

Long-Term Anticipation of Activities with Cycle Consistency

Yazan Abu Farha¹, QiuHong Ke², Bernt Schiele³, and Juergen Gall¹

¹University of Bonn, Germany

²The University of Melbourne, Australia

³MPI Informatics, Germany

Abstract. With the success of deep learning methods in analyzing activities in videos, more attention has recently been focused towards anticipating future activities. However, most of the work on anticipation either analyzes a partially observed activity or predicts the next action class. Recently, new approaches have been proposed to extend the prediction horizon up to several minutes in the future and that anticipate a sequence of future activities including their durations. While these works decouple the semantic interpretation of the observed sequence from the anticipation task, we propose a framework for anticipating future activities directly from the features of the observed frames and train it in an end-to-end fashion. Furthermore, we introduce a cycle consistency loss over time by predicting the past activities given the predicted future. Our framework achieves state-of-the-art results on two datasets: the Breakfast dataset and 50Salads.

1 Introduction

Humans spend a significant time of their life thinking about the future. Whether thinking about their future dream job, or planning for the next research project. Even unconsciously, people tend to anticipate future trajectories of moving agents in the surrounding environment and the activities that they will be doing in the near future. Such anticipation capability is considered a sign of intelligence and an important factor in determining how we interact with the environment and how we make decisions.

Since anticipation is an important intrinsic capability of human beings, researchers have recently tried to model this capability and embed it in intelligent and robotic systems. For example, several approaches have been proposed to anticipate future trajectories of pedestrians [18, 4], or semantic segmentation of future frames in video [25, 5]. These approaches have many applications in autonomous driving and navigation. Another line of research focuses on anticipating future activities [21, 19, 44], which has potential applications in surveillance and human-robot interaction.

While anticipating the next action a few seconds in the future has been addressed in [44, 46, 11], such short time horizon is insufficient for many applications. Service robots, for example, where a robot is continuously interacting




Fig. 1. (Left) Overview of the proposed approach, which is trained end-to-end and includes a cycle consistency module. (Right) Effect of the cycle consistency module. Without the cycle consistency, the network anticipates actions that are plausible based on the previous actions. However, in some cases an essential action is missing. In this example *pour oil*. By using the cycle consistency, we enforce the network to verify if all required actions have been done before. For the action *fry pancake*, *pouring oil* into the pan is required and the cycle consistency resolves this issue.

with a human, require anticipating a longer time horizon, which rather includes a sequence of future activities than only predicting the next action. By anticipating longer in the future, such robots can plan ahead and accomplish their tasks more efficiently. Recent approaches, therefore, focused on increasing the prediction horizon up to several minutes and predict multiple action segments in the future [3, 10, 16]. While these approaches successfully predict future activities and their duration, they decouple the semantic interpretation of the observed sequence from the anticipation task using a two-step approach.

Separating the understanding of the past and the anticipation of future has several disadvantages. First, the model is not trained end-to-end, which means that the approach for temporal action segmentation is not optimized for the anticipation task. Second, if there are any mistakes in the temporal action segmentation, these mistakes will be propagated and effect the anticipated activities. Finally, the action labels do not represent all information in the observed video that is relevant for anticipating the future. In contrast to these approaches, we propose a sequence-to-sequence model that directly maps the sequence of observed frames to a sequence of future activities and their duration. We then cast the understanding of the past as an auxiliary task by proposing a recognition module, which consists of a temporal convolutional neural network and a recognition loss, that is combined with the encoder.

Furthermore, as we humans can reason about the past given the future, previous works only aim to predict the future, and it is intuitive that forcing the network to predict the past as well is helpful. To this end, we propose a cycle consistency module that predicts the past activities given the predicted future. This module verifies if, for the predicted future actions, all required actions have been done before as illustrated in Fig. 1. In this example, the actions *pour dough to pan* and *fry pancake* are plausible given the previous actions, but the action *pour oil* has been missed. The cycle consistency module then predicts from the anticipated actions, the observed actions. Since *pour dough to pan* and

fry pancake are the inputs for the cycle consistency module, it will predict all the required preceding actions such as *pour milk*, *stir dough* and *pour oil*. However, as *pour oil* is not part of the observations, the cycle consistency module will have a high error, which steers the network to predict *pour oil* in the future actions.

Our contribution is thus three folded: First, we propose an end-to-end model for anticipating a sequence of future activities and their durations. Second, we show that the proposed recognition module improves the sequence-to-sequence model. Third, we propose a cycle consistency module that verifies the predictions.

We evaluate our model on two datasets with untrimmed videos containing many action segments: the Breakfast dataset [20] and 50Salads [41]. Our model is able to predict the future activities and their duration accurately achieving superior results compared to the state-of-the-art methods.

2 Related Work

Early action detection: The early action detection task tries to recognize an ongoing activity given only a partial observation of that activity. An initial work addressing this task by Ryoo [36] is based on a probabilistic formulation using dynamic bag-of-words of spatio-temporal features. Hoai and De la Torre [14] used a max-margin formulation. More recent approaches use recurrent neural networks with special loss functions to encourage early activity detection [26, 37]. In contrast to these approaches, we anticipate future activities without even any partial observations.

Future prediction: Predicting the future has become one of the main research topics in computer vision. Many approaches have been proposed to predict future frames [30, 24], future human trajectories [18, 4], future human poses [15, 29, 35], image semantic segmentation [25, 5], or even full sentences describing future frames or steps in recipes [38, 27]. However, low level representations like pixels of frames or very high level natural language sentences cannot be used directly. A complementary research direction is concentrated on anticipating activity labels. Lan *et al.* [21] predicted future actions using hierarchical representations in a max-margin framework. Koppula and Saxena [19] used a spatio-temporal graph representation to encode the observed activities and then predict object affordances, trajectories, and sub-activities. Instead of directly predicting the future action labels, several approaches were proposed to predict future visual representations and then a classifier is trained on top of the predicted representations to predict the future labels [44, 12, 46, 39, 34, 11]. Predicting future representations has also been used in the literature for unsupervised representations learning [40]. Other approaches use multi-task learning to predict the future activity and its starting time [28, 31], or the future activity and its location [23, 42]. Miech *et al.* [32] modeled the transition probabilities between actions and combine it with a predictive model to anticipate future action labels. There is a parallel line of research addressing the anticipation task in egocentric videos [8, 9]. However, these approaches are limited to very few seconds in the future. In

contrast to these approaches, we address the anticipation task for a longer time horizon.

Recently, more effort has been dedicated to increase the anticipation horizon. Several methods have been proposed to anticipate activities several minutes in the future using RNNs [2, 3], temporal convolutions with time-variable [16], or memory networks [10]. While these approaches manage to anticipate activities for a longer time horizon, their performance is limited. The approach of [10] relies on the ground-truth action labels of the observations, whereas the methods in [2, 3, 16] follow a two-step approach. *I.e.* they first infer the activities in the observed part, then anticipate the future activities with their corresponding duration. These two steps are trained separately, which prevents the model from utilizing the visual cues in the observed frames. In contrast to these approaches, our model is trained in one step in an end-to-end fashion.

Cycle consistency: Cycle consistency has been widely used in computer vision. It was used to learn dense correspondence [47], image-to-image translation [48], and depth estimation [13]. Recent approaches used cycle consistency in the temporal dimension [7, 45]. Dwibedi *et al.* [7] introduced an approach to learn representations using video alignment as a proxy task and cycle consistency for training. In [13], the appearance consistency between consecutive video frames is used to learn representations that generalize to different tasks. Motivated by the success of cycle consistency in various applications, we apply it as an additional supervisory signal to predict the future activities.

3 The Anticipation Framework

Given a partially observed video with many activities, we want to predict all the activities that will be happening in the remainder of that video with their corresponding duration. Assuming that the observed part consists of t_o frames $X_{1:t_o} = (x_1, \dots, x_{t_o})$ corresponding to n activities $A_{1:n} = (A_1, \dots, A_n)$, our goal is to predict the future activities $A_{n+1:N} = (A_{n+1}, \dots, A_N)$ and their corresponding duration $\ell_{n+1:N} = (\ell_{n+1}, \dots, \ell_N)$, where N is the total number of activities in that video. In contrast to the previous approaches that use only the action labels of the observed frames for anticipating the future, we propose to anticipate the future activities directly from the observed frames. First, we propose a sequence-to-sequence model that maps the sequence of features from the observations to a sequence of future activities and their duration. Then, we introduce a cycle consistency module that predicts the past activities given the predicted future. The motivation of this module is to force the sequence-to-sequence module to encode all the relevant information in the observed frames and verify if the predictions are plausible. Finally, we extend the sequence-to-sequence model with a recognition module that generates discriminative features that capture the relevant information for anticipating the future. An overview of the proposed model is illustrated in Fig. 2. Our framework is hence trained using an anticipation loss, a recognition loss, and a cycle consistency loss. In the following sections, we describe in detail these modules.

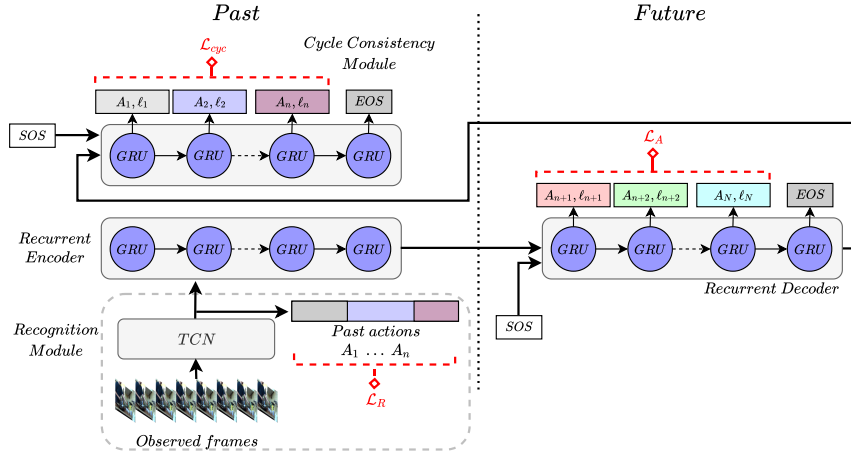


Fig. 2. Overview of the anticipation framework. The observed frames are passed through a TCN-based recognition module which produces discriminative features for the sequence-to-sequence model. The sequence-to-sequence model predicts the future activities with their duration. In addition, we enforce cycle consistency over time by predicting the past activities given the predicted future.

3.1 Sequence-to-Sequence Model

The sequence-to-sequence model maps the sequence of observed frames to a sequence of future activities and their duration. For this model, we use a recurrent encoder-decoder architecture based on gated recurrent units (GRUs).

Recurrent Encoder. The purpose of the recurrent encoder is to encode the observed frames in a single vector which will be used to decode the future activities. Given the input features X , the recurrent encoder passes these features through a single layer with a gated recurrent unit (GRU)

$$h_t^e = GRU(x_t, h_{t-1}^e), \quad (1)$$

where x_t is the input feature for frame t , and h_{t-1}^e is the hidden state at the previous time step. The hidden state at the last time step $h_{t_o}^e$ encodes the observed frames and will be used as an initial state for the decoder.

Recurrent Decoder. Given the output of the encoder $h_{t_o}^e$, the recurrent decoder predicts the future activities and their relative duration. The decoder also consists of a single layer with a gated recurrent unit (GRU). The hidden state at each step is updated using the GRU update rules

$$h_m^d = GRU(A_{m-1}, h_{m-1}^d), \quad (2)$$

where the input at each time step A_{m-1} is the predicted activity label for the previous step. At training time, the ground-truth label is used as input. Whereas the predicted label is used during inference time. For the first time step, a special start-of-sequence (SOS) symbol is used as input. Given the hidden state h_m^d at

each time step, the future activity label A_m and its relative duration ℓ_m are predicted using a fully connected layer, *i.e.*

$$\tilde{A}_m = W_A h_m^d + b_A, \quad (3)$$

$$\hat{\ell}_m = W_\ell h_m^d + b_\ell, \quad (4)$$

where \tilde{A}_m is the predicted logits for the future activity, and $\hat{\ell}_m$ is the predicted duration in log space. To get the relative duration, we apply softmax over the time steps

$$\tilde{\ell}_m = \frac{e^{\hat{\ell}_m}}{\sum_k e^{\hat{\ell}_k}}. \quad (5)$$

The decoder keeps predicting future activity labels and the corresponding duration until a special end-of-sequence (EOS) symbol is predicted. As a loss function, we use a combination of a cross entropy loss for the activity label and a mean squared error (MSE) for the predicted relative duration

$$\mathcal{L}_A = \mathcal{L}_{CE} + \mathcal{L}_{MSE}, \quad (6)$$

$$\mathcal{L}_{CE} = \frac{1}{N-n} \sum_{m=n+1}^N -\log(\tilde{a}_{m,A}), \quad (7)$$

$$\mathcal{L}_{MSE} = \frac{1}{N-n} \sum_{m=n+1}^N (\tilde{\ell}_m - \ell_m)^2, \quad (8)$$

where \mathcal{L}_A is the anticipation loss, $\tilde{a}_{m,A}$ is the the predicted probability for the ground truth activity label at step m , n is the number of observed action segments and N is the total number of action segments in the video.

Since the input to the encoder are frame-wise features which might be very long, the output of the encoder $h_{t_o}^e$ might not be able to capture all the relevant information. To alleviate this problem we combine the decoder with an attention mechanism using a multi-head attention module [43]. Additional details are given in the supplementary material.

3.2 Cycle Consistency

Since predicting the future from the past and the past from the future should be consistent, we propose an additional cycle consistency loss. Given the predicted future activities, we want to predict the past activities and their duration. This requires the predicted future to be good enough to predict the past. The cycle consistency loss has two benefits. First, it verifies if, for the predicted future actions, the actions that have been previously observed are plausible. Second, it encourages the recurrent decoder to keep the most important information of the observed sequence until the end, instead of storing only the information of the previous anticipated activity. In this way, a wrong prediction does not necessary propagate since the observed sequence is kept in memory.

The cycle consistency module is similar to the recurrent decoder and consists of a single layer with GRU, however, it predicts the past instead of the future. The hidden state of this GRU is initialized with the last hidden state of the recurrent decoder and at each step the hidden state is updated as follows

$$h_m^{cyc} = GRU(A_{m-1}, h_{m-1}^{cyc}). \quad (9)$$

Given the hidden state h_m^{cyc} at step m , an activity label of the observations and its relative duration are predicted using a fully connected layer. The loss function is also similar to the recurrent decoder

$$\mathcal{L}_{cyc} = \mathcal{L}_{CE} + \mathcal{L}_{MSE}, \quad (10)$$

$$\mathcal{L}_{CE} = \frac{1}{n} \sum_{m=1}^n -\log(\tilde{a}_{m,A}), \quad (11)$$

$$\mathcal{L}_{MSE} = \frac{1}{n} \sum_{m=1}^n (\tilde{\ell}_m - \ell_m)^2, \quad (12)$$

where \mathcal{L}_{cyc} is the cycle consistency loss, \mathcal{L}_{CE} and \mathcal{L}_{MSE} are the cross entropy loss and MSE loss applied on the past activity labels and their relative duration.

While the mapping from past to future can be multi-modal, this does not limit the applicability of the cycle consistency module. Since the cycle consistency module is conditioned on the predicted future, no matter what mode is predicted, the cycle consistency makes sure it is plausible. This also applies to the inverse mapping. As there is a path from the observed frames to the cycle consistency module through the sequence-to-sequence model, there is no ambiguity in which past activities have been observed.

3.3 Recognition Module

In the sequence-to-sequence model, the input of the recurrent encoder are the frame-wise features. However, directly passing the features to the encoder is sub-optimal as the encoder might struggle to capture all the relevant information for anticipating the future activities. As past activities provide a strong signal for anticipating future activities, we use a recognition module that learns discriminative features of the observed frames. These features will then serve as an input for the sequence-to-sequence model to anticipate the future activities. Given the success of temporal convolutional networks (TCNs) in analyzing activities in videos [22, 1], we use a similar network for our recognition module. Besides being a strong model for analyzing videos, TCNs are fully differentiable and can be integrated in our framework without preventing end-to-end training. For our module, we use a TCN similar to the one proposed in [1]. The TCN consists of several layers of dilated 1D convolutions where the dilation factor is doubled at each layer. The operations at each layer can be formally described as follows

$$\hat{F}_l = ReLU(W_1 * F_{l-1} + b_1), \quad (13)$$

$$F_l = F_{l-1} + W_2 * \hat{F}_l + b_2, \quad (14)$$

where F_l is the output of layer l , $*$ is the convolution operator, $W_1 \in \mathbb{R}^{3 \times K \times K}$ are the weights of the dilated convolutional filters with kernel size 3 and K is the number of the filters, $W_2 \in \mathbb{R}^{1 \times K \times K}$ are the weights of a 1×1 convolution, and $b_1, b_2 \in \mathbb{R}^K$ are bias vectors. The input of the first layer F_0 is obtained by applying a 1×1 convolution over the input features X .

The output of the last dilated convolutional layer serves as input to the subsequent modules. To make sure that these features are discriminative enough, we add a classification layer that predicts the action label at each observed frame

$$\tilde{Y}_t = \text{Softmax}(Wf_{L,t} + b), \quad (15)$$

where \tilde{Y}_t contains the class probabilities at time t , $f_{L,t} \in \mathbb{R}^K$ is the output of the last dilated convolutional layer at time t , $W \in \mathbb{R}^{C \times K}$ and $b \in \mathbb{R}^C$ are the weights and bias for the 1×1 convolutional layer, where C is the number of action classes. To train this module we use a cross entropy loss

$$\mathcal{L}_R = \frac{1}{t_o} \sum_{t=1}^{t_o} -\log(\tilde{y}_{t,c}), \quad (16)$$

where $\tilde{y}_{t,c}$ is the predicted probability for the ground truth label c at time t , and t_o is the number of observed frames.

3.4 Loss Function

To train our framework, we sum up all the three mentioned losses

$$\mathcal{L} = \mathcal{L}_A + \mathcal{L}_R + \mathcal{L}_{cyc}, \quad (17)$$

where \mathcal{L}_A is the anticipation loss, \mathcal{L}_R is the recognition loss, and \mathcal{L}_{cyc} is the cycle consistency loss.

4 Experiments

We evaluate the proposed model on two datasets: the Breakfast dataset [20] and 50Salads [41]. In all experiments, we report the average of three runs.

The **Breakfast** dataset is a collection of 1,712 videos with overall 66.7 hours and roughly 3.6 million frames. Each video belongs to one out of ten breakfast related activities, such as make tea or pancakes. The video frames are annotated with fine-grained action labels like *pour milk* or *fry egg*. Overall, there are 48 different actions. On average, each video contains 6 action instances and is 2.3 minutes long. For evaluation, we use the standard 4 splits as proposed in [20] and report the average.

The **50Salads** dataset contains 50 videos showing people preparing different kinds of salad. These videos are relatively long with an average of 6.4 minutes and 20 action instances per video. The video frames are annotated with 17 fine-grained action labels like *cut tomato* or *peel cucumber*. For evaluation, we use five-fold cross-validation and report the average as in [41].

We follow the state-of-the-art evaluation protocol and report the mean over classes (MoC) accuracy for different observation/prediction percentages [3, 16].

Table 1. Ablation study on the Breakfast dataset. Numbers represent mean over classes (MoC) accuracy.

Observation %	20%				30%			
	10%	20%	30%	50%	10%	20%	30%	50%
Prediction %								
S2S	23.22	20.92	20.10	20.05	26.43	24.38	24.13	23.38
S2S + TCN	14.52	13.55	13.15	12.70	14.83	14.11	13.61	13.01
S2S + TCN + \mathcal{L}_R	24.72	22.43	21.70	21.77	28.35	26.29	25.01	24.47
S2S + TCN + \mathcal{L}_R + \mathcal{L}_{cyc}	25.16	22.73	22.22	22.01	28.07	26.25	25.12	24.81
S2S + TCN + \mathcal{L}_R + \mathcal{L}_{cyc} + attn.	25.88	23.42	22.42	21.54	29.66	27.37	25.58	25.20

Implementation Details. For the recognition module, we used a TCN with 10 layers and 64 filters in each layer. The number of units in the GRU cells is set to 512. For each training video, we generate two training examples with 20% and 30% observation percentage. The prediction percentage is always set to 50%. All the models are trained for 80 epochs using Adam optimizer [17]. We set the learning rate to 0.001 and reduce it every 20 epochs with a factor of 0.8. For both datasets, we extract I3D [6] features for the video frames using both RGB and flow streams and sub-sample them at 5 frames per second.

4.1 Ablation Analysis

In this section, we analyze the impact of the different modules in our framework on the anticipation performance. This analysis is conducted on the Breakfast dataset and the results are shown in Table 1. Additional ablation experiments to study the impact of the input to the recurrent encoder and decoder are provided in the supplementary material.

Impact of the recognition module: The recognition module consists of two parts: a TCN and a recognition loss \mathcal{L}_R . Starting from only the sequence-to-sequence module (S2S), we can achieve a good accuracy in the range 20% – 27%. By combining the sequence-to-sequence module with the recognition module (S2S + TCN + \mathcal{L}_R), we gain an improvement of 1% – 2% for each observation-prediction percentage. This indicates that recognition helps the anticipation task. We also evaluate the performance when the TCN is combined with the sequence-to-sequence module without the recognition loss (S2S + TCN). As shown in Table 1, the results are worse than using only the sequence-to-sequence module if we do not apply the recognition loss. This can be explained by the structure of the network. The recurrent encoder maps the features extracted by the TCN from all frames to a single vector. Without additional loss for the recognition module, the gradient vanishes and the parameters of the TCN are not well estimated. Nevertheless, by just applying the recognition loss, we get enough supervisory signal to train the TCN and improve the overall anticipation accuracy. This also highlights that the improvements from the recognition module are due to the additional recognition task and not because of having more parameters.




Fig. 3. Impact of the cycle consistency loss. Cycle consistency verifies if, for the predicted future actions, all required actions have been done before and no essential action is missing (a), and it further encourages the decoder to keep the important information from the observations until the end, which results in better predictions (b).

Table 2. Comparison between the two-step approach and ours on the Breakfast dataset. Numbers represent mean over classes (MoC) accuracy.

Observation %	20%				30%				
	Prediction %	10%	20%	30%	50%	10%	20%	30%	50%
Two-Step		23.58	21.90	20.80	20.18	27.80	25.21	23.32	22.96
Ours		25.88	23.42	22.42	21.54	29.66	27.37	25.58	25.20

Impact of the cycle consistency loss: The cycle consistency module predicts the past activities from the predicted future. The intuition is that to be able to predict the past activities, the predicted future activities have to be correct. As shown in Table 1, using the cycle consistency loss gives an additional improvement on the anticipation accuracy. The cycle consistency loss verifies if, for the predicted future actions, all required actions have been done before and no essential action is missing. For example in Fig. 3 (a), the model observes *spoon flour*, *crack egg*, *pour milk*, and *stir dough*. Without the cycle consistency the network did not predict the action *pour oil*, which is required to *fry pancake*. By using cycle consistency this issue is resolved. Cycle consistency also forces the decoder to remember all observed activities. As illustrated in Fig. 3 (b), the model observes *spoon flour*, *crack egg*, *pour milk*, *butter pan*, and *stir dough*. Without the cycle consistency, the network predicts the action *stir fry egg*, which would have been plausible if *spoon flour* and *pour milk* were not part of the observations. Since the cycle consistency encourages the decoder to use all observed actions for anticipation, the activity *fry pancake* is correctly anticipated.

Impact of the attention module: Finally, using the full model by combining the recurrent decoder with the multi-head attention module further improves the results by roughly 1%. As shown in Table 1, the gain from using the attention module is higher when the observation percentage is 30%. This is mainly because of the encoder module. Given the observed frames, the encoder tries to encode them in a single vector. This means that the encoder has to throw away more information from the long sequences compared to shorter sequences. In this case, the attention module can help in capturing some of the lost information in the encoder output by attending on the relevant information in the observations.

Table 3. Comparison with the state-of-the-art. Numbers represent MoC accuracy.

Observation %	20%				30%			
	10%	20%	30%	50%	10%	20%	30%	50%
Breakfast								
RNN model [3]	18.11	17.20	15.94	15.81	21.64	20.02	19.73	19.21
CNN model [3]	17.90	16.35	15.37	14.54	22.44	20.12	19.69	18.76
RNN [3] + TCN	05.93	05.68	05.52	05.11	08.87	08.90	07.62	07.69
CNN [3] + TCN	09.85	09.17	09.06	08.87	17.59	17.13	16.13	14.42
UAAA (mode) [2]	16.71	15.40	14.47	14.20	20.73	18.27	18.42	16.86
Time-Cond. [16]	18.41	17.21	16.42	15.84	22.75	20.44	19.64	19.75
Ours	25.88	23.42	22.42	21.54	29.66	27.37	25.58	25.20
50Salads								
RNN model [3]	30.06	25.43	18.74	13.49	30.77	17.19	14.79	09.77
CNN model [3]	21.24	19.03	15.98	09.87	29.14	20.14	17.46	10.86
RNN [3] + TCN	32.31	25.51	19.10	14.15	26.14	17.69	16.33	12.97
CNN [3] + TCN	16.02	14.68	12.09	09.89	19.23	14.68	13.18	11.20
UAAA (mode) [2]	24.86	22.37	19.88	12.82	29.10	20.50	15.28	12.31
Time-Cond. [16]	32.51	27.61	21.26	15.99	35.12	27.05	22.05	15.59
Ours	34.76	28.41	21.82	15.25	34.39	23.70	18.95	15.89

4.2 End-to-End vs. Two-Step Approach

To illustrate the benefits of end-to-end learning over two-step approaches, we compare our framework with its two-step counterpart. For this comparison, we first train the recognition module from our framework and then fix the weights of the TCN and train the remaining components of our model with the anticipation loss and the cycle consistency loss. Table 2 shows the results of our framework compared to the two-step approach on the Breakfast dataset with different observation and prediction percentages. As shown in the table, our framework outperforms the two-step approach with a large margin of up to 2.3%. This highlights the benefits of end-to-end approaches where the model can capture the relevant information in the observed frames to anticipate the future. On the contrary, two-step approaches can only utilize the label information of the observed frames that are not optimized for the anticipation task which is sub-optimal.

4.3 Comparison with the State-of-the-Art

In this section, we compare our framework with the state-of-the-art methods on both the Breakfast dataset and 50Salads. We follow the same protocol and report results for different observation and prediction percentages. Table 3 shows the results on both datasets. All the previous approaches follow the two-step approach by inferring the action labels of the observed frames first and then use these labels to anticipate the future activities. As shown in Table 3, our framework outperforms all the state-of-the-art methods by a large margin of roughly 5% – 8% for each observation-prediction percentage pair on the Breakfast dataset. An interesting observation is that all the previous approaches achieve comparable results despite the fact that they are using different network architectures based on RNNs [2, 3], CNNs [3], or even temporal convolution [16]. On the contrary,

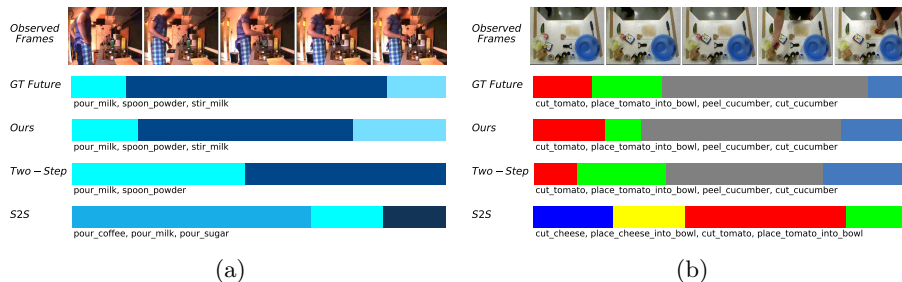


Fig. 4. Qualitative results for anticipating future activities. (a) An example from the Breakfast dataset for the case of observing 20% of the video and predicting the activities in the following 50%. (b) An example from the 50Salads dataset for the case of observing 20% of the video and predicting the activities in the following 20%.

our framework clearly outperforms these approaches with the advantage that it was trained in an end-to-end fashion.

For 50Salads, our model outperforms the state-of-the-art in 50% of the cases. This is mainly because 50Salads is a small dataset. Since our model is trained end-to-end, it requires more data to show the benefits over two-step approaches.

Since the state-of-the-art methods like [3] use an RNN-HMM model [33] for recognition, we also report the results of [3] with our TCN as a recognition model. The results are shown in Table 3 (RNN [3] + TCN and CNN [3] + TCN). Our model outperforms these methods even when they are combined with TCN. This highlights that the improvements in our model are not only due to the TCN, but mainly because of the joint optimization of all modules for the anticipation task in an end-to-end fashion and the introduced cycle consistency loss.

Qualitative results for our model on both datasets are illustrated in Fig. 4. As shown in the figure, our model can generate accurate predictions of the future activities and their duration. We also show the results of the sequence-to-sequence (S2S) and the two-step baselines. Our model anticipates the activities better.

5 Conclusion

In this paper, we introduced a model for anticipating future activities from a partially observed video. In contrast to the state-of-the-art methods which rely on the action labels of the observations, our model directly predicts the future activities from the observed frames. We train the proposed model in an end-to-end fashion and show a superior performance compared to the previous approaches. Additionally, we introduced a cycle consistency loss for the anticipation task which further boosts the performance. Our framework achieves state-of-the-art results on two publicly available datasets.

Acknowledgments: The work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GA 1927/4-1 (FOR 2535 Anticipating Human Behavior) and the ERC Starting Grant ARCA (677650).

References

1. Abu Farha, Y., Gall, J.: MS-TCN: Multi-stage temporal convolutional network for action segmentation. In: CVPR (2019)
2. Abu Farha, Y., Gall, J.: Uncertainty-aware anticipation of activities. In: ICCV Workshops (2019)
3. Abu Farha, Y., Richard, A., Gall, J.: When will you do what?-Anticipating temporal occurrences of activities. In: CVPR (2018)
4. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social LSTM: Human trajectory prediction in crowded spaces. In: CVPR (2016)
5. Bhattacharyya, A., Fritz, M., Schiele, B.: Bayesian prediction of future street scenes using synthetic likelihoods. In: ICLR (2019)
6. Carreira, J., Zisserman, A.: Quo vadis, action recognition? A new model and the kinetics dataset. In: CVPR (2017)
7. Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A.: Temporal cycle-consistency learning. In: CVPR (2019)
8. Furnari, A., Battiato, S., Farinella, G.M.: Leveraging uncertainty to rethink loss functions and evaluation measures for egocentric action anticipation. In: ECCV Workshops. Springer (2018)
9. Furnari, A., Farinella, G.M.: What would you expect? Anticipating egocentric actions with rolling-unrolling LSTMs and modality attention. In: ICCV (2019)
10. Gammulle, H., Denman, S., Sridharan, S., Fookes, C.: Forecasting future action sequences with neural memory networks. In: BMVC (2019)
11. Gammulle, H., Denman, S., Sridharan, S., Fookes, C.: Predicting the future: A jointly learnt model for action anticipation. In: ICCV (2019)
12. Gao, J., Yang, Z., Nevatia, R.: RED: Reinforced encoder-decoder networks for action anticipation. In: BMVC (2017)
13. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR (2017)
14. Hoai, M., De la Torre, F.: Max-margin early event detectors. IJCV 107(2), 191–202 (2014)
15. Jain, A., Zamir, A.R., Savarese, S., Saxena, A.: Structural-RNN: Deep learning on spatio-temporal graphs. In: CVPR (2016)
16. Ke, Q., Fritz, M., Schiele, B.: Time-conditioned action anticipation in one shot. In: CVPR (2019)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
18. Kitani, K.M., Ziebart, B.D., Bagnell, J.A., Hebert, M.: Activity forecasting. In: ECCV. Springer (2012)
19. Koppula, H.S., Saxena, A.: Anticipating human activities using object affordances for reactive robotic response. TPAMI 38(1), 14–29 (2016)
20. Kuehne, H., Arslan, A., Serre, T.: The language of actions: Recovering the syntax and semantics of goal-directed human activities. In: CVPR (2014)
21. Lan, T., Chen, T.C., Savarese, S.: A hierarchical representation for future action prediction. In: ECCV. Springer (2014)
22. Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: CVPR (2017)
23. Liang, J., Jiang, L., Niebles, J.C., Hauptmann, A.G., Fei-Fei, L.: Peeking into the future: Predicting future person activities and locations in videos. In: CVPR (2019)

24. Liang, X., Lee, L., Dai, W., Xing, E.P.: Dual motion GAN for future-flow embedded video prediction. In: ICCV (2017)
25. Luc, P., Neverova, N., Couprie, C., Verbeek, J., LeCun, Y.: Predicting deeper into the future of semantic segmentation. In: ICCV (2017)
26. Ma, S., Sigal, L., Sclaroff, S.: Learning activity progression in LSTMs for activity detection and early detection. In: CVPR (2016)
27. Mahmud, T., Billah, M., Hasan, M., Roy-Chowdhury, A.K.: Captioning near-future activity sequences. arXiv (2019)
28. Mahmud, T., Hasan, M., Roy-Chowdhury, A.K.: Joint prediction of activity labels and starting times in untrimmed videos. In: ICCV (2017)
29. Martinez, J., Black, M.J., Romero, J.: On human motion prediction using recurrent neural networks. In: CVPR (2017)
30. Mathieu, M., Couprie, C., LeCun, Y.: Deep multi-scale video prediction beyond mean square error. In: ICLR (2016)
31. Mehrotra, N., Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: A variational auto-encoder model for stochastic point processes. In: CVPR (2019)
32. Miech, A., Laptev, I., Sivic, J., Wang, H., Torresani, L., Tran, D.: Leveraging the present to anticipate the future in videos. In: CVPR Workshops (2019)
33. Richard, A., Kuehne, H., Gall, J.: Weakly supervised action learning with RNN based fine-to-coarse modeling. In: CVPR (2017)
34. Rodriguez, C., Fernando, B., Li, H.: Action anticipation by predicting future dynamic images. In: ECCV Workshops. Springer (2018)
35. Ruiz, A.H., Gall, J., Moreno-Noguer, F.: Human motion prediction via spatio-temporal inpainting. In: ICCV (2019)
36. Ryoo, M.S.: Human activity prediction: Early recognition of ongoing activities from streaming videos. In: ICCV (2011)
37. Sadegh Aliakbarian, M., Sadat Saleh, F., Salzmann, M., Fernando, B., Petersson, L., Andersson, L.: Encouraging LSTMs to anticipate actions very early. In: ICCV (2017)
38. Sener, F., Yao, A.: Zero-shot anticipation for instructional activities. In: ICCV (2019)
39. Shi, Y., Fernando, B., Hartley, R.: Action anticipation with RBF kernelized feature mapping RNN. In: ECCV (2018)
40. Srivastava, N., Mansimov, E., Salakhudinov, R.: Unsupervised learning of video representations using LSTMs. In: ICML. pp. 843–852 (2015)
41. Stein, S., McKenna, S.J.: Combining embedded accelerometers with computer vision for recognizing food preparation activities. In: UbiComp (2013)
42. Sun, C., Shrivastava, A., Vondrick, C., Sukthankar, R., Murphy, K., Schmid, C.: Relational action forecasting. In: CVPR (2019)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NIPS (2017)
44. Vondrick, C., Pirsivash, H., Torralba, A.: Anticipating visual representations from unlabeled video. In: CVPR (2016)
45. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: CVPR (2019)
46. Zeng, K.H., Shen, W.B., Huang, D.A., Sun, M., Carlos Niebles, J.: Visual forecasting by imitating dynamics in natural sequences. In: ICCV (2017)
47. Zhou, T., Krahenbuhl, P., Aubry, M., Huang, Q., Efros, A.A.: Learning dense correspondence via 3d-guided cycle consistency. In: CVPR (2016)
48. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)