

Encouraging LSTMs to Anticipate Actions Very Early

Mohammad Sadegh Aliakbarian^{1,3}, Fatemeh Sadat Saleh^{1,3}, Mathieu Salzmann², Basura Fernando¹, Lars Petersson^{1,3}, Lars Andersson³

¹Australian National University, ²CVLab, EPFL, Switzerland, ³Smart Vision Systems, CSIRO

firstname.lastname@data61.csiro.au, mathieu.salzmann@epfl.ch, basura.fernando@anu.edu.au

Abstract

In contrast to the widely studied problem of recognizing an action given a complete sequence, action anticipation aims to identify the action from only partially available videos. As such, it is therefore key to the success of computer vision applications requiring to react as early as possible, such as autonomous navigation. In this paper, we propose a new action anticipation method that achieves high prediction accuracy even in the presence of a very small percentage of a video sequence. To this end, we develop a multi-stage LSTM architecture that leverages context-aware and action-aware features, and introduce a novel loss function that encourages the model to predict the correct class as early as possible. Our experiments on standard benchmark datasets evidence the benefits of our approach; We outperform the state-of-the-art action anticipation methods for early prediction by a relative increase in accuracy of 22.0% on JHMDB-21, 14.0% on UT-Interaction and 49.9% on UCF-101.

1. Introduction

Understanding actions from videos is key to the success of many real-world applications, such as autonomous navigation and sports analysis. While great progress has been made to recognize actions from complete sequences [16, 38, 7, 9, 30, 41, 31, 10, 2] in the past decade, action anticipation [25, 26, 24, 45, 32, 33, 21], which aims to predict the observed action as early as possible, has become a popular research problem only recently. Anticipation is crucial in scenarios where one needs to react before the action is finalized, such as to avoid hitting pedestrians with an autonomous car, or to forecast dangerous situations in surveillance scenarios.

The key difference between recognition and anticipation lies in the fact that the methods tackling the latter should predict the correct class as early as possible, given only a few frames from the beginning of the video sequence. To address this, several approaches have introduced new train-

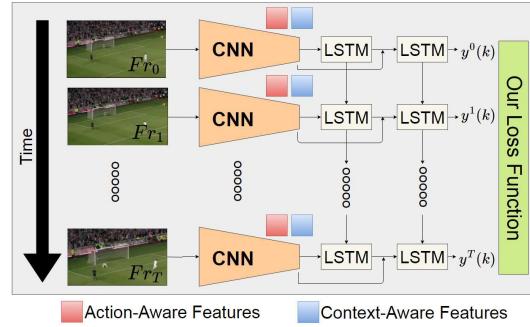


Figure 1. **Overview of our approach.** Given a small portion of sequential data, our approach is able to predict the action category with very high performance. For instance, in UCF-101, our approach anticipates actions with more than 80% accuracy given only the first 1% of the video. To achieve this, we design a model that leverages action- and context-aware features together with a new loss function that encourages the model to make correct predictions as early as possible.

ing losses encouraging the score [32] or the rank [21] of the correct action to increase with time, or penalizing increasingly strongly the classification mistakes [11]. In practice, however, the effectiveness of these losses remains limited for very early prediction, such as from 1% of the sequence.

In this paper, we introduce a novel loss that encourages making correct predictions very early. Specifically, our loss models the intuition that some actions, such as running and high jump, are highly ambiguous after seeing only the first few frames, and false positives should therefore not be penalized too strongly in the early stages. By contrast, we would like to predict a high probability for the correct class as early as possible, and thus penalize false negatives from the beginning of the sequence. Our experiments demonstrate that, for a given model, our new loss yields significantly higher accuracy than existing ones on the task of early prediction.

In particular, in this paper, we also contribute a novel multi-stage Long Short Term Memory (LSTM) architecture for action anticipation. This model effectively extracts and jointly exploits context- and action-aware features (see

Fig. 1). This is in contrast to existing methods that typically extract either global representations for the entire image [6, 41, 42, 7] or video sequence [35, 14], thus not focusing on the action itself, or localize the feature extraction process to the action itself via dense trajectories [39, 38, 9], optical flow [8, 41, 18] or actionness [40, 3, 12, 44, 36, 22], thus failing to exploit contextual information. To the best of our knowledge, only two-stream networks [28, 8, 4, 20] have attempted to jointly leverage both information types by making use of RGB frames in conjunction with optical flow to localize the action. Exploiting optical flow, however, does not allow these methods to explicitly leverage appearance in the localization process. Furthermore, computing optical flow is typically expensive, thus significantly increasing the runtime of these methods. By not relying on optical flow, our method is significantly more efficient: On a single GPU, our model analyzes a short video (e.g., 50 frames) 14 times faster than [28] and [8].

Our model is depicted in Fig. 2. In a first stage, it focuses on the global, context-aware information by extracting features from the entire RGB image. The second stage then combines these context-aware features with action-aware ones obtained by exploiting class-specific activations, typically corresponding to regions where the action occurs. In short, our model first extracts the contextual information, and then merges it with the localized one.

As evidenced by our experiments, our approach significantly outperforms the state-of-the-art action anticipation methods on all the standard benchmark datasets that we evaluated on, including UCF-101 [34], UT-Interaction [24], and JHMDB21 [13]. In the supplementary material, we further show that our combination of context- and action-aware features is also beneficial for the more traditional task of action recognition. Moreover, we evaluate the effect of optical flow features for both action recognition and anticipation.

2. Related Work

The focus of this paper is twofold: Action anticipation, with a new loss that encourages correct prediction as early as possible, and action modeling, with a model that combines context- and action-aware information using multi-stage LSTMs. Below, we discuss the most relevant approaches for these two aspects.

2.1. Action Anticipation

The idea of action anticipation was introduced by [26], which models causal relationships to predict human activities. This was followed by several attempts to model the dynamics of the observed actions, such as by introducing integral and dynamic bag-of-words [25], using spatial-temporal implicit shape models [45], extracting human body movements via skeleton information [47], and accounting for the complete and partial history of observed features [15].

More recently, [33] proposed to make use of binary SVMs to classify video snippets into sub-action categories and obtain the final class label in an online manner using dynamic programming. To overcome the need to train one classifier per sub-action, [32] extended this approach to using a structural SVM. Importantly, this work further introduced a new objective function to encourage the score of the correct action to increase as time progresses.

While the above-mentioned work made use of hand-crafted features, recent advances have naturally led to the development of deep learning approaches to action anticipation. In this context, [21] proposed to combine a Convolutional Neural Network (CNN) with an LSTM to model both spatial and temporal information. The authors further introduced new ranking losses whose goal is to enforce either the score of the correct class or the margin between the score of the correct class and that of the best score to be non-decreasing over time. Similarly, in [11], a new loss that penalizes classification mistakes increasingly strongly over time was introduced in an LSTM-based framework that used multiple modalities. While the two above-mentioned methods indeed aim at improving classification accuracy over time, they do not explicitly encourage making correct predictions as early as possible. By contrast, while accounting for ambiguities in early stages, our new loss still aims to prevent false negatives from the beginning of the sequence.

Instead of predicting the future class label, in [37], the authors proposed to predict the future visual representation. However, the main motivation for this was to work with unlabeled videos, and the learned representation is therefore not always related to the action itself.

2.2. Action Modeling

Most recent action approaches extract global representations for the entire image [6, 42, 7] or video sequence [35, 14]. As such, these methods do not truly focus on the actions of interest, but rather compute a *context-aware* representation. Unfortunately, context does not always bring reliable information about the action. For example, one can play guitar in a bedroom, a concert hall or a yard. To overcome this, some methods localize the feature extraction process by exploiting dense trajectories [39, 38, 9] or optical flow [18]. Inspired by objectness, the notion of actionness [40, 3, 12, 44, 36, 22] has recently also been proposed to localize the regions where a generic action occurs. The resulting methods can then be thought of as extracting *action-aware* representations. In other words, these methods go to the other extreme and completely discard the notion of context which can be useful for some actions, such as playing football on a grass field.

There is nevertheless a third class of methods that aim to leverage these two types of information [28, 8, 4, 20, 41]. By combining RGB frames and optical flow in two-stream

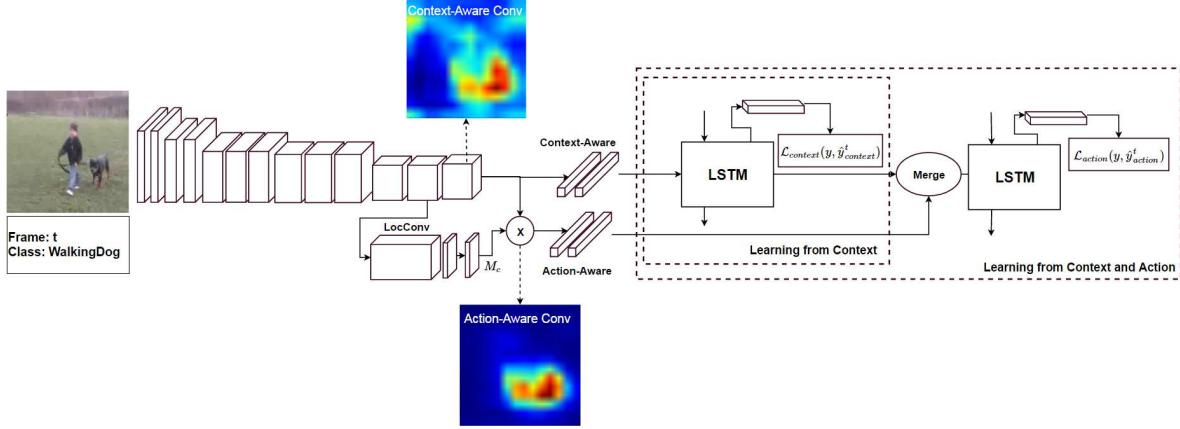


Figure 2. **Overview of our approach.** We propose to extract context-aware features, encoding global information about the scene, and combine them with action-aware ones, which focus on the action itself. To this end, we introduce a multi-stage LSTM architecture that leverages the two types of features to predict the action or forecast it. Note that, for the sake of visualization, the color maps were obtained from 3D tensors ($512 \times W \times H$) via an average pooling operation over the 512 channels.

architectures, these methods truly exploit context and motion, from which the action can be localized by learning to distinguish relevant motion. This localization, however, does not directly exploit appearance. Here, inspired by the success of these methods, we develop a novel multi-stage network that also leverages context- and action-aware information. However, we introduce a new action-aware representation that exploits the RGB data to localize the action. As such, our approach effectively leverages appearance for action-aware modeling, and, by avoiding the expensive optical flow computation, is much more efficient than the above-mentioned two-stream models. In particular, our model is about 14 times faster than the state-of-the-art two-stream network of [8] and has less parameters. The reduction in number of parameters is due to the fact that [8] and [28] rely on two VGG-like networks (one for each stream) with a few additional layers (including 3D convolutions for [8]). By contrast, our model has, in essence, a single VGG-like architecture, with some additional LSTM layers, which only have few parameters. Moreover, our work constitutes the first attempt at explicitly leveraging context- and action-aware information for action anticipation. Finally, we introduce a novel multi-stage LSTM fusion strategy to integrate action and context aware features.

3. Our Approach

Our goal is to predict the class of an action as early as possible, that is, after having seen only a very small portion of the video sequence. To this end, we first introduce a new loss function that encourages making correct predictions very early. We then develop a multi-stage LSTM model that makes use of this loss function while leveraging both context- and action-aware information.

3.1. A New Loss for Action Anticipation

As argued above, a loss for action anticipation should encourage having a high score for the correct class as early as possible. However, it should also account for the fact that, early in the sequence, there can often be ambiguities between several actions, such as running and high jump. Below, we introduce a new anticipation loss that follows these two intuitions.

Specifically, let $y^t(k)$ encode the true activity label at time t , i.e., $y^t(k) = 1$ if the sample belongs to class k and 0 otherwise, and $\hat{y}^t(k)$ denote the corresponding label predicted by a given model. We define our new loss as

$$\mathcal{L}(y, \hat{y}) = -\frac{1}{N} \sum_{k=1}^N \sum_{t=1}^T \left[y^t(k) \log(\hat{y}^t(k)) + \frac{t(1-y^t(k))}{T} \log(1-\hat{y}^t(k)) \right], \quad (1)$$

where N is the number of action classes and T the length (number of frames) of the input sequence.

This loss function consists of two terms. The first one penalizes false negatives with the same strength at any point in time. By contrast, the second one focuses on false positives, and its strength increases linearly over time, to reach the same weight as that on false negatives. Therefore, the relative weight of the first term compared to the second one is larger at the beginning of the sequence. Altogether, this encourages predicting a high score for the correct class as early as possible, i.e., preventing false negatives, while accounting for the potential ambiguities at the beginning of the sequence, which give rise to false positives. As we see more frames, however, the ambiguities are removed, and these false positives are encouraged to disappear.

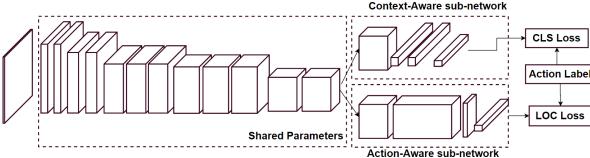


Figure 3. Our feature extraction network. Our CNN model for feature extraction is based on the VGG-16 structure with some modifications. Up to conv5-2, the network is the same as VGG-16. The output of this layer is connected to two sub-models. The first one extracts context-aware features by providing a global image representation. The second one relies on another network to extract action-aware features.

Our new loss matches the desirable properties of an action anticipation loss. In the next section, we introduce a novel multi-stage architecture that makes use of this loss.

3.2. Multi-stage LSTM Architecture

To tackle action anticipation, we develop the novel multi-stage recurrent architecture based on LSTMs depicted by Fig. 2. This architecture consists of a stage-wise combination of context- and action-aware information. Below, we first discuss our approach to extracting these two types of information, and then present our complete multi-stage recurrent network.

3.2.1 Context- and Action-aware Modeling

To model the context- and action-aware information, we introduce the two-stream architecture shown in Fig. 3. The first part of this network is shared by both streams and, up to conv5-2, corresponds to the VGG-16 network [29], pre-trained on ImageNet for object recognition. The output of this layer is connected to two sub-models: One for context-aware features and the other for action-aware ones. We then train these two sub-models for the same task of action recognition from a single image using a cross-entropy loss function defined on the output of each stream. In practice, we found that training the entire model in an end-to-end manner did not yield a significant improvement over training only the two sub-models. In our experiments, we therefore opted for this latter strategy, which is less expensive computationally and memory-wise. Below, we first discuss the context-aware sub-network and then turn to the action-aware one.

Context-Aware Feature Extraction. This sub-model is similar to VGG-16 from conv5-3 up to the last fully-connected layer, with the number of units in the last fully-connected layer changed from 1000 (original 1000-way ImageNet classification model) to the number of activities N .

In essence, this sub-model focuses on extracting a deep representation of the whole scene for each activity and thus

incorporates context. We then take the output of its fc7 layer as our context-aware features.

Action-Aware Feature Extraction. As mentioned before, the context of an action does not always correlate with the action itself. Our second sub-model therefore aims at extracting features that focus on the action itself. To this end, we draw inspiration from the object classification work of [48]. At the core of this work lies the idea of Class Activation Maps (CAMs). In our context, a CAM indicates the regions in the input image that contribute most to predicting each class label. In other words, it provides information about the location of an action. Importantly, this is achieved without requiring any additional annotations.

More specifically, CAMs are extracted from the activations in the last convolutional layer in the following manner. Let $f_l(x, y)$ represent the activation of unit l in the last convolutional layer at spatial location (x, y) . A score S_k for each class k can be obtained by performing global average pooling [19] to obtain, for each unit l , a feature $F^l = \sum_{x,y} f_l(x, y)$, followed by a linear layer with weights $\{w_l^k\}$. That is, $S_k = \sum_l w_l^k F^l$. A CAM for class k at location (x, y) can then be computed as

$$M_k(x, y) = \sum_l w_l^k f_l(x, y), \quad (2)$$

which indicates the importance of the activations at location (x, y) in the final score for class k .

Here, we propose to make use of the CAMs to extract action-aware features. To this end, we use the CAMs in conjunction with the output of the conv5-3 layer of the model. The intuition behind this is that conv5-3 extracts high-level features that provide a very rich representation of the image [46] and typically correspond to the most discriminative parts of the object [1, 27], or, in our case, the action. Therefore, we incorporate a new layer to our sub-model, whose output can be expressed as

$$A_k(x, y) = \text{conv5-3}(x, y) \times \text{ReLU}(M_k(x, y)), \quad (3)$$

where $\text{ReLU}(M_k(x, y)) = \max(0, M_k(x, y))$. As shown in Fig. 4, this new layer is then followed by fully-connected ones, and we take our action-aware features as the output of the corresponding fc7 layer.

3.2.2 Sequence Learning for Action Anticipation

To effectively combine the information contained in the context-aware and action-aware features described above, we design the novel multi-stage LSTM model depicted by Fig. 2. This model first focuses on the context-aware features, which encode global information about the entire image. It then combines the output of this first stage with our action-aware features to provide a refined class prediction.

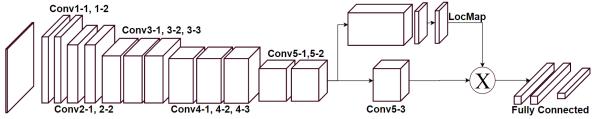


Figure 4. **Action-aware feature extraction.** Given the fine-tuned feature extraction network, we introduce a new layer that alters the output of conv5-3. This lets us filter out the conv5-3 features that are irrelevant, to focus on the action itself. Our action-aware features are taken as the output of the last fully-connected layer.

To train this model for action anticipation, we make use of our new loss introduced in Section 3.1. Therefore, ultimately, our network models long-range temporal information and yields increasingly accurate predictions as it processes more frames.

Specifically, we write the overall loss of our model as

$$\mathcal{L}_o = \frac{1}{V} \sum_{i=1}^V (\mathcal{L}_{c,i} + \mathcal{L}_{a,i}) , \quad (4)$$

where V is the total number of training sequences. This loss function combines losses corresponding to the context-aware stage and to the action-aware stage, respectively. Below, we discuss these two stages in more detail.

Learning Context. The first stage of our model takes as input our context-aware features, and passes them through a layer of LSTM cells followed by a fully-connected layer that, via a softmax operation, outputs a probability for each action class. Let $\hat{y}_{c,i}$ be the vector of probabilities for all classes and all time steps predicted by the first stage for sample i . We then define the loss for a single sample as

$$\mathcal{L}_{c,i} = \mathcal{L}(y_i, \hat{y}_{c,i}) , \quad (5)$$

where $\mathcal{L}(\cdot)$ is our new loss defined in Eq. 1, and y_i is the ground-truth class label for sample i .

Learning Context and Action. The second stage of our model aims at combining context-aware and action-aware information. Its structure is the same as that of the first stage, i.e., a layer of LSTM cells followed by a fully-connected layer to output class probabilities via a softmax operation. However, its input merges the output of the first stage with our action-aware features. This is achieved by concatenating the hidden activations of the LSTM layer with our action-aware features. We then make use of the same loss function as before, but defined on the final prediction. For sample i , this can be expressed as

$$\mathcal{L}_{a,i} = \mathcal{L}(y_i, \hat{y}_{a,i}) , \quad (6)$$

where \hat{y}_a is the vector of probabilities for all classes predicted by the second stage.

Inference. At inference, the input RGB frames are forward-propagated through our model. We therefore obtain a probability vector for each class at each frame. While one could simply take the probabilities in the current frame t to obtain the class label at time t , via *argmax*, we propose to increase robustness by leveraging the predictions of all the frames up to time t . To this end, we make use of an average pooling of these predictions over time.

4. Experiments

In this section, we first compare our method with state-of-the-art techniques on the task of action anticipation, and then analyze various aspects of our model, such as the influence of the loss function and of the different feature types. In the supplementary material, we provide additional experiments to analyze the effectiveness of different LSTM architectures, and the influence of the number of hidden units and of our temporal average pooling strategy. We also report the performance of our method on the task of action recognition from complete videos with and without optical flow, and action anticipation with optical flow.

4.1. Datasets

For our experiments, we made use of the standard UCF-101 [34], UT-Interaction [24], and JHMDB-21 [13] benchmarks, which we briefly describe below.

The UCF-101 dataset consists of 13,320 videos (each contains a single action) of 101 action classes including a broad set of activities such as sports, playing musical instruments and human-object interaction, with an average length of 7.2 seconds. UCF-101 is one of the most challenging datasets due to its large diversity in terms of actions and to the presence of large variations in camera motion, cluttered background and illumination conditions. There are three standard training/test splits for this dataset. In our comparisons to the state-of-the-art for both action anticipation and recognition, we report the average accuracy over the three splits. For the detailed analysis of our model, however, we rely on the first split only.

The JHMDB-21 dataset is another challenging dataset of realistic videos from various sources, such as movies and web videos, containing 928 videos and 21 action classes. Similarly to UCF-101, in our comparison to the state-of-the-art, we report the average accuracy over the three standard splits of data. Similar to UCF-101 dataset, each video contains one action starting from the beginning of the video.

The UT-Interaction dataset contains videos of continuous executions of 6 human-human interaction classes: shake-hands, point, hug, push, kick and punch. The dataset contains 20 video sequences whose length is about 1 minute each. Each video contains at least one execution of each interaction type, providing us with 8 executions of human activities per video on average. Following the recommended

experimental setup, we used 10-fold leave-one-out cross validation for each of the standard two sets of 10 videos. That is, within each set, we leave one sequence for testing and use the remaining 9 for training. Following standard practice, we also made use of the annotations provided with the dataset to split each video into sequences containing individual actions.

4.2. Implementation Details

CNN and LSTM Configuration. The parameters of the CNN were optimized using stochastic gradient descent with a fixed learning rate of 0.001, a momentum of 0.9, a weight decay of 0.0005, and mini-batches of size 32. To train our LSTMs, we similarly used stochastic gradient descent with a fixed learning rate of 0.001, a momentum of 0.9, and mini-batch size of 32. For all LSTMs, we used 2048 hidden units. To implement our method, we used Python and Keras [5]. We will make our code publicly available.

Training Procedure. To fine-tune the network on each dataset, we augment the data, so as to reduce the effect of over-fitting. The input images were randomly flipped horizontally and rotated by a random amount in the range -8 to 8 degrees. We then extracted crops according to the following procedure: **(1)** Compute the maximum cropping rectangle with a given aspect ratio (320/240) that fits inside the input image. **(2)** Scale the width and height of the cropping rectangle by a factor randomly selected in the range 0.8–1. **(3)** Select a random location for the cropping rectangle within the original input image and extract the corresponding subimage. **(4)** Scale the subimage to 224×224 .

After these geometric transformations, we further applied RGB channel shifting [43], followed by randomly adjusting image brightness, contrast and saturation with a factor $\alpha = 0.3$. The operations are: for brightness, $\alpha \times Image$, for contrast, $Image \times \alpha + (1.0 - \alpha) \times mean(grey(Image))$, and for saturation, $Image \times \alpha + (1.0 - \alpha) \times grey(Image)$.

4.3. Comparison to the State-of-the-Art

We compare our approach to the state-of-the-art action anticipation results reported on each of the three datasets discussed above. We further complement these state-of-the-art results with additional baselines that make use of our context-aware features with the loss of either [21] or [11]. Note that a detailed comparison of different losses within our model is provided in Section 4.4.1.

Following standard practice, we report the so-called *earliest* and *latest* prediction accuracies. Note, however, that there is no real agreement on the proportion of frames that the *earliest* setting corresponds to. For each dataset, we make use of the proportion that has been employed by the baselines (i.e., either 20% or 50%). Note also that our approach relies on at most T frames (with $T = 50$ in practice). Therefore, in the *latest* setting, where the baselines

Table 1. Comparison with state-of-the-art baselines on the task of action anticipation on the JHMDB-21 dataset. Note that our approach outperforms all baselines significantly in both settings.

Method	Earliest	Latest
DP-SVM [32]	5%	46%
S-SVM [32]	5%	43%
Where/What [33]	10%	43%
Ranking Loss [21]	29%	43%
Context-Aware+Loss of [11]	28 %	43%
Context-Aware+Loss of [21]	33%	39%
Ours	55%	58%

Table 2. Comparison with state-of-the-art baselines on the task of action anticipation on the UT-Interaction dataset.

Method	Earliest	Latest
D-BoW [25]	70.0%	85.0%
I-BoW [25]	65.0%	81.7%
CuboidSVM [24]	31.7%	85.0%
BP-SVM [17]	65.0%	83.3%
CuboidBayes [25]	25.0%	71.7%
DP-SVM [32]	13.0%	14.6%
S-SVM [32]	11.0%	13.4%
Context-Aware+Loss of [11]	45.0 %	65.0%
Context-Aware+Loss of [21]	48.0%	60.0%
Ours	84.0%	90.0%

Table 3. Action Anticipation on the UCF-101 dataset.

Method	Earliest	Latest
Context-Aware+Loss of [11]	30.6 %	71.1%
Context-Aware+Loss of [21]	22.6%	73.1%
Ours	80.5%	83.4%

rely on the complete sequences, we only exploit the first T frames. We believe that the fact that our method significantly outperforms the state-of-the-art in this setting despite using less information further evidences the effectiveness of our approach.

JHMDB-21. The results for the JHMDB-21 dataset are provided in Table 1. In this case, following the baselines, earliest prediction corresponds to observing the first 20% of the sequence. Note that we clearly outperform all the baselines by a significant margin in both the earliest and latest settings. Remarkably, we also outperform the methods that rely on additional information as input, such as optical flow [32, 33, 21] and Fisher vector features based on Improved Dense Trajectories [32]. This clearly demonstrates the benefits of our approach for anticipation.

UT-Interaction. We provide the results for the UT-Interaction dataset in Table 2. Here, following standard practice, 50% of the sequence was observed for earliest prediction, and the entire sequence for latest prediction. Recall that our approach uses at most $T = 50$ frames for prediction in both settings, while the average length of a complete sequence is around 120 frames. Therefore, as evidenced by the results, our approach yields significantly higher accuracy despite using considerably less data as input.

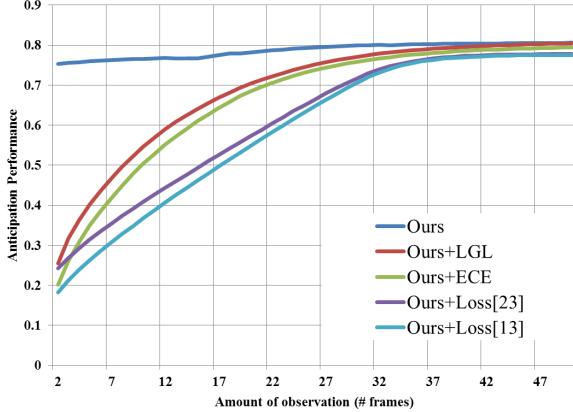


Figure 5. Comparison of different losses for action anticipation on UCF-101. We evaluate the accuracy of our model trained with different losses as a function of the number of frames observed. This plot clearly shows the superiority of our loss function.

UCF-101. We finally compare our approach with our two baselines on the UCF-101 dataset. While this is not a standard benchmark for action anticipation, this experiment is motivated by the fact that this dataset is relatively large, has many classes, with similarity across different classes, and contains variations in video capture conditions. Altogether, this makes it a challenging dataset to anticipate actions, especially when only a small amount of data is available. The results on this dataset are provided in Table 3. Here, the earliest setting corresponds to using the first 2 frames of the sequences, which corresponds to around 1% of the data. Again, we clearly outperform the two baselines consisting of exploiting context-aware features with the loss of either [21] or [11]. We believe that this further evidences the benefits of our approach, which leverages both context- and action-aware features with our new anticipation loss. A detailed evaluation of the influence of the different feature types and losses is provided in the next section.

4.4. Analysis

In this section, we provide a more detailed analysis of the influence of our loss function and of the different feature types on anticipation accuracy. Finally, we also provide a visualization of our different feature types, to illustrate their respective contributions.

4.4.1 Influence of the Loss Function

Throughout the paper, we have argued that our novel loss, introduced in Section 3.1, is better-suited to action anticipation than existing ones. To evaluate this, we trained several versions of our model with different losses. In particular, as already done in the comparison to the state-of-the-art above, we replaced our loss with the ranking loss of [21] (ranking loss on detection score) and the loss of [11], but this time

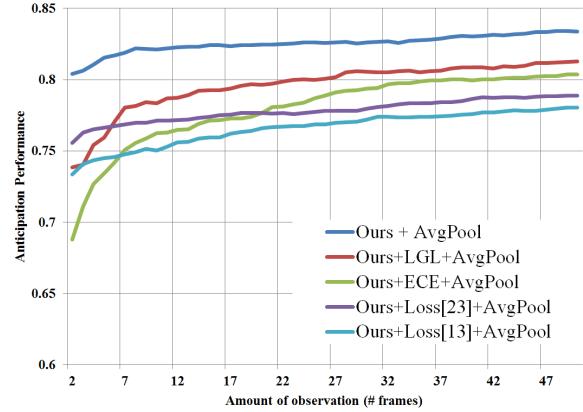


Figure 6. Influence of our average pooling strategy. Our simple yet effective average pooling leverages the predictions of all the frames up to time t . As shown on the UCF-101 dataset, this increases anticipation performance, especially at very early stages.

within our complete multi-stage model, with both context- and action-aware features.

Furthermore, we made use of the standard cross-entropy (**CE**) loss, which only accounts for one activity label for each sequence (at time T). This loss can be expressed as

$$\mathcal{L}_{CE} = \sum_{k=1}^N [y^T(k) \log(\hat{y}^T(k)) + (1 - y^T(k)) \log(1 - \hat{y}^T(k))] . \quad (7)$$

We then also modified the loss of [11], which consists of an exponentially weighted softmax, with an exponentially weighted cross-entropy loss (**ECE**), written as

$$\mathcal{L}_{ECE} = \sum_{t=1}^T -e^{-(T-t)} \sum_{k=1}^N [y^t(k) \log(\hat{y}^t(k)) + (1 - y^t(k)) \log(1 - \hat{y}^t(k))] . \quad (8)$$

The main drawback of this loss comes from the fact that it does not strongly encourage the model to make correct predictions as early as possible. To address this issue, we also introduce a linearly growing loss (**LGL**), defined as

$$\mathcal{L}_{LGL} = \sum_{t=1}^T -\frac{t}{T} \sum_{k=1}^N [y^t(k) \log(y^t(k)) + (1 - y^t(k)) \log(1 - \hat{y}^t(k))] . \quad (9)$$

While our new loss, introduced in Section 3.1, also makes use of a linearly-increasing term, it corresponds to the false positives in our case, as opposed to the false negatives in the LGL. Since some actions are ambiguous in the first few frames, we find it more intuitive not to penalize false positives too strongly at the beginning of the sequence. This intuition is supported by our results below, which show that our loss yields better results than the LGL.

Table 4. Importance of the different feature types.

Feature	Model	Earliest K=1	Latest K=50
Context-Aware	LSTM	62.80%	72.71%
Action-Aware	LSTM	69.33%	77.86%
Context+Action	MS-LSTM	80.5%	83.37%

In Fig. 5, we report the accuracy of the corresponding models as a function of the number of observed frames on the UCF-101 dataset. Note that our new loss yields much higher accuracies than the other ones, particularly when only a few frames of the sequence are observed; With only 2 frames observed, our loss yields an accuracy similar to the other losses with 30–40 frames. With 30fps, this essentially means that we can predict the action 1 second earlier than other methods. The importance of this result is exemplified by research showing that a large proportion of vehicle accidents are due to mistakes/misinterpretations of the scene in the immediate time leading up to the crash [11, 23].

Moreover, in Fig. 6, we report the performance of the corresponding models as a function of the number of observed frames when using our average pooling strategy. Note that this strategy can generally be applied to any action anticipation method and, as shown by comparing Figs. 5 and 6, increases accuracy, especially at very early stages, which clearly demonstrates its effectiveness. Note that using it in conjunction with our loss still yields the best results by a significant margin.

4.4.2 Influence of the Features

We then evaluate the importance of the different feature types, context-aware and action-aware, on anticipation accuracy. To this end, we compare models trained using each feature type individually with our model that uses them jointly. For the models using a single feature type, we made use of a single LSTM to model temporal information. By contrast, our approach relies on a multi-stage LSTM, which we denote by *MS-LSTM*. Note that all models were trained using our new anticipation loss. The results of this experiment on the UCF-101 dataset are provided in Table 4. These results clearly evidence the importance of using both feature types, which consistently outperforms individual ones.

Since we extract action-aware features, and not motion-aware ones, our approach will not be affected by irrelevant motion. The CNN that extracts these features learns to focus on the discriminative parts of the images, thus discarding the irrelevant information. To confirm this, we conducted an experiment on some classes of UCF-101 that contain irrelevant motions/multiple actors, such as Baseball pitch, Basketball, Cricket Shot and Ice dancing. The results of our action-aware and context-aware frameworks for these classes are: 66.1% vs. 58.2% for Baseball pitch, 83% vs. 76.4% for Basketball, 65.1% vs. 58% for Cricket Shot,

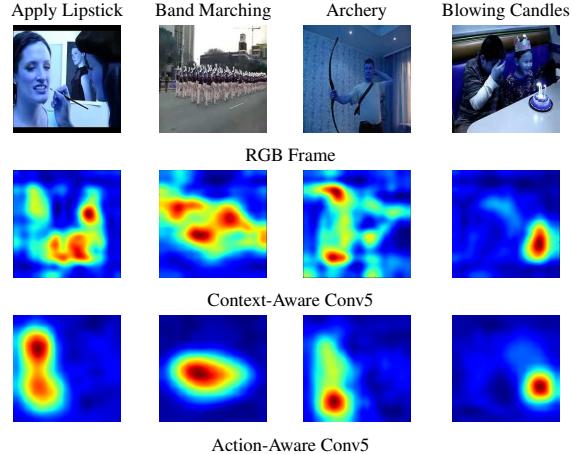


Figure 7. Visualization of action-aware and context-aware features on UCF-101. These samples are representative of the data.

and 92.3% vs. 91.7% for Ice dancing. This shows that our action-aware features can effectively discard irrelevant motion/actors to focus on the relevant one(s).

4.4.3 Visualization

Finally, we provide a better intuition of the kind of information each of our feature types encode (see Fig. 7). This visualization was computed by average pooling over the 512 channels of Conv5-3 (of both the context-aware and action-aware sub-networks). As can be observed in the figure, our context-aware features have high activations on regions corresponding to any relevant object in the scene (context). By contrast, in our action-aware features, high activations correctly correspond to the focus of the action. Therefore, they can reasonably localize the parts of the frame that most strongly participate in the action happening in the video and reduce the noise coming from context.

5. Conclusion

In this paper, we have introduced a novel loss function to address very early action anticipation. Our loss encourages the model to make correct predictions as early as possible in the input sequence, thus making it particularly well-suited to action anticipation. Furthermore, we have introduced a new multi-stage LSTM model that effectively combines context-aware and action-aware features. Our experiments have evidenced the benefits of our new loss function over existing ones. Furthermore, they have shown the importance of exploiting both context- and action-aware information. Altogether, our approach significantly outperforms the state-of-the-art in action anticipation on all the datasets we applied it to. In the future, we intend to study new ways to incorporate additional sources of information, such as dense trajectories and human skeletons in our framework.

References

- [1] G. Bertasius, J. Shi, and L. Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4380–4389, 2015.
- [2] H. Bilen, B. Fernando, E. Gavves, A. Vedaldi, and S. Gould. Dynamic image networks for action recognition. In *IEEE International Conference on Computer Vision and Pattern Recognition CVPR*, 2016.
- [3] W. Chen, C. Xiong, R. Xu, and J. J. Corso. Actionness ranking with lattice conditional ordinal random fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 748–755, 2014.
- [4] G. Chéron, I. Laptev, and C. Schmid. P-cnn: Pose-based cnn features for action recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3218–3226, 2015.
- [5] F. Chollet. keras. <https://github.com/fchollet/keras>, 2015.
- [6] A. Diba, A. Mohammad Pazandeh, H. Pirsiavash, and L. Van Gool. Deepcamp: Deep convolutional action & attribute mid-level patterns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3557–3565, 2016.
- [7] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2625–2634, 2015.
- [8] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [9] B. Fernando, P. Anderson, M. Hutter, and S. Gould. Discriminative hierarchical rank pooling for activity recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [10] B. Fernando, E. Gavves, J. Oramas, A. Ghodrati, and T. Tuytelaars. Rank pooling for action recognition. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.
- [11] A. Jain, A. Singh, H. S. Koppula, S. Soh, and A. Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3118–3125. IEEE, 2016.
- [12] M. Jain, J. Van Gemert, H. Jégou, P. Bouthemy, and C. G. Snoek. Action localization with tubelets from motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 740–747, 2014.
- [13] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black. Towards understanding action recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3192–3199, 2013.
- [14] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.
- [15] Y. Kong, D. Kit, and Y. Fu. A discriminative model with multiple temporal scales for action prediction. In *European Conference on Computer Vision*, pages 596–611. Springer, 2014.
- [16] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005.
- [17] K. Laviers, G. Sukthankar, D. W. Aha, M. Molineaux, C. Darken, et al. Improving offensive performance through opponent modeling. In *Proceedings of AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 58–63.
- [18] Y. Li, W. Li, V. Mahadevan, and N. Vasconcelos. Vlad3: Encoding dynamics of deep features for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1951–1960, 2016.
- [19] M. Lin, Q. Chen, and S. Yan. Network in network. In *arXiv preprint arXiv:1312.4400*, 2013.
- [20] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *European Conference on Computer Vision*, pages 816–833. Springer, 2016.
- [21] S. Ma, L. Sigal, and S. Sclaroff. Learning activity progression in lstms for activity detection and early detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1942–1950, 2016.
- [22] D. Oneata, J. Revaud, J. Verbeek, and C. Schmid. Spatio-temporal object detection proposals. In *European Conference on Computer Vision*, pages 737–752. Springer, 2014.
- [23] P. Rueda-Domingo, T. Lardelli-Claret, J. Luna del Castillo, J. Jimenez-Moleon, M. Garcia-Martín, and A. Bueno-Cavanillas. The influence of passengers on the risk of the driver causing a car collision in spain: Analysis of collisions from 1990 to 1999. *Accident Analysis and Prevention*, 2004.
- [24] M. Ryoo, C.-C. Chen, J. Aggarwal, and A. Roy-Chowdhury. An overview of contest on semantic description of human activities (sdha) 2010. In *Recognizing Patterns in Signals, Speech, Images and Videos*, pages 270–285. Springer, 2010.
- [25] M. S. Ryoo. Human activity prediction: Early recognition of ongoing activities from streaming videos. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 1036–1043. IEEE, 2011.
- [26] M. S. Ryoo and J. K. Aggarwal. Spatio-temporal relationship match: Video structure comparison for recognition of complex human activities. In *Computer vision, 2009 ieee 12th international conference on*, pages 1593–1600. IEEE, 2009.
- [27] F. Saleh, M. S. Aliakbarian, M. Salzmann, L. Petersson, S. Gould, and J. M. Alvarez. Built-in foreground/background prior for weakly-supervised semantic segmentation. In *European Conference on Computer Vision*, pages 413–432. Springer, 2016.
- [28] 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.

- [29] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *arXiv preprint arXiv:1409.1556*, 2014.
- [30] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1961–1970, 2016.
- [31] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1961–1970, 2016.
- [32] K. Soomro, H. Idrees, and M. Shah. Online localization and prediction of actions and interactions. *arXiv preprint arXiv:1612.01194*, 2016.
- [33] K. Soomro, H. Idrees, and M. Shah. Predicting the where and what of actors and actions through online action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2648–2657, 2016.
- [34] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. 2012.
- [35] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4489–4497. IEEE, 2015.
- [36] M. Van den Bergh, G. Roig, X. Boix, S. Manen, and L. Van Gool. Online video seeds for temporal window objectness. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 377–384, 2013.
- [37] C. Vondrick, H. Pirsiavash, and A. Torralba. Anticipating visual representations from unlabeled video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 98–106, 2016.
- [38] H. Wang and C. Schmid. Action recognition with improved trajectories. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3551–3558, 2013.
- [39] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4305–4314, 2015.
- [40] L. Wang, Y. Qiao, X. Tang, and L. Van Gool. Actionness estimation using hybrid fully convolutional networks. *arXiv preprint arXiv:1604.07279*, 2016.
- [41] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: towards good practices for deep action recognition. In *European Conference on Computer Vision*, pages 20–36. Springer, 2016.
- [42] X. Wang, A. Farhadi, and A. Gupta. Actions ~ transformations. In *IEEE International Conference on Computer Vision and Pattern Recognition CVPR*. IEEE, 2016.
- [43] R. Wu, S. Yan, Y. Shan, Q. Dang, and G. Sun. Deep image: Scaling up image recognition. *arXiv preprint arXiv:1501.02876*, 7(8), 2015.
- [44] G. Yu and J. Yuan. Fast action proposals for human action detection and search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1302–1311, 2015.
- [45] G. Yu, J. Yuan, and Z. Liu. Predicting human activities using spatio-temporal structure of interest points. In *Proceedings of the 20th ACM international conference on Multimedia*, pages 1049–1052. ACM, 2012.
- [46] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014.
- [47] X. Zhao, X. Li, C. Pang, X. Zhu, and Q. Z. Sheng. Online human gesture recognition from motion data streams. In *Proceedings of the 21st ACM international conference on Multimedia*, pages 23–32. ACM, 2013.
- [48] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *IEEE International Conference on Computer Vision and Pattern Recognition CVPR*, 2016.