

Predicting the Future: A Jointly Learnt Model for Action Anticipation

Harshala Gammulle Simon Denman Sridha Sridharan Clinton Fookes
 Image and Video Research Lab, SAIVT, Queensland University of Technology (QUT), Australia
 {pranali.gammulle, s.denman, s.sridharan, c.fookes}@qut.edu.au

Abstract

Inspired by human neurological structures for action anticipation, we present an action anticipation model that enables the prediction of plausible future actions by forecasting both the visual and temporal future. In contrast to current state-of-the-art methods which first learn a model to predict future video features and then perform action anticipation using these features, the proposed framework jointly learns to perform the two tasks, future visual and temporal representation synthesis, and early action anticipation. The joint learning framework ensures that the predicted future embeddings are informative to the action anticipation task. Furthermore, through extensive experimental evaluations we demonstrate the utility of using both visual and temporal semantics of the scene, and illustrate how this representation synthesis could be achieved through a recurrent Generative Adversarial Network (GAN) framework. Our model outperforms the current state-of-the-art methods on multiple datasets: UCF101, UCF101-24, UT-Interaction and TV Human Interaction.¹

1. Introduction

We propose an action anticipation model that uses visual and temporal data to predict future behaviour, while also predicting a frame-wise future representation to support the learning. Unlike action recognition where the recognition is carried out after the event, by observing the full video sequence (Fig. 1(a)), the aim of action anticipation (Fig. 1(b)) is to predict the future action as early as possible by observing only a portion of the action [3]. Therefore, for the prediction we only have partial information in the form of a small number of frames, so the available information is scarce. Fig. 1(c) shows the intuition behind our proposed model. The action anticipation task is accomplished by jointly learning to predict the future embeddings (both visual and temporal) along with the action anticipation task, where the anticipation task provides cues to help compen-



Figure 1. Action anticipation through future embedding prediction. Action recognition approaches (a) carry out the recognition task via fully observed video sequences while the typical action anticipation methods (b) are based on predicting the action from a small portion of the frames. In our proposed model (c) we jointly learn the future frame embeddings to support the anticipation task.

sate for the missing information from the unobserved frame features. We demonstrate that joint learning of the two tasks complements each other.

This approach is inspired by recent theories of how humans achieve the action predictive ability. Recent psychology literature has shown that humans build a mental image of the future, including future actions and interactions (such as interactions between objects) before initiating muscle movements or motor controls [10, 17, 31]. These representations capture both the visual and temporal information of the expected future. Mimicking this biological process,

¹This research was supported by an Australian Research Council (ARC) Linkage grant LP140100221

our action anticipation method jointly learns to anticipate future scene representations while predicting the future action, and outperforms current state-of-the-art methods.

In contrast to recent works [3, 45, 50] which rely solely on visual inputs, and inspired by [10, 17, 31], we propose a joint learning process which attends to salient components of both visual and temporal streams, and builds a highly informative context descriptor for future action anticipation. In [50] the authors demonstrate that the context semantics, which capture high level action related concepts including environmental details, objects, and historical actions and interactions, are more important when anticipating actions than the actual pixel values of future frames. Furthermore, the semantics captured through pre-trained deep learning models show robustness to background and illumination changes as they tend to capture the overall meaning of the input frame rather than simply using pixel values [8, 26]. Hence in the proposed architecture we extract deep visual and temporal representations from the inputs streams and predict the future representations of those streams.

Motivated by recent advances in Generative Adversarial Networks (GAN) [1, 16, 33] and their ability to automatically learn a task specific loss function, we employ a GAN learning framework in our approach as it provides the capability to predict a plausible future action sequence.

Although there exist individual GAN models for anticipation [32, 56], we take a step further in this work. The main contribution is the joint learning of a context descriptor for two tasks, action anticipation and representation prediction, through the joint training of two GANs.

Fig. 2 shows the architecture of our proposed Action Anticipation GAN (AA-GAN) model. The model receives the video frames and optical flow streams as the visual and temporal representations of the scene. We extract a semantic representation of the individual streams by passing them through a pre-trained feature extractor, and fuse them through an attention mechanism. This allows us to provide a varying level of focus to each stream and effectively embed the vital components for different action categories. Through this process low level feature representations are mapped to a high-level context descriptor which is then used by both the future representation synthesis and classification procedures. By coupling the GANs (visual and temporal synthesisers) through a common context descriptor, we optimally utilise all available information and learn a descriptor which better describes the given scene.

Our main contributions are as follow:

- We propose a joint learning framework for early action anticipation and synthesis of the future representations.
- We demonstrate how attention can efficiently determine the salient components from the multi-modal in-

formation, and generate a single context descriptor which is informative for both tasks.

- We introduce a novel regularisation method based on the exponential cosine distance, which effectively guides the generator networks in the prediction task.
- We perform evaluations on several challenging datasets, and through a thorough ablation study, demonstrate the relative importance of each component of the proposed model.

2. Previous Work

Human action recognition is an active research area that has great importance in multiple domains [7, 9, 23]. Since the inception of the field researchers have focused on improving the applicability of methods to tally with real world scenarios. The aim of early works was to develop discrete action recognition methods using image [6, 25] or video inputs [15, 22, 46], and these have been extended to detect actions in fine-grained videos [29, 37]. Although these methods have shown impressive performance, they are still limited for real-world applications as they rely on fully completed action sequences. This motivates the development of action anticipation methods, which can accurately predict future actions utilising a limited number of early frames, and thereby providing the ability to predict actions that are in progress.

In [50], a deep network is proposed to predict a representation of the future. The predicted representation is used to classify future actions. However [50] requires the progress level of the ongoing action to be provided during testing, limiting applicability [19]. Hu et al. [19] introduced a soft regression framework to predict ongoing actions. This method [19] learns soft labels for regression on the subsequences containing partial action executions. Lee et al. [30] proposed a human activity representation method, termed sub-volume co-occurrence matrix, and developed a method to predict partially observed actions with the aid of a pre-trained CNN. The deep network approach of Aliakbarian et al. [3] used a multi-stage LSTM architecture that incorporates context-aware and action-aware features to predict classes as early as possible. The CNN based action anticipation model of [40] predicts the most plausible future motion, and was improved via an effective loss function based on dynamic and classification losses. The dynamic loss is obtained through a dynamic image generator trained to generate class specific dynamic images. However, performance is limited due to the hand-crafted loss function. A GAN based model can overcome this limitation as it can automatically learn a loss function and has shown promising performance in recent research [35, 38, 54].

In our work we utilise a conditional GAN [12, 13, 36] for deep future representation generation. A limited number

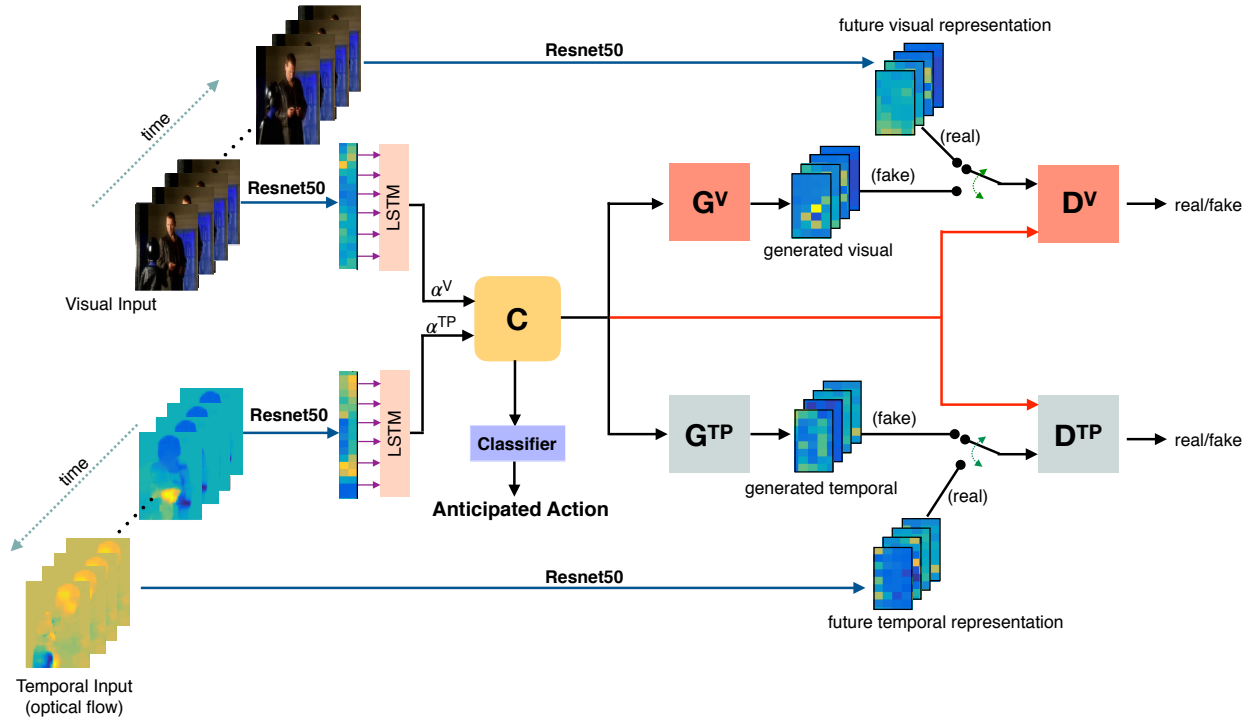


Figure 2. Action Anticipation GAN (AA-GAN): The model receives RGB and optical flow streams as the visual and temporal representations of the given scene. Rather than utilising the raw streams we extract the semantic representation of the individual streams by passing them through a pre-trained feature extractor. These streams are merged via an attention mechanism which embeds these low-level feature representations in a high-level context descriptor. This context representation is utilised by two GANs: one for future visual representation synthesis and one for future temporal representation synthesis; and the anticipated future action is obtained by utilising the context descriptor. Hence context descriptor learning is influenced by both the future representation prediction, and the action anticipation task.

of GAN approaches can be found for human action recognition [1, 33]. In [33], a GAN is used to generate masks to detect the actors in input frame and action classification is done via a CNN. This method is prone to difficulties with the loss function as noted previously. Considering other GAN methods, [32, 56] require human skeletal data which is not readily available; [56] only synthesises the future skeletal representation; and [55] considers the task of synthesising future gaze points using a single generator and discriminator pair and directly extracting spatio-temporal features from a 3D CNN. In contrast to these, we analyse two modes and utilise an attention mechanism to embed the salient components of each mode into a context descriptor which can be used for multiple tasks; and we learn this descriptor through joint training of two GANs and a classifier.

The authors of [45] have adapted the model of [50] to a GAN setting; using GANs to predict the future visual feature representation. Upon training this representation, they train a classifier on the predicted features to anticipate the future action class. We argue that the approach of [45] is suboptimal, as there is no guarantee that the future action representation is well suited to predicting the action due to the two stage learning. Our approach, which learns the tasks jointly, ensures that a rich multi-modal embedding is learnt

that captures the salient information needed for both tasks. Furthermore, by extending this to a multi-modal setting, we demonstrate the importance of attending to both visual and temporal features for the action anticipation task.

3. Action Anticipation Model

Our action anticipation model is designed to predict the future while classifying future actions. The model aims to generate embeddings for future frames, to obtain a complete notion of the ongoing action and to understand how best to classify the action. In Sec. 3.1, we discuss how the context descriptor is generated using the visual and temporal input streams while Sec. 3.2 describes the use of the GAN in the descriptor generation process. The future action classification procedure is described in Sec. 3.3 and we further improve this process with the addition of the cosine distance based regularisation method presented in Sec. 3.4.

3.1. Context Descriptor Formulation

Inputs to our model are two fold: visual and temporal. The visual inputs are the RGB frames and the temporal inputs are the corresponding optical flow images (computed using [4]). If the number of input video frames is T , then

both the visual input (I^V) and the temporal input (I^{TP}) can be represented as follows,

$$\begin{aligned} I^V &= \{I_1^V, I_2^V, \dots, I_T^V\}, \\ I^{TP} &= \{I_1^{TP}, I_2^{TP}, \dots, I_T^{TP}\}. \end{aligned} \quad (1)$$

These inputs are passed through a pre-trained feature extractor which extracts features θ^V and θ^{TP} frame wise,

$$\begin{aligned} \theta^V &= \{\theta_1^V, \theta_2^V, \dots, \theta_T^V\}, \\ \theta^{TP} &= \{\theta_1^{TP}, \theta_2^{TP}, \dots, \theta_T^{TP}\}. \end{aligned} \quad (2)$$

Then θ^V and θ^{TP} are sent through separate LSTM networks to capture the temporal structure of the input features. The LSTM outputs are defined as,

$$h_t^V = LSTM(\theta_t^V), h_t^{TP} = LSTM(\theta_t^{TP}). \quad (3)$$

Attention values are generated for each frame such that,

$$e_t^V = \tanh(a^V[h_t^V]^\top), e_t^{TP} = \tanh(a^{TP}[h_t^{TP}]^\top), \quad (4)$$

where a^V and a^{TP} are multilayer perceptrons trained together with the rest of the network, and are passed through a sigmoid function to get the score values,

$$\alpha_t^V = \sigma([e_t^V, e_t^{TP}]), \alpha_t^{TP} = 1 - \alpha_t^V. \quad (5)$$

Then, an attention weighted output vector is generated,

$$\tilde{\mu}_t^V = \alpha_t^V h_t^V, \tilde{\mu}_t^{TP} = \alpha_t^{TP} h_t^{TP} \quad (6)$$

Finally these output vectors are concatenated (denoted by $[\cdot]$) to generate the context descriptor (C_t),

$$C_t = [\tilde{\mu}_t^V, \tilde{\mu}_t^{TP}]. \quad (7)$$

C_t encodes the recent history of both inputs, and thus is used to predict future behaviour.

3.2. Visual and Temporal GANs

GAN based models are capable of learning an output that is difficult to discriminate from real examples. They learn a mapping from the input to this realistic output while learning a loss function to train the mapping. The context descriptor, C_t , is the input for both GANs (visual and temporal synthesisers, see Fig. 2). The ground truth future visual and temporal frames are denoted F^V and F^{TP} , and are given by,

$$\begin{aligned} F^V &= \{F_1^V, F_2^V, \dots, F_T^V\}, \\ F^{TP} &= \{F_1^{TP}, F_2^{TP}, \dots, F_T^{TP}\}. \end{aligned} \quad (8)$$

We extract features for F^V and F^{TP} similar to Eq. 2,

$$\begin{aligned} \beta^V &= \{\beta_1^V, \beta_2^V, \dots, \beta_T^V\}, \\ \beta^{TP} &= \{\beta_1^{TP}, \beta_2^{TP}, \dots, \beta_T^{TP}\}. \end{aligned} \quad (9)$$

These features, β^V and β^{TP} , are utilised during GAN training. The aim of the generator (G^V or G^{TP}) of each GAN is to synthesise the future deep feature sequence that is sufficiently realistic to fool the discriminator (D^V or D^{TP}). It should be noted that the GAN models do not learn to predict the future frames, but the deep features of the frames (visual or temporal). As observed in [50] this allows the model to recognise higher-level concepts in the present and anticipate their relationships with future actions. This is learnt through the following loss functions,

$$\begin{aligned} \mathcal{L}^V(G^V, D^V) &= \sum_{t=1}^T \log D^V(C_t, \beta^V) + \\ &\quad \sum_{t=1}^T \log(1 - D^V(C_t, G^V(C_t))), \end{aligned} \quad (10)$$

$$\begin{aligned} \mathcal{L}^{TP}(G^{TP}, D^{TP}) &= \sum_{t=1}^T \log D^{TP}(C_t, \beta^{TP}) + \\ &\quad \sum_{t=1}^T \log(1 - D^{TP}(C_t, G^{TP}(C_t))). \end{aligned} \quad (11)$$

3.3. Classification

The deep future sequences are learnt through the two GAN models as described in Sec. 3.2. A naive way to perform the future action classification is using the trained future feature predictor and passing the synthesised future features to the classifier. However, this is sub-optimal as G^V and G^{TP} have no knowledge of this task, and thus features are likely sub-optimal for it. As such, in this work we investigate joint learning of the embedding prediction and future action anticipation, allowing the model to learn the salient features that are required for action anticipation. Hence, the GANs are able to support learning salient features for both processes. We perform future action classification for the action anticipation task through a classifier, the input for which is C_t . Then the classification loss can be defined as,

$$\mathcal{L}^C = - \sum_{t=1}^T y_t \log f^C(C_t). \quad (12)$$

It is important to note that the context descriptor C_t is influenced by both the classification loss, \mathcal{L}^C , and the GAN losses, \mathcal{L}^V and \mathcal{L}^{TP} , as G^V and G^{TP} utilise the context descriptor to synthesise the future representations.

3.4. Regularisation

To stabilise GAN learning a regularisation method such as the L_2 loss is often used [20]. However the cosine distance has been shown to be more effective when comparing deep embeddings [51, 52]. Furthermore when generating future sequence forecasts it is more challenging to forecast representations in the distant future than the near future. However, the semantics from the distant future are more informative for the action class anticipation problem, as they carry more information about what the agents are likely to do. Hence we propose a temporal regularisation mechanism which compares the predicted embeddings with the ground truth future embeddings using the cosine distance, and encourages the model to focus more on generating accurate embeddings for the distant future,

$$\mathcal{L}^R = \sum_{t=1}^T -e^t d(G^V(C_t), \beta_t^V) + \sum_{t=1}^T -e^t d(G^{TP}(C_t), \beta_t^{TP}), \quad (13)$$

where d represents the cosine distance function. Motivated by [3] we introduce the exponential term, e^t , encouraging more accurate prediction of distant future embeddings.

Then, the loss for the final model that learns the context descriptor C_t and is reinforced by both deep future sequence synthesisers (GAN models), and the future action classification can be written as,

$$\mathcal{L} = w^V \mathcal{L}^V + w^{TP} \mathcal{L}^{TP} + w^c \mathcal{L}^C + w^R \mathcal{L}^R, \quad (14)$$

where w^V, w^{TP}, w^c and w^R are hyper-parameters which control the contribution of the respective losses.

4. Evaluations

4.1. Datasets

Related works on action anticipation or early action prediction typically use discrete action datasets. The four datasets we use to evaluate our work are outlined below.

UCF101 [49] has been widely used for discrete action recognition and recent works for action anticipation due to its size and variety. The dataset includes 101 action classes from 13,320 videos with an average length of 7.2 seconds. In order to perform comparison to the state-of-the-art methods, we utilise the provided three training/testing splits and report the average accuracy over three splits.

UCF101-24 [47] is a subset of the UCF101 dataset. It is composed of 24 action classes in 3207 videos. In order to compare action anticipation results to the state-of-the-art we utilise only the data provided in set1.

UT-Interaction (UTI) [43] is a human interaction dataset, which contains videos of two or more people performing interactions such as handshake, punch etc. in a sequential and/or concurrent manner. The dataset has total of 120 videos. For the state-of-the-art comparison we utilise a 10-fold leave-one-out cross validation on each set and the mean performance over all sets is obtained, as per [3].

TV Human Interaction (TV-HI) [39] dataset is a collection of 300 video clips collected from 20 different TV shows. It is composed of four action classes of people performing interactions such as handshake, highfive, hug and kiss, and a fifth action class called ‘none’ which does not contain any of the four actions. The provided train/ test splits are utilised with a 25-fold cross validation, as per [50].

4.2. Network Architecture and Training

Considering related literature for different datasets, different numbers of observed frames [3, 45] are used. Let T be the number of observed frames, then we extract frames $T + 1$ to $T + \hat{T}$ as future frames, where \hat{T} is the number of future frames for embedding prediction. As the temporal input, similar to [46] we use dense optical flow displacements computed using [4]. In addition to horizontal and vertical components we also use the mean displacement of the horizontal and vertical flow. Both visual and temporal inputs are individually passed through a pre-trained ResNet50 [18] trained on ImageNet [41], and activations from the ‘activation_23’ layer are used as the input feature representation.

The network of the generator is composed of two LSTM layers followed by a fully connected layer. The generator is fed only with the context input while the discriminator is fed with both the context and the real/fake feature representations. The two inputs of the discriminator are passed through separate LSTM layers and then the merged output is passed through two fully connected layers. The classifier is composed of a single LSTM layer followed by a single fully connected layer. For clarity, we provide model diagrams in the supplementary materials. For all LSTMs, 300 hidden units are used. For the model training procedure we follow the approach of [20], alternating between one gradient decent pass for the discriminators, and the generators and the classifier using 32 samples per mini batch. The Adam optimiser [24] is used with a learning rate of 0.0002 and a decay of 8×10^{-9} , and is trained for 40 epochs. Hyper-parameters w^V, w^{TP}, w^c, w^R are evaluated experimentally and set to 25, 20, 43 and 15, respectively. Please refer to supplementary material for these evaluations. When training the proposed model for the UTI and TV-HI datasets, due to the limited availability of training examples we first train the model on UCF101 training data and fine-tuned it on the training data from the specific datasets.

For the implementation of our proposed method we utilised Keras [5] with Theano [2] as the backend.

4.3. Performance Evaluation

4.3.1 Evaluation Protocol

To evaluate our model on each dataset, where possible we consider two settings for the number of input frames, namely the ‘Earliest’ and ‘Latest’ settings. For UCF101 and UTI, similar to [3] we consider 20% and 50% of the frames for the ‘Earliest’ and ‘Latest’ settings, respectively; following [3] we do not use more than 50 frames for the ‘Latest’ setting. For each dataset and setting, we resample the input videos such that all sequences have a constant number of frames. Due to unavailability of baseline results and following [45], for UCF101-24 we report evaluate using 50% of the frames from each video and for the TV-HI dataset, as in [14, 27], we consider only 1 seconds worth frames.

4.3.2 Comparison to the state-of-the-art methods

Evaluations for UCF101, UCF101-24, UTI and TV-HI datasets are presented in Tables 1, to 4 respectively. Considering the results, the authors of Multi_stage LSTM [3] and RED [14] have introduced a new hand engineered loss that encourages the early prediction of the action class. The authors of RBF-RNN [45] use a GAN learning process where the loss function is also automatically learnt. Similar to the proposed architecture, the RBF-RNN [45] model also utilises the spatial representation of the scene through a Deep CNN model and tries to predict the future scene representations. However in contrast to the proposed architecture this method does not utilise temporal features, or joint learning. We learn a context descriptor which effectively combines both spatial and temporal representations which not only aids the action classification but also anticipates the future representations more accurately. This led us to obtain superior results. In Tab. 2, the results for UCF101-24 shows that our model is able to outperform RBF-RNN [45] by 0.9% while in Tab. 3 we outperform [45] on the UTI dataset by 1.3% at the earliest setting.

When comparing the performance gap between the earliest and latest settings, our model has a smaller performance drop compared to the baseline models. The gap for UCF101 on our model is 1.4% while the gap for the Multi_stage LSTM model [3] is 2.9%. G^V and G^{TP} synthesise the future representation of both visual and temporal streams while considering the current context. As such, the proposed model is able to better anticipate future actions, even with fewer frames. Our evaluations on multiple benchmarks further illustrate the generalisability of the proposed architecture, with varying video lengths and dataset sizes.

4.4. Ablation Experiments

To further demonstrate the proposed AA-GAN method, we conducted an ablation study by strategically removing

Method	Earliest	Latest
Context-aware + loss in [21]	30.6	71.1
Context-aware + loss in [34]	22.6	73.1
Multi_stage LSTM [3]	80.5	83.4
Proposed	84.2	85.6

Table 1. Action anticipation results for UCF101 considering the ‘Earliest’ 20% of frames and ‘Latest’ 50% of frames.

Method	Accuracy
Temporal Fusion [11]	86.0
ROAD [47]	92.0
ROAD + BroxFlow [47]	90.0
RBF-RNN [45]	98.0
Proposed	98.9

Table 2. Action anticipation results for UCF101-24 considering 50% of frames from each video.

Method	Earliest	Latest
S-SVM [48]	11.0	13.4
DP-SVM [48]	13.0	14.6
CuboidBayes [42]	25.0	71.7
CuboidSVM [44]	31.7	85.0
Context-aware+ loss in [21]	45.0	65.0
Context-aware + loss in [34]	48.0	60.0
I-BoW [42]	65.0	81.7
BP-SVM [28]	65.0	83.3
D-BoW [42]	70.0	85.0
multi-stageLSTM [3]	84.0	90.0
Future-dynamic [40]	89.2	91.9
RBF-RNN [45]	97.0	NA
Proposed	98.3	99.2

Table 3. Action anticipation results for UTI ‘Earliest’ 20% of frames and ‘Latest’ 50% of frames.

Method	Accuracy
Vondrick et. al [50]	43.6
RED [14]	50.2
Proposed	55.7

Table 4. Action anticipation results for TV Human Interaction dataset considering 1 second worth of frames from each video.

components of the proposed system. We evaluated seven non-GAN based model variants and ten GAN-based variants of the proposed AA-GAN model. Non-GAN based models are further broken into two categories: models with and without future representation generators. Similarly, the GAN based models fall into two categories: those that do and do not learn tasks jointly. Diagrams of these ablation models are available in the supplementary materials.

Non-GAN based models: These models do not utilise any future representation generators, and are only trained through classification loss.

- (a) $\eta^{C,V}$: A model trained to classify using the context

feature extracted only from the visual input stream (V).

- (b) $\eta^{C,TP}$: As per model (a), but using the temporal input stream (TP).
- (c) $\eta^{C,(V+TP)}$: As per (a), but using both data streams to create the context embedding.

Non-GAN based models with future representation generators: Here, we add future embedding generators to the previous set of models. The generators are trained through mean squared error (i.e. no discriminator and no adversarial loss) while the classification is learnt through categorical cross entropy loss. The purpose of these models is to show how the joint learning can improve performance, and how a common embedding can serve both tasks.

- (d) $\eta^{C,V} + G^V$: Model with the future visual representation generator (G^V) and fed only with the visual input stream to train the classifier
- (e) $\eta^{C,TP} + G^{TP}$: As per (d), but receiving and predicting the temporal input stream.
- (f) $\eta^{C,(V+TP)} + G^V + G^{TP}$: The model is composed of both generators, G^V and G^{TP} , and fed with both visual and temporal input streams.
- (g) $\eta^{C,(V+TP)} + G^V + G^{TP} + \text{Att}$: As per (f) but with the use of attention to combine the streams.

GAN based models without joint training: These methods are based on the GAN framework that generates future representations and a classifier that anticipates the action where these two tasks are learnt separately. We first train the GAN model using the adversarial loss and once this model is trained, using the generated future embeddings the classifier anticipates the action.

- (h) $\eta^{C,V} + GAN^V \setminus \text{Joint}$: Use the GAN learning framework with only the visual input stream and cosine distance based regularisation is used.
- (i) $\eta^{C,TP} + GAN^{TP} \setminus \text{Joint}$: As per (h), but with the temporal input stream
- (j) AA-GAN $\setminus \text{Joint}$: Use the GAN learning framework with both the visual and temporal input streams.

GAN based models with joint training: These models train the deep future representation generators adversarially. The stated model variants are introduced by removing the different components from the proposed model.

- (k) $\eta^{C,V} + GAN^V \setminus (\mathcal{L}^R)$: The proposed approach with only the visual input stream and without cosine distance based regularisation.

- (l) $\eta^{C,TP} + GAN^{TP} \setminus (\mathcal{L}^R)$: The proposed approach with only the temporal input stream and without cosine distance based regularisation.
- (m) $\eta^{C,V} + GAN^V$: The proposed approach with only the visual input stream. Cosine distance based regularisation is used.
- (n) $\eta^{C,TP} + GAN^{TP}$: The proposed approach with only the temporal input stream. Cosine distance based regularisation is used.
- (o) AA-GAN $\setminus (\mathcal{L}^R)$: Proposed model without cosine distance based regularisation.
- (p) AA-GAN $\setminus (DR)$: Similar to the proposed model, however G^V and G^{TP} predict pixel values for future visual and temporal frames instead of representations extracted from the pre-trained feature extractor.

Method	Accuracy
(a) $\eta^{C,V}$	45.1
(b) $\eta^{C,TP}$	39.8
(c) $\eta^{C,(V+TP)}$	52.0
(d) $\eta^{C,V} + G^V$	54.7
(e) $\eta^{C,TP} + G^{TP}$	52.4
(f) $\eta^{C,(V+TP)} + G^V + G^{TP}$	68.1
(g) $\eta^{C,(V+TP)} + G^V + G^{TP} + \text{Att}$	68.8
(h) $\eta^{C,V} + GAN^V \setminus \text{Joint}$	98.1
(i) $\eta^{C,TP} + GAN^{TP} \setminus \text{Joint}$	97.9
(j) AA-GAN $\setminus \text{Joint}$	98.3
(k) $\eta^{C,V} + GAN^V \setminus (\mathcal{L}^R)$	96.0
(l) $\eta^{C,TP} + GAN^{TP} \setminus (\mathcal{L}^R)$	95.4
(m) $\eta^{C,V} + GAN^V$	98.4
(n) $\eta^{C,TP} + GAN^{TP}$	98.1
(o) AA-GAN $\setminus (\mathcal{L}^R)$	98.7
(p) AA-GAN $\setminus (DR)$	95.9
AA-GAN (proposed)	98.9

Table 5. Ablation results for UCF101-24 dataset for the ‘Latest’ setting, which uses 50% of the frames from each video.

The evaluation results of the ablation models on the UCF101-24 test set are presented in Tab. 5.

Non-GAN based models (a to g): Model performance clearly improves when using both data streams together over either one individually (see (c) vs (a) and (b); and (f) vs (d) and (e)). Hence, it is clear that both streams provide different information cues to facilitate the prediction. Comparing the results of models that do not utilise the future representation generators to (d), we see that overseeing future representation does improve the results.

GAN based models without joint training (h to j): Comparing the non-GAN based methods with ablation

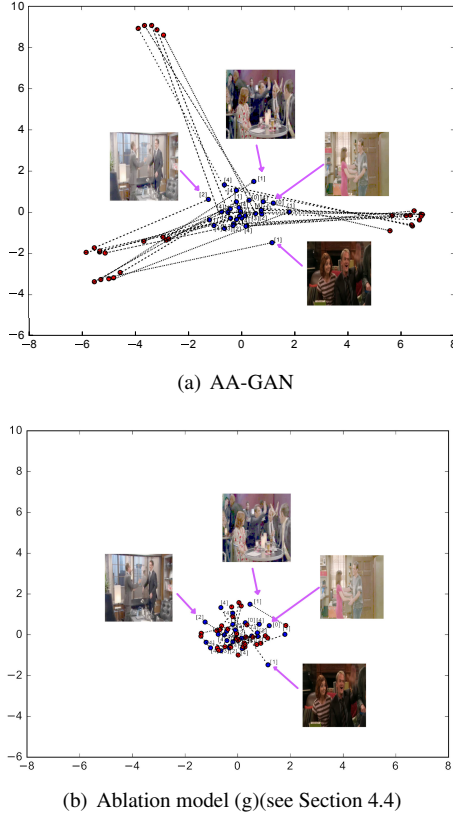


Figure 3. Projections of the discriminator hidden states for the AA-GAN (a) and ablation model (g) in (b) before (in blue) and after (in red) training. Ground truth action classes are in brackets. Insert indicates sample frames from the respective videos.

model (h), we see that a major performance boost is achieved through the GAN learning process, denoting the importance of the automated loss function learning. Comparing the performance of visual and temporal streams, we observe that the visual stream is dominant, however combining both streams through the proposed attention mechanism captures complimentary information.

GAN based models with joint training (k to p): Comparing models (h) and (i), which are single modal models that do not use joint training, with models (m) and (n) which do, we can see the clear benefit offered by learning the two complementary tasks together. This contradicts the observation reported in [45], who use a classifier which was connected to the predicted future embeddings. We speculate that by learning a compressed context representation for both tasks we effectively propagate the effect of the action anticipation error through the encoding mechanisms, allowing this representation to be informative for both tasks. Finally, by coupling the GAN loss together with \mathcal{L}^R , where the cosine distance based regularisation is combined with the exponential loss to encourage accurate long-term predictions, we achieve state-of-the-art results. Furthermore

we compare the proposed AA-GAN model, where G^V and G^{TP} synthesise future visual and temporal representations, against ablation model (p) where G^V and G^{TP} synthesise pixel values for future frames. It is evident that the latter model fails to capture the semantic relationships between the low-level pixel features and the action class, leading to the derived context descriptor being less informative for action classification, reducing performance.

To demonstrate the discriminative nature of the learnt context embeddings, Fig. 3 (a) visualises the embedding space before (in blue) and after (in red) training of the proposed context descriptor for 30 randomly selected examples of the TV-HI test set. We extracted the learned context descriptor, C_t , and applied PCA [53] to generate 2D vectors. Ground truth action classes are indicated in brackets.

This clearly shows that the proposed context descriptor learns embeddings which are informative for both future representation generation and the segregation of action classes. From the inserts which show sample frames from the videos, visual similarities exist between the classes, hence the overlap in the embedding space before training. However after learning, the context descriptor has been able to maximise the interclass distance while minimising the distance within the class. Fig. 3 (b) shows an equivalent plot for the ablation model (g). Given the cluttered nature of the embeddings before and after learning, it is clear that the proposed GAN learning process makes a significant contribution to learning discriminative embeddings ²

4.5. Time Complexity

We evaluate the computational demands of the proposed AA-GAN model for the UTI dataset’s ‘Earliest’ setting. The model contains 43M trainable parameters, and generates 500 predictions (including future visual and temporal predictions and the action anticipation) in 1.64 seconds using a single core of an Intel E5-2680 2.50 GHz CPU.

5. Conclusion

In this paper we propose a framework which jointly learns to anticipate an action while also synthesising future scene embeddings. We learn a context descriptor which facilitates both of these tasks by systematically attending to individual input streams and effectively extracts the salient features. This method exhibits traits analogous to human neurological behaviour in synthesising the future, and renders an end to end learning platform. Additionally, we introduced a cosine distance based regularisation method to guide the generators in the synthesis task. Our evaluations demonstrate the superior performance of the proposed method on multiple public benchmarks.

²Additional qualitative evaluations showing generated future visual and temporal representations are in the supplementary material.

References

- [1] Unaiza Ahsan, Chen Sun, and Irfan Essa. Discrimnet: Semi-supervised action recognition from videos using generative adversarial networks. *arXiv preprint arXiv:1801.07230*, 2018.
- [2] Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, Christof Angermueller, Dzmitry Bahdanau, Nicolas Ballas, Frédéric Bastien, Justin Bayer, Anatoly Belikov, Alexander Belopolsky, et al. Theano: A python framework for fast computation of mathematical expressions. *arXiv preprint arXiv:1605.02688*, 472:473, 2016.
- [3] Mohammad Sadegh Aliakbarian, F Sadat Saleh, Mathieu Salzmann, Basura Fernando, Lars Petersson, and Lars Andersson. Encouraging lstms to anticipate actions very early. In *IEEE International Conference on Computer Vision (ICCV)*, volume 1, 2017.
- [4] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision*, pages 25–36. Springer, 2004.
- [5] François Chollet et al. Keras. <https://keras.io>, 2015.
- [6] Vincent Delaitre, Ivan Laptev, and Josef Sivic. Recognizing human actions in still images: a study of bag-of-features and part-based representations. In *British Machine Vision (BMVC) Conference*, 2010. updated version, available at <http://www.di.ens.fr/willow/research/stillactions/>.
- [7] Alan Dix. Human-computer interaction. In *Encyclopedia of database systems*, pages 1327–1331. Springer, 2009.
- [8] Jeff Donahue, Y Jia, O Vinyals, J Hoffman, N Zhang, E Tzeng, and T Darrell. A deep convolutional activation feature for generic visual recognition. arxiv preprint. *arXiv preprint arXiv:1310.1531*, 2013.
- [9] Ahmet Ekin, A Murat Tekalp, and Rajiv Mehrotra. Automatic soccer video analysis and summarization. *IEEE Transactions on Image processing*, 12(7):796–807, 2003.
- [10] Birgit Elsner and Bernhard Hommel. Effect anticipation and action control. *Journal of experimental psychology: human perception and performance*, 27(1):229, 2001.
- [11] Zhaoxuan Fan, Tianwei Lin, Xu Zhao, Wanli Jiang, Tao Xu, and Ming Yang. An online approach for gesture recognition toward real-world applications. In *International Conference on Image and Graphics*, pages 262–272. Springer, 2017.
- [12] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Multi-level sequence gan for group activity recognition. In *Asian Conference on Computer Vision*, pages 331–346. Springer, 2018.
- [13] Harshala Gammulle, Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Coupled generative adversarial network for continuous fine-grained action segmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 200–209. IEEE, 2019.
- [14] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. Red: Reinforced encoder-decoder networks for action anticipation. *arXiv preprint arXiv:1707.04818*, 2017.
- [15] Georgia Gkioxari and Jitendra Malik. Finding action tubes. In *The IEEE Conference on Computer Vision and Pattern recognition (CVPR)*, 2015.
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [17] Peter F Greve. The role of prediction in mental processing: A process approach. *New Ideas in Psychology*, 39:45–52, 2015.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [19] Jian-Fang Hu, Wei-Shi Zheng, Lianyang Ma, Gang Wang, Jian-Huang Lai, and Jianguo Zhang. Early action prediction by soft regression. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [21] Ashesh Jain, Avi Singh, Hema S Koppula, Shane Soh, and Ashutosh Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pages 3118–3125. IEEE, 2016.
- [22] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):221–231, 2013.
- [23] Christoph G Keller and Dariu M Gavrilă. Will the pedestrian cross? a study on pedestrian path prediction. *IEEE Transactions on Intelligent Transportation Systems*, 15(2):494–506, 2014.
- [24] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015.
- [25] ByoungChul Ko, JuneHyeok Hong, and Jae-Yeal Nam. Human action recognition in still images using action poselets and a two-layer classification model. *J. Vis. Lang. Comput.*, 28(C):163–175, June 2015.
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [27] Tian Lan, Tsung-Chuan Chen, and Silvio Savarese. A hierarchical representation for future action prediction. In *European Conference on Computer Vision*, pages 689–704. Springer, 2014.
- [28] Kennard Laviers, Gita Sukthankar, David W Aha, Matthew Molineaux, C Darken, et al. Improving offensive performance through opponent modeling. In *AIIDE*, 2009.
- [29] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] Dong-Gyu Lee and Seong-Wan Lee. Prediction of partially observed human activity based on pre-trained deep representation. *Pattern Recognition*, 85:198–206, 2019.

- [31] Moritz Lehne and Stefan Koelsch. Toward a general psychological model of tension and suspense. *Frontiers in Psychology*, 6:79, 2015.
- [32] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5226–5234, 2018.
- [33] Xinyu Li, Yanyi Zhang, Jianyu Zhang, Yueyang Chen, Huangcan Li, Ivan Marsic, and Randall S Burd. Region-based activity recognition using conditional gan. In *Proceedings of the 2017 ACM on Multimedia Conference*, pages 1059–1067. ACM, 2017.
- [34] Shugao Ma, Leonid Sigal, and Stan 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.
- [35] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *International Conference on Learning Representations (ICLR)*, 2016.
- [36] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [37] Bingbing Ni, Xiaokang Yang, and Shenghua Gao. Progressively parsing interactional objects for fine grained action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1020–1028, 2016.
- [38] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.
- [39] Alonso Patron-Perez, Marcin Marszalek, Ian Reid, and Andrew Zisserman. Structured learning of human interactions in tv shows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(12):2441–2453, 2012.
- [40] Cristian Rodriguez, Basura Fernando, and Hongdong Li. Action anticipation by predicting future dynamic images. In *ECCV’18 workshop on Anticipating Human Behavior*, 2018.
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [42] Michael 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.
- [43] Michael S. Ryoo and J. K. Aggarwal. UT-Interaction Dataset, ICPR contest on Semantic Description of Human Activities (SDHA), 2010.
- [44] Michael S. Ryoo, Chia-Chih Chen, JK Aggarwal, and Amit 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.
- [45] Yuge Shi, Basura Fernando, and Richard Hartley. Action anticipation with rbf kernelized feature mapping rnn. In *European Conference on Computer Vision*, pages 305–322. Springer, 2018.
- [46] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *CoRR*, abs/1406.2199, 2014.
- [47] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip HS Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *ICCV*, pages 3657–3666, 2017.
- [48] Khurram Soomro, Haroon Idrees, and Mubarak Shah. Online localization and prediction of actions and interactions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [49] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [50] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 98–106, 2016.
- [51] Jiabao Wang, Yang Li, Zhuang Miao, Yulong Xu, and Gang Tao. Learning deep discriminative features based on cosine loss function. *Electronics Letters*, 53(14):918–920, 2017.
- [52] Nicolai Wojke and Alex Bewley. Deep cosine metric learning for person re-identification. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 748–756. IEEE, 2018.
- [53] Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.
- [54] Donggeun Yoo, Namil Kim, Sunggyun Park, Anthony S Paek, and In So Kweon. Pixel-level domain transfer. In *European Conference on Computer Vision*, pages 517–532. Springer, 2016.
- [55] Kuo-Hao Zeng, William B Shen, De-An Huang, Min Sun, and Juan Carlos Nieves. Visual forecasting by imitating dynamics in natural sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2999–3008, 2017.
- [56] Mengmi Zhang, Keng Teck Ma, Joo Hwee Lim, Qi Zhao, and Jiashi Feng. Deep future gaze: Gaze anticipation on egocentric videos using adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4372–4381, 2017.