18.04	18.91	19.97	21.42	22.37	23.49	24.18	0.51
RU (OBJ)	24.56	25.60	26.61	28.32	29.89	30.85	31.82	33.39	0.67
Early Fusion	25.58	27.25	28.58	29.59	31.88	32.78	33.99	35.62	0.72
Late Fusion	29.10	29.77	31.72	33.09	34.23	35.28	36.10	37.61	0.73
MATT	29.44	30.73	32.24	33.41	35.32	36.34	37.37	38.98	0.76
Imp.	+0.34	+0.96	+0.52	+0.32	+1.09	+1.06	+1.27	+1.37	+0.03

(b) Modality Attention Fusion Mechanism.

w/o SCP	29.22	30.43	32.34	33.37	34.75	35.84	36.79	37.93	0.75
with SCP	29.44	30.73	32.24	33.41	35.32	36.34	37.37	38.98	0.76
Imp. of SCP	+0.22	+0.30	-0.10	+0.04	+0.57	+0.50	+0.58	+1.05	+0.01

(c) Sequence-Completion Pre-training.

BL (Fusion)	27.96	28.76	29.99	31.09	32.02	33.09	34.13	34.92	0.66
RU (Fusion)	29.44	30.73	32.24	33.41	35.32	36.34	37.37	38.98	0.76
Imp. (Fusion)	+1.48	+1.97	+2.25	+2.32	+3.30	+3.25	+3.24	+4.06	+0.1

(d) Overall comparison wrt strong baseline.

Table 4. Ablation study on EPIC-KITCHENS.




Figure 6. Correlations between modality attention weights

tion times. Under each frame we report top-4 predictions, whereas modality weights computed by MATT are reported in percentage on the right. Green bounding boxes are shown around the detected objects and the optical flow is illustrated in orange. In the first example (top), the model can predict “close door” based on context and past actions (e.g., taking objects out of the cupboard), hence it assigns large weights to RGB and Flow and low weights to OBJ. In the second example (bottom), the model initially predicts “squeeze lime” at $\tau_a = 2s$. Later, the prediction is corrected to “squeeze lemon” as soon as the lemon can be reliably detected. Note that in this case the network assigns larger weights to OBJ as compared to the previous example.⁴

Early Action Recognition We also observe that the proposed method generalizes to the task of early action recognition. We adapt our processing scheme by sampling 8 video snippets from each action segment uniformly in time and set $S_{enc} = 0$, $S_{ant} = 8$. The snippets are fed to the model, which produces predictions at each time-step, corresponding to the following observation rates: 12.5%, 25%, 37.5%, 50%, 62.5%, 75%, 87.5%, 100%. Branches are fused by late fusion in this case. We compare the proposed method with respect to FN, EL and RL.

⁴See supp. and <https://iplab.dmi.unict.it/rulstm/> for additional examples and videos.

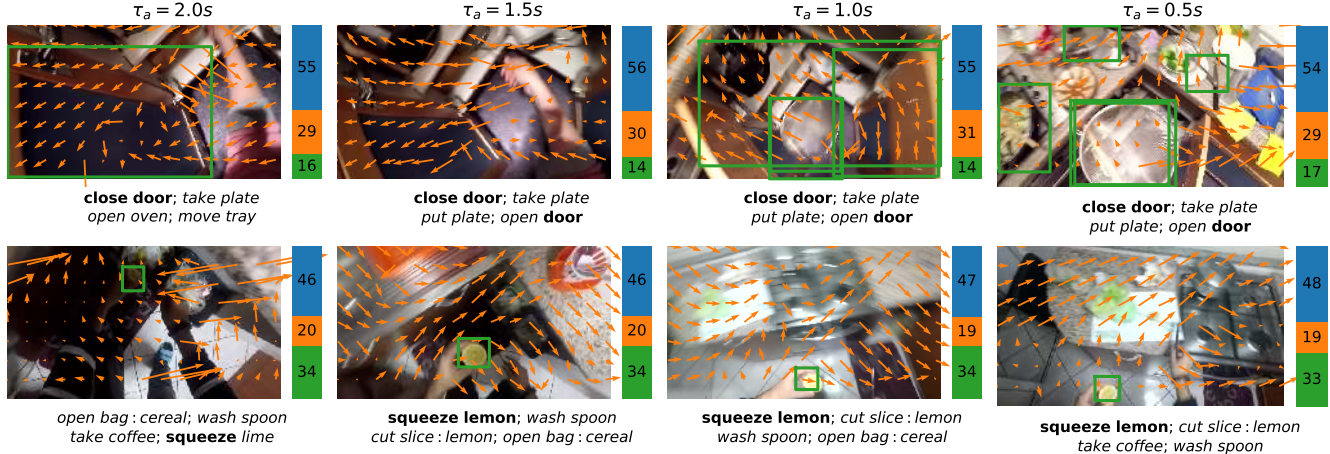


Figure 7. Qualitative examples (best seen on screen). Legend for attention weights: blue - RGB, orange - Flow, green - objects.

Top-1 ACTION Accuracy% @ different observation rates								
	12.5%	25.0%	37.5%	50.0%	62.5%	75.0%	87.5%	100%
FN [19]	19.61	23.85	25.66	26.85	27.47	28.34	28.26	28.38
EL [25]	19.69	23.27	26.03	27.49	29.06	29.97	30.91	31.46
RL [36]	<u>22.53</u>	<u>25.08</u>	<u>27.19</u>	<u>28.64</u>	<u>29.57</u>	<u>30.13</u>	30.45	30.47
RU	24.48	27.63	29.44	30.93	32.16	33.09	33.63	34.07
Imp.	+1.95	+2.55	+2.25	+2.29	+2.59	+2.96	+2.72	+2.61

Table 5. Early recognition results on EPIC-KITCHENS.

Top-1 ACTION Accuracy% @ different observation rates								
	12.5%	25.0%	37.5%	50.0%	62.5%	75.0%	87.5%	100%
FN [19]	44.02	50.32	53.34	55.10	56.58	57.31	57.95	57.72
EL [25]	40.31	48.08	51.84	54.71	56.93	58.45	<u>59.55</u>	60.18
RL [36]	<u>45.42</u>	<u>51.00</u>	<u>54.20</u>	<u>56.54</u>	<u>58.09</u>	<u>58.93</u>	59.29	59.50
RU	45.94	51.84	54.39	57.05	58.15	59.31	60.10	60.20
Imp.	+0.51	+0.84	+0.20	+0.51	+0.06	+0.38	+0.55	+0.02

Table 6. Early recognition results on EGTEA Gaze+.

Table 5 reports Top-1 early action recognition accuracy results obtained by the compared methods on our validation set of EPIC-Kitchens. The proposed method consistently outperforms the competitors at all observation rates. Interestingly, RU achieves an early action recognition accuracy of 33.09% when observing only 75% of the action, which is already comparable to the accuracy of 34.07% achieved when the full action is observed. This indicates that RU can timely recognize actions before they are completed. RL achieves second best results up to observation rate 75%, whereas EL achieves second best results when more than 75% of the action is observed, which confirms the lack of a clear second-best performer. Table 6 reports Top-1 accuracy results obtained on EGTEA Gaze+. The proposed RU outperforms the competitors for all observation rates by small but consistent margins. Coherently with Table 5, second best results are obtained by RL and EL.

Egocentric Action Recognition Results The proposed method can be used to perform egocentric action recognition by considering the predictions obtained for the observation rate of 100%. Table 7 compares the performance of the proposed method with other egocentric action recognition methods on the two test sets of EPIC-Kitchens.⁵

⁵See supp. for precision and recall results.

		Top-1 Accuracy%			Top-5 Accuracy%		
		VERB	NOUN	ACTION	VERB	NOUN	ACTION
S1	2SCNN [8]	42.16	29.14	13.23	80.58	53.70	30.36
	TSN [8]	48.23	36.71	20.54	84.09	<u>62.32</u>	39.79
	LSTA [51]	59.55	38.35	<u>30.33</u>	85.77	61.49	<u>49.97</u>
	MCE [16]	54.22	<u>38.85</u>	29.00	85.22	61.80	49.62
	RU	<u>56.93</u>	43.05	33.06	<u>85.68</u>	67.12	55.32
Imp.	-2.62	+4.20	+2.73	-0.09	+4.80	+5.35	
S2	2SCNN [8]	36.16	18.03	07.31	71.97	38.41	19.49
	TSN [8]	39.40	22.70	10.89	<u>74.29</u>	<u>45.72</u>	25.26
	LSTA [51]	47.32	22.16	<u>16.63</u>	77.02	43.15	30.93
	MCE [16]	40.90	<u>23.46</u>	16.39	72.11	43.05	<u>31.34</u>
	RU	<u>43.67</u>	26.77	19.49	73.30	48.28	37.15
Imp.	-3.65	+3.31	+2.86	-3.72	+2.56	+5.81	

Table 7. Recognition results on the EPIC-KITCHENS test sets.

Our RU outperforms all competitors in recognizing actions and nouns on both sets by significant margins, whereas it achieves second-best results in most cases for verb recognition. RU obtains an action recognition accuracy of 60.2% on EGTEA Gaze+. Despite being designed for action anticipation, RU outperforms recent approaches, such as Li et al. [33] (+6.9% wrt 53.3%) and Zhang et al. [61] (+3.19% wrt 57.01% - reported from [51]), and obtaining performances comparable to state-of-the-art approaches such as Sudhakaran and Lanz [52] (-0.56% wrt 60.76) and Sudhakaran et al. [51] (-1.66% wrt 61.86%).⁶

6. Conclusion

We presented RU-LSTM, a learning architecture which processes RGB, optical flow and object-based features using two LSTMs and a modality attention mechanism to anticipate actions from egocentric video. Experiments on two datasets show the superiority of the approach with respect to prior art and highlight generalization over datasets and tasks: anticipation, early recognition, and recognition.

Acknowledgment

This research is supported by Piano della Ricerca 2016-2018, linea di Intervento 2 of DMI, University of Catania.

⁶See supp. for the full table.

References

- [1] Y. Abu Farha, A. Richard, and J. Gall. When will you do what?-anticipating temporal occurrences of activities. In *Computer Vision and Pattern Recognition*, pages 5343–5352, 2018.
- [2] M. S. Aliakbarian, F. S. Saleh, M. Salzmann, B. Fernando, L. Petersson, and L. Andersson. Encouraging lstms to anticipate actions very early. In *International Conference on Computer Vision*, volume 1, pages 280–289, 2017.
- [3] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations*, 2015.
- [4] F. Becattini, T. Uricchio, L. Seidenari, A. Del Bimbo, and L. Ballan. Am I done? predicting action progress in videos. *British Machine Vision Conference*, 2017.
- [5] S. Z. Bokhari and K. M. Kitani. Long-term activity forecasting using first-person vision. In *Asian Conference on Computer Vision*, pages 346–360, 2016.
- [6] Y. Cao, D. Barrett, A. Barbu, S. Narayanaswamy, H. Yu, A. Michaux, Y. Lin, S. Dickinson, J. Mark Siskind, and S. Wang. Recognize human activities from partially observed videos. In *Computer Vision and Pattern Recognition*, pages 2658–2665, 2013.
- [7] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Computer Vision and Pattern Recognition*, pages 4724–4733, 2017.
- [8] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision*, pages 720–736, 2018.
- [9] R. De Geest, E. Gavves, A. Ghodrati, Z. Li, C. Snoek, and T. Tuytelaars. Online action detection. In *European Conference on Computer Vision*, pages 269–284, 2016.
- [10] C. Fan, J. Lee, and M. S. Ryoo. Forecasting hand and object locations in future frames. *CoRR*, abs/1705.07328, 2017.
- [11] A. Fathi, A. Farhadi, and J. M. Rehg. Understanding egocentric activities. In *International Conference on Computer Vision*, pages 407–414, 2011.
- [12] A. Fathi, Y. Li, and J. Rehg. Learning to recognize daily actions using gaze. In *European Conference on Computer Vision*, pages 314–327, 2012.
- [13] C. Feichtenhofer, A. Pinz, and R. P. Wildes. Spatiotemporal multiplier networks for video action recognition. In *Computer Vision and Pattern Recognition*, pages 4768–4777, 2017.
- [14] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Computer Vision and Pattern Recognition*, pages 1933–1941, 2016.
- [15] P. Felsen, P. Agrawal, and J. Malik. What will happen next? forecasting player moves in sports videos. In *International Conference on Computer Vision*, pages 3342–3351, 2017.
- [16] A. Furnari, S. Battiato, and G. M. Farinella. Leveraging uncertainty to rethink loss functions and evaluation measures for egocentric action anticipation. In *European Conference on Computer Vision Workshops*, 2018.
- [17] A. Furnari, S. Battiato, K. Grauman, and G. M. Farinella. Next-active-object prediction from egocentric videos. *Journal of Visual Communication and Image Representation*, 49:401–411, 2017.
- [18] Y. Gao, Z. Yang, and R. Nevatia. RED: Reinforced encoder-decoder networks for action anticipation. *British Machine Vision Conference*, 2017.
- [19] R. D. Geest and T. Tuytelaars. Modeling temporal structure with lstm for online action detection. In *Winter Conference on Applications in Computer Vision*, 2018.
- [20] M. Hoai and F. D. la Torre. Max-margin early event detectors. *International Journal of Computer Vision*, 107(2):191–202, 2014.
- [21] D. Huang, S. Yao, Y. Wang, and F. De La Torre. Sequential max-margin event detectors. In *European conference on computer vision*, pages 410–424, 2014.
- [22] D.-A. Huang and K. M. Kitani. Action-reaction: Forecasting the dynamics of human interaction. In *European Conference on Computer Vision*, pages 489–504, 2014.
- [23] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456, 2015.
- [24] A. Jain, H. S. Koppula, B. Raghavan, S. Soh, and A. Saxena. Car that knows before you do: Anticipating maneuvers via learning temporal driving models. In *International Conference on Computer Vision*, pages 3182–3190, 2015.
- [25] A. Jain, A. Singh, H. S. Koppula, S. Soh, and A. Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *International Conference on Robotics and Automation*, pages 3118–3125. IEEE, 2016.
- [26] T. Kanade and M. Hebert. First-person vision. *Proceedings of the IEEE*, 100(8):2442–2453, 2012.
- [27] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.
- [28] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert. Activity forecasting. In *European Conference on Computer Vision*, pages 201–214. Springer, 2012.
- [29] H. S. Koppula and A. Saxena. Anticipating human activities using object affordances for reactive robotic response. *Pattern Analysis and Machine Intelligence*, 38(1):14–29, 2016.
- [30] T. Lan, T.-S. Chen, and S. Savarese. A hierarchical representation for future action prediction. In *European Conference on Computer Vision*, pages 689–704, 2014.
- [31] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005.
- [32] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *Computer Vision and Pattern Recognition*, 2008.
- [33] Y. Li, M. Liu, and J. M. Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *European Conference on Computer Vision*, 2018.
- [34] Y. Li, Z. Ye, and J. M. Rehg. Delving into egocentric actions. In *Computer Vision and Pattern Recognition*, pages 287–295, 2015.

- [35] M. Ma, H. Fan, and K. M. Kitani. Going deeper into first-person activity recognition. In *Computer Vision and Pattern Recognition*, pages 1894–1903, 2016.
- [36] S. Ma, L. Sigal, and S. Sclaroff. Learning activity progression in lstms for activity detection and early detection. In *Computer Vision and Pattern Recognition*, 2016.
- [37] T. Mahmud, M. Hasan, and A. K. Roy-Chowdhury. Joint prediction of activity labels and starting times in untrimmed videos. In *International Conference on Computer Vision*, pages 5773–5782, 2017.
- [38] O. Mees, A. Eitel, and W. Burgard. Choosing smartly: Adaptive multimodal fusion for object detection in changing environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Deajeon, South Korea, Oct. 2016.
- [39] H. S. Park, J.-J. Hwang, Y. Niu, and J. Shi. Egocentric future localization. In *Computer Vision and Pattern Recognition*, pages 4697–4705, 2016.
- [40] H. Pirsivash and D. Ramanan. Detecting activities of daily living in first-person camera views. In *Computer Vision and Pattern Recognition*, pages 2847–2854, 2012.
- [41] N. Rhinehart and K. M. Kitani. First-person activity forecasting with online inverse reinforcement learning. In *International Conference on Computer Vision*, pages 3696–3705, 2017.
- [42] M. S. Ryoo. Human activity prediction: Early recognition of ongoing activities from streaming videos. In *International Conference on Computer Vision*, pages 1036–1043, 2011.
- [43] M. S. Ryoo, T. J. Fuchs, L. Xia, J. K. Aggarwal, and L. Matthies. Robot-centric activity prediction from first-person videos: What will they do to me? In *International Conference on Human-Robot Interaction*, pages 295–302, 2015.
- [44] M. S. Ryoo, B. Rothrock, and L. Matthies. Pooled motion features for first-person videos. In *Computer Vision and Pattern Recognition*, pages 896–904, 2015.
- [45] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, pages 568–576, 2014.
- [46] K. K. Singh, K. Fatahalian, and A. A. Efros. Krishnacam: Using a longitudinal, single-person, egocentric dataset for scene understanding tasks. In *Winter Conference on Applications of Computer Vision*, pages 1–9, 2016.
- [47] S. Singh, C. Arora, and C. Jawahar. First person action recognition using deep learned descriptors. In *Computer Vision and Pattern Recognition*, pages 2620–2628, 2016.
- [48] S. Singh, C. Arora, and C. Jawahar. Trajectory aligned features for first person action recognition. *Pattern Recognition*, 62:45–55, 2017.
- [49] B. Soran, A. Farhadi, and L. Shapiro. Generating notifications for missing actions: Don’t forget to turn the lights off! In *International Conference on Computer Vision*, pages 4669–4677, 2015.
- [50] E. H. Spriggs, F. De La Torre, and M. Hebert. Temporal segmentation and activity classification from first-person sensing. In *Computer Vision and Pattern Recognition Workshops*, pages 17–24, 2009.
- [51] S. Sudhakaran, S. Escalera, and O. Lanz. Lsta: Long short-term attention for egocentric action recognition. *arXiv preprint arXiv:1811.10698*, 2018.
- [52] S. Sudhakaran and O. Lanz. Attention is all we need: Nailing down object-centric attention for egocentric activity recognition. *British Machine Vision Conference*, 2018.
- [53] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *International Conference on Computer Vision*, pages 4489–4497, 2015.
- [54] C. Vondrick, H. Pirsivash, and A. Torralba. Anticipating visual representations from unlabeled video. In *Computer Vision and Pattern Recognition*, pages 98–106, 2016.
- [55] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Dense trajectories and motion boundary descriptors for action recognition. *International Journal of Computer Vision*, 103(1):60–79, 2013.
- [56] H. Wang and C. Schmid. Action recognition with improved trajectories. In *International Conference on Computer Vision*, pages 3551–3558, 2013.
- [57] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. V. Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision*, pages 20–36, 2016.
- [58] 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 *International conference on machine learning*, pages 2048–2057, 2015.
- [59] K.-H. Zeng, W. B. Shen, D.-A. Huang, M. Sun, and J. Carlos Nieves. Visual forecasting by imitating dynamics in natural sequences. In *International Conference on Computer Vision*, pages 2999–3008, 2017.
- [60] M. Zhang, K. T. Ma, J.-H. Lim, Q. Zhao, and J. Feng. Deep future gaze: Gaze anticipation on egocentric videos using adversarial networks. In *Computer Vision and Pattern Recognition*, pages 3539–3548, 2017.
- [61] P. Zhang, J. Xue, C. Lan, W. Zeng, Z. Gao, and N. Zheng. Adding attentiveness to the neurons in recurrent neural networks. In *European Conference on Computer Vision*, pages 135–151, 2018.
- [62] B. Zhou, A. Andonian, A. Oliva, and A. Torralba. Temporal relational reasoning in videos. In *European Conference on Computer Vision*, pages 803–818, 2018.
- [63] Y. Zhou and T. L. Berg. Temporal perception and prediction in ego-centric video. In *International Conference on Computer Vision*, pages 4498–4506, 2016.