Learning to Anticipate Egocentric Actions by Imagination

Yu Wu, Linchao Zhu, Xiaohan Wang, Yi Yang, Fei Wu

Abstract-Anticipating actions before they are executed is crucial for a wide range of practical applications, including autonomous driving and robotics. In this paper, we study the egocentric action anticipation task, which predicts future action seconds before it is performed for egocentric videos. Previous approaches focus on summarizing the observed content and directly predicting future action based on past observations. We believe it would benefit the action anticipation if we could mine some cues to compensate for the missing information of the unobserved frames. We then propose to decompose the action anticipation into a series of future feature predictions. We imagine how the visual feature changes in the near future and then predicts future action labels based on these imagined representations. Differently, our ImagineRNN is optimized in a contrastive learning way instead of feature regression. We utilize a proxy task to train the ImagineRNN, i.e., selecting the correct future states from distractors. We further improve ImagineRNN by residual anticipation, i.e., changing its target to predicting the feature difference of adjacent frames instead of the frame content. This promotes the network to focus on our target, *i.e.*, the future action, as the difference between adjacent frame features is more important for forecasting the future. Extensive experiments on two large-scale egocentric action datasets validate the effectiveness of our method. Our method significantly outperforms previous methods on both the seen test set and the unseen test set of the EPIC Kitchens Action Anticipation Challenge.

Index Terms—Action Anticipation, Action Prediction, Egocentric videos

I. INTRODUCTION

NTICIPATING the near future is a natural task that has drawn increasing research attention [1], [2]. It has a wide range of applications in the intelligent systems when it needs to react before an action gets executed. For instance, it is critical to anticipate if a car would stop or a pedestrian would cross the road in the autonomous driving task. The prediction is supposed to be seconds before the action is actually taken place, so that the autonomous vehicle could have time to react to avoid an accident. Under these circumstances, recent works are proposed to predict activities a few seconds in the future [3], [4], which is practical for real-world applications.

In this paper, we focus on the problem of egocentric action anticipation defined in [4]. Egocentric (First Person Vision) videos [5], [6] offers an interesting scenario to study the action



Fig. 1. In the action anticipation task, the model needs to predict the future action that happens T seconds later. Predicting the intermediate future features [7], [8], [9], [10], [11] would benefit the action anticipation task. Our study focuses on how to learn a better imagined intermediate feature.

anticipation problem. Given an egocentric video sequence denoted as observed video, we aim to predict the future action that happens after a time period of T seconds, whereas the time T is known as the anticipation time.

Most previous approaches [3], [4] focus on summarizing the past observed frames, and then directly predict the future action that takes place T seconds later. These methods overlook the temporal gap between the past observations and the future action that is supposed to be predicted. However, frames in this temporal period are closer to the future, thus containing more useful evidence for the next action. If we could mine cues to compensate for missing information of unobserved frames, it would be easier for anticipation models to predict the future.

In this paper, we propose to tackle this issue by imagining the near future. First, we decompose the long-time action anticipation into a series of future feature predictions. We imagine how the visual feature changes in the very near future and then predict the future action labels based on these imagined representations. Specifically, we design the ImagineRNN to predict the next visual representation based on past observations in a step-wise manner. Since our target is to predict the future action, it is unnecessary to waste model capacity on resolving the stochasticity of frame appearance changes due to camera motion and shadows in egocentric videos. Thus in ImagineRNN, we only generate the visual representation instead of raw pixels. The final anticipation is built on both the observed content and visual representation that we imagined within the anticipation time T.

Recently, some works [7], [8], [9], [10], [11] also propose to generate intermediate future frames or future content features using RNN or GAN architectures. Most of these works

⁽Corresponding Author: Yi Yang)

Y. Wu, L. Zhu, X. Wang, and Y. Yang are with the Australian Artificial Intelligence Institute, University of Technology Sydney, Ultimo 2007, NSW, Australia. (E-mail: yu.wu-3@student.uts.edu.au; linchao.zhu@uts.edu.au; xiaohan.wang-3@student.uts.edu.au; yi.yang@uts.edu.au).

F. Wu is with Zhejiang University, Hangzhou 310027, China (Email:wufei@cs.zju.edu.cn).

use regression loss functions (*e.g.*, l_2 loss or cosine loss) or discriminator (justifying real or fake) to optimize their generator model. However, these optimization methods are too deterministic in training the generator model. There are only positive targets in these loss functions, leading to biased or sub-optimal optimization on the predicted future features. In addition, since actions are changed very quickly in egocentric videos, the predicted future states should be distinguishable in time sequences. Optimization with only positive targets would overlook the state changes in the future time period.

Our ImagineRNN differs from existing works in two aspects. First is that our ImagineRNN is optimized in the contrastive learning manner instead of feature regression. We propose a proxy task to train the ImagineRNN by selecting the correct future states from distractors. For the predicted future feature, we first build a set of candidates containing both the positive target (the ground truth future feature) and negative distractors (features from other time steps). Then we encourage the model to learn to identify the correct future state from candidates given the observed context. In this way, our ImagineRNN could essentially learn the change of future features. We found the new optimization method significantly improves the generalisability on the unseen test set.

Second, we further improve ImagineRNN by residual anticipation, *i.e.*, changing its target to predicting the feature difference of adjacent frames, instead of the entire frame feature. Different from [10], [11] that predict the entire optical flow frames or dynamic image, we only predict the feature changes between adjacent frames. The motivation is in threefolds. First, the difference between adjacent frame features is more important for forecasting the future. Predicting the video difference promotes the network to focus on the change of intermediate features, leading to better results on the future action anticipation. Second, it reduces the load of the ImagineRNN and thus saves the model capacity. In this way, the information the ImagineRNN has to predict is minimized, while the unchanged feature channels are directly carried forward. Third, the unchanged content plays a role of shortcut connection, avoiding the noise accumulation and the gradient vanishing. To the best of our knowledge, we are the first to forecast the difference of frames in generating future features.

We conduct extensive experiments on two large-scale egocentric video datasets EPIC-KITCHENS [5] and EGTEA Gaze+ [12]. Results from the leaderboard of the EPIC-KITCHENS action anticipation challenge clearly show our model beats other existing single models. To summarize, our contributions are summarized as follows:

- We propose ImagineRNN that breaks down the longtime action anticipation into a series of step-wise feature predictions of short periods, and then predicts the future action labels upon these imagined features.
- We reformulate the future feature prediction problem, and propose to optimize the ImagineRNN by picking the correct future states from lots of distractors, which essentially learns the change of future features compared to the traditional regression loss functions.
- We further replace the ImagineRNN's target by predicting the difference between adjacent frames, which helps the

II. RELATED WORK

A. Video Understanding and Action Recognition

Deep learning methods have achieved promising performance on the video classification task. Simonyan et al. [13] proposed Two-Stream to utilize both RGB frames and optical flow as the 2D CNN input to modeling appearance and motion, respectively. Temporal Segment Networks (TSN) [14] extended the two-stream CNN by extracting features from multiple temporal segments. Tran et al. [15] proposed a 3D CNN to learn the spatial-temporal information. Moreover, Recurrent Neural Networks (RNNs) are also effective in temporal modeling and have been found useful for video classification in [16], [17]. More recently, some researchers study the egocentric action recognition problem [18], [19], [20], [21]. Sudhakaran et al. [19] proposed a Long Short-Term Attention model to focus on features from relevant spatial parts. Wang et al. [21] proposed a Symbiotic Attention mechanism to enable the communications between motion features and object features in egocentric videos. Our method builds on these methods and uses TSN as a base framework to train CNNs for action recognition.

B. Early Action Recognition

The early action recognition task [22], [23], [24], [1] is to recognize the ongoing action as early as possible from partial observations. In this task, the model is only allowed to observe a part of the action videos, and predict the action based on the video segment [25], [26]. This task is closed to our target, the action anticipating task. Differently from these works, in the egocentric anticipating task, the action should be recognized before it starts, so we cannot partially observe the action frames at the time of prediction.

C. Action Anticipation

Predicting the near future has been widely studied recently [27], [28], [29], [30]. Action anticipation is to predict an action before it occurs [31], [32], [33]. Previous works investigated different forms of action and activity anticipation [34], [35], [36], [4], [37], [38], [39], [40]. We share a similar idea with past works and use the recurrent neural networks to summarize the past observations [34], [31]. Very recently, RULSTM [4] consists of two LSTMs to anticipate actions from egocentric video, where one LSTM is used to summarize the past, and the other is used to predict future actions based on the past future directly. Miech et al. [3] proposed to directly anticipate future action based on the combination of past visual inputs and past action recognition results. Concurrent to us, Sener *et al.* [41] propose a multi-scale temporal aggregate method for action anticipation by relating recent to long-range observations. It computes recent and spanning representations pooled from snippets that are related via coupled attention mechanisms. The experiments shows great advantages brought by ensembles of multiple scales.

There are also some other interesting researches for the anticipation task. In [34], the authors study the problem of anticipating a sequence of activities within time horizons of up to 5 minutes, in contrast to other works that anticipate the next action within several seconds. [42] studies anticipating future actions in the long-time period. Given an untrimmed video containing a long composite activity, the proposed topological affordance graphs could predict the actions that will likely occur in the future to complete it. Ego-OMG [43] proposes to structure the long video clips into a discrete set of states, where each state represents the objects presently in contact or anticipated to soon be in contact.

Given past observation, it might have many possible future actions due to future uncertainty. Future uncertainty (alternative future) is important in the action anticipation task. Furnari *et al.* [32] study how to explicitly incorporate the uncertainty in the loss functions. Canuto *et al.* [44] propose to minimize the model uncertainty instead of maximizing its class probabilities, which could be used as the online decisionmaking criterion for action anticipation. In [45], both an action model and a length model are trained to capture the uncertainty of future activities. In this paper, we do not explicitly model the future uncertainty in our method. Given existing video data, we only optimize the model to predict the exact next future action that happens in the video. It is a limitation of our method. We hope to handle uncertainty in our future works.

Some recent works [10], [11] propose to predict the optical flow frames or dynamic image in the future, which has a similar motivation with our designed residual anticipation, *i.e.*, predicting low-entropy signals (the frame-feature difference) However, different from [10], [11] that predict the entire optical flow frames or dynamic image, we only predict the feature changes between adjacent frames, which avoids wasting model capacity on resolving the stochasticity of frame changes due to camera motion in egocentric videos.

D. Contrastive Learning

Contrastive learning aims at optimizing models by distinguishing similar and dissimilar data pairs. Recent works [46], [47], [48] proposed to utilize contrastive learning for selfsupervised learning. Contrastive Predictive Coding (CPC) [48] proposed to learn representation by encoding predictions over future observations from the past. MoCo [49] designed a momentum encoder and maintained a queue of representations to conduct contrastive learning. SimCLR [46] experiments with different combinations of data augmentation methods for paired samples in contrastive learning. Very recently, Han et al. [47] proposed to introduce contrastive learning into the action recognition task. The model is optimized by a predictive attention mechanism over the compressed memories that predicts future representations based on recent observation. Different these methods, we focus on the action anticipation task rather than representation learning. We found the contrastive learning helps to learn the change of future features, which can be used to obtain better intermediate imaginary data in our ImagineRNN framework.

III. PROPOSED APPROACH

A. Egocentric Action Anticipation

Task definition. In the EPIC-Kitchens anticipation challenge [5], the egocentric action anticipation task is defined to predict the future action one second before it happens. In a more general task definition [4], the video is input in an on-line fashion, with a short video snippet consumed every α seconds, *i.e.*, the video is divided into segments of length α . For an action occurring at time τ_s , the model should anticipate the action by observing the video frames before $\tau_s - T$. In our framework, our model is allowed to observe the video segment of length (l-T) starting at time $(\tau_s - l)$ and ending at time $(\tau_s - T)$. Following [4], we use the same task setting and set l = 3.5s and $\alpha = 0.25s$. We also validate our model under different anticipate time, *i.e.*, $T \in \{1.75s, 1.5s, 1.25s, 1s, 0.75s, 0.5s, 0.25s\}$. Note that it is more general compared to the task defined in [5], which only validates the model under anticipate time T = 1.

CNN pretraining. The input of our model is the frame-level feature provided by the pre-trained TSN model. In action anticipation, the anticipation targets (objects and actions) do not always appear in the input video, making it hard to learn good representations for CNN models in an end-to-end manner. To avoid over-fitting and make the CNN model more meaningful, we follow [4] and pre-train the TSN model on the action recognition task. Then the pre-trained CNN weights are fixed during the following training on our action anticipation task. We pre-process the videos and obtain different modalities features by pre-trained CNN models, *i.e.*, RGB frame features, optical flow frame features, and the object features.

Encoder. We take a Long Short-Term Memory (LSTM) [50] model as the temporal encoder. At each time step, the encoder takes as input the visual content that is being observed. Specifically, at each time-step t, we use the pre-trained TSN model to get the current frame feature \mathbf{f}_t . Then we input the feature \mathbf{f}_t to update the memory. The new encoding hidden state \mathbf{h}_{t+1}^E is obtained by updating the LSTM unit as follows:

$$\mathbf{h}_{t+1}^E = \texttt{Encoder}(\mathbf{f}_t, \mathbf{h}_t^E), \tag{1}$$

where \mathbf{h}_t^E is the hidden state from the previous forward. We initialize the hidden state as zeros. To save memory and avoid noises, we only input the frames several seconds before the action occurring time τ_s . Following [4], we take the frames from $(\tau_s - 4)$ s to $(\tau_s - 2.5)$ s as the input for the encoder.

Decoder. The decoder is an LSTM model that performs anticipation. It takes the observed information extracted from the EncodingRNN as the initial hidden states, and then recurrently takes the last observed frame as input. Based on the last output of the DecodingRNN, we use a fully-connected layer as the classifier for the action anticipation prediction.

B. Bridging the gap between past and future

In the egocentric action anticipation task, it is hard to train a meaningful model due to the clear gap between past observations and future action. We alleviate this issue by decomposing the long-time prediction into a series of shortterm forecasts. Then we design ImagineRNN to fill in the



Fig. 2. The framework of our method. ImagineRNN predicts the next visual representation based on past observations in a step-wise manner. The imaginary features are input to the decoder to improve the anticipation performance. We propose to better optimize the ImagineRNN with the contrastive learning task. We further improve the ImagineRNN by forecasting the features difference between frames, instead of generating the entire frame representations.

gap by producing the future visual representation. In this way, the long-time reasoning is simplified by predicting the action based on past observations and future imaginary data.

Specifically, we break down the T seconds anticipation into several short-term anticipations with each lasting α seconds ($\alpha < T$). Given the visual feature \mathbf{f}_t at time t, the ImagineRNN is designed to generate the future visual feature $\hat{\mathbf{f}}_{t+1}$ by,

$$\mathbf{h}_{t+1}^{I} = \text{ImagineRNN}(\mathbf{f}_{t}, \mathbf{h}_{t}^{I}), \qquad (2)$$

$$\hat{\mathbf{f}}_{t+1} = \phi(\mathbf{h}_{t+1}^I),\tag{3}$$

where \mathbf{h}_t^I is the hidden state of ImagineRNN at time step t. $\phi(\cdot)$ is a transformation layer that maps the hidden state space to the visual feature space. The generated visual feature $\mathbf{\hat{f}}_{t+1}$ is supposed to fill in the gap between the past and future. In the framework, we input the output of ImagineRNN to the decoder to predict future action. Thus the prediction of ImagineRNN should be consistent with the ground truth visual content. Next, we illustrate how we optimize the ImagineRNN model efficiently in the action anticipation framework.

C. Optimization of ImagineRNN

In egocentric videos, the action states usually change very quickly. Thus the predicted future from ImagineRNN should be substantially different along with the anticipation time. The commonly used regression loss functions, such as l_2 loss, can hardly optimize the ImagineRNN to perceive the changes of action states. Differently, we propose a more effective optimization for the ImagineRNN by introducing the contrastive learning task, where the model is asked to pick the correct future states from lots of distractors. We use Noise Contrastive Estimation (NCE) [51] to encourage the predicted future feature $\hat{\mathbf{f}}_{t+1}$ to be close to the ground truth future state $\hat{\mathbf{f}}_{t+1}$. Compared to the regression losses, NCE does not require to resolve the low-level stochasticity strictly. Specifically, for the imagined future feature $\hat{\mathbf{f}}_t$ at time t, the only positive target

is the ground truth feature \mathbf{f}_t . We then build a set of candidates as distractors for the ground truth feature \mathbf{f}_t at time t.

Distractors. The distractors contain easy negatives and hard negatives. The easy negatives contain the frame features from the other videos instead of the target video. We use the framelevel features from the other videos in the same mini-batch as the easy negatives for simplicity in the calculation. These candidates are easy to distinguish since these frames usually look different from the current video.

The hard negatives contain the frames from the same video but at different time steps, \mathbf{f}'_t where $t' \neq t$. These candidates are hard to distinguish since they are very close to the ground truth frame feature \mathbf{f}_t . Distinguishing the hard negatives encourages ImagineRNN to generate essential intermediate features and capture the change of a series of future states.

Contrastive Learning. With the positive targets and these distractors, we can take the contrastive learning as a proxy task for better optimizing the ImagineRNN. Inspired by recent representation learning work [52], [53], we first calculate the cosine similarity between the predicted feature and the candidates, $\mathbf{v}_j^{T} \hat{\mathbf{f}}_t$, where \mathbf{v}_j denotes the *j*-th distractors. Here we enforce all vectors to be L2-normalized feature embeddings, *i.e.*, $||\mathbf{v}_j|| = 1$, $||\hat{\mathbf{f}}_t|| = 1$, and $||\mathbf{f}_t|| = 1$. Thus we have the following objective function at the time step *t*,

$$\mathcal{L}_{c} = -\log \frac{\exp(\mathbf{f}_{t}^{\mathsf{T}} \hat{\mathbf{f}}_{t}/\tau)}{\sum_{j} \exp(\mathbf{v}_{j}^{\mathsf{T}} \hat{\mathbf{f}}_{t}/\tau) + \exp(\mathbf{f}_{t}^{\mathsf{T}} \hat{\mathbf{f}}_{t}/\tau)}, \qquad (4)$$

where τ is a temperature parameter that controls the concentration level of the distribution. Higher τ leads to a softer probability distribution. We set $\tau = 0.2$ in our experiments.

With Eqn. (4), we optimize the ImagineRNN with a crossentropy loss (negative log-likelihood), instead of the commonly used regression loss functions. During optimization, the loss function encourages the predicted feature $\hat{\mathbf{f}}_t$ to be close to ground truth target \mathbf{f}_t , and also push the predicted feature $\hat{\mathbf{f}}_t$ to be distinct from these distractors. Thus the trained ImagineRNN could catch the change of action states at different times, which is essential in action anticipation.

2) The *future intention*. In addition, following [9], [10], [11], we also take the future intentions as additional supervision. The future intention is the purpose of the currently observed actions (the next future action), which explains the visual changes that happen during the unseen temporal region T. The intuition behinds it is that the generated visual representation should also benefit the anticipation task. Specifically, we input the generated visual feature to the decoder for several time steps during the anticipation time period. The decoder's last hidden state is further input to the action classifier for recognizing the future action. Then we use the Cross-Entropy loss on the final action anticipation to optimize the ImagineRNN. Denote the Cross-Entropy loss of the classifier as \mathcal{L}_f , the final loss is the sum of the two losses,

$$\mathcal{L} = \mathcal{L}_c + \mathcal{L}_f. \tag{5}$$

D. Forecasting the difference between frames

However, the visual features of adjacent frames would be close since the backgrounds in frames are the same. Directly predicting the visual feature of future frames might waste model capacity in generating the unchanged background information. In addition, ImagineRNN might not essentially learn the change during future frames. Thus we propose to improve ImagineRNN by explicitly force it to predict the feature difference of adjacent frames, instead of the entire frame feature.

Specifically, we optimize ImagineRNN by learning to produce the difference between the current visual feature and the next one. The output of ImagineRNN is to forecast future *changes* of the visual feature given the current observation. Thus we change Eqn. (3) to be,

$$\hat{\mathbf{f}}_{t+1} = \phi(\mathbf{h}_{t+1}^I) + \mathbf{f}_t. \tag{6}$$

Since we are designed to predict a series of intermediate frame features before anticipating the future action, we repeatedly use Eqn. (6) to generate a series of future frame features in an auto-regressive way. Suppose frame t to be the last observed frame, we can obtain the imagined feature $\hat{\mathbf{f}}_{t+n}$ of future frame t + n by,

$$\hat{\mathbf{f}}_{t+n} = \phi(\mathbf{h}_{t+n}^{I}) + \phi(\mathbf{h}_{t+n-1}^{I}) + \dots + \phi(\mathbf{h}_{t+1}^{I}) + \mathbf{f}_{t}.$$
 (7)

As can be seen in the above equation, predicting the difference sets up a shortcut connection between step-wise reconstructions, which helps ease the optimization of ImagineRNN and avoids noise accumulation during the autoregressive future feature generation in testing. In addition, Predicting the frame difference promotes the model to focus on the change of intermediate features, which might be the core of future action anticipation.

IV. EXPERIMENTS

We first discuss the experimental setups and then compare our method with the state-of-the-art methods on two largescale egocentric action datasets, EPIC-Kitchens and EGTEA Gaze+. Ablation studies and qualitative results are provided to show the effectiveness of our method.

A. Experimental Settings

Datasets. We perform experiments on two large-scale datasets of egocentric videos: EPIC-Kitchens [5] and EGTEA Gaze+ [12]. EPIC-Kitchens is the largest dataset in firstperson vision so far. It consists of 55 hours of recordings capturing all daily activities in the kitchens. The activities performed are non-scripted, which makes the dataset very challenging and close to real-world data. This dataset is densely annotated with timestamps for each action so that it is ready for the action anticipation task. Actions in the EPIC-Kitchens dataset is annotated in the format of (verb, noun) pairs. The dataset contains 39, 596 action annotations, 125 verbs, and 352 nouns. We considered all unique (verb, noun) pairs in the public training set, thus obtaining 2,513 unique actions. We use the same split as [4] and split the public training set of EPIC-Kitchens (28, 472 action segments) into training (23, 493 segments) and validation (4,979 segments) sets. EGTEA Gaze+ contains 19 verbs, 51 nouns and 106 unique actions. We report the average performance across the three official splits provided by the authors of the dataset.

Evaluation Metrics. Following [4], we use the Top-k accuracy to evaluate our method. Under this evaluation metric, the prediction is deemed to be correct if the ground truth action falls in the top-k predictions. This metric is appropriate due to the uncertainty of future predictions [32], [54]. Many possible actions can follow an observation. We use the Top-5 accuracy as a class-agnostic measure. We also report Mean Top-5 Recall [32] as a class aware metric. Top-5 recall for a given class c is defined as the fraction of samples of ground truth class c for which the class c is in the list of the top-5 anticipated actions. Mean Top-5 Recall averages Top-5 recall values over classes. In [32], Top-5 Recall is averaged over the provided list of many-shot verbs, nouns, and actions. Performances are evaluated for verb, noun, and action predictions. Following [4], in training the only targets are the action labels, and our model is optimized to predict the action prediction. In the testing, we obtain the predictions for verb and noun by the marginalization on action predictions.

Implementation Details. We use Pytorch [55] to implement our framework. For the pre-trained action recognition model, we use a BNInception CNN [56] with the TSN framework to train the action recognition model. After pre-training, we resize the frame to 456×256 pixels and input it to the CNN model. The output (1024-dimensional vectors) of the last global average pooling layer is used as the frame-level feature. The encoder, decoder, and the ImagineRNN are all single-layer LSTMs with 1024 hidden units. We use Stochastic Gradient Descent (SGD) to train the framework with a mini-batch size of 128 and a learning rate of 0.01 and momentum equal to 0.9. We train 100 epochs and apply early stopping at each training stage the same as [4]. This is done by choosing the intermediate and final models' iterations, which obtain the best Top-5 action anticipation accuracy for the anticipation time T = 1s on the validation set. Following [4], we use the RGB frames, optical flow frames, and the object detection features as input for our model. We first train the model with each modality individually and then obtain the final prediction by a

TABLE I

Egocentric action anticipation results on the Seen (S1) test set of the EPIC-KITCHENS Action Anticipation Challenge [5] with anticipation time T = 1 second. All values are reported as percentage (%).

Methods	Verb		Nc	oun	Action	
	Top-1 Acc	Top-5 Acc	Top-1 Acc	Top-5 Acc	Top-1 Acc	Top-5 Acc
2SCNN [5]	29.76	76.03	15.15	38.56	04.32	15.21
ATSN [5]	31.81	76.56	16.22	42.15	06.00	28.21
ED [31]	29.35	74.49	16.07	38.83	08.08	18.19
MCE [32]	27.92	73.59	16.09	39.32	10.76	25.28
Transitional [3]	30.74	76.21	16.47	42.72	09.74	25.44
RULSTM [4]	33.04	79.55	22.78	50.95	14.39	33.73
Ours $(l_2 \text{ loss})$	35.26	79.66	22.57	52.04	15.07	34.66
Ours (Contrastive)	35.44	79.72	22.79	52.09	14.66	34.98

TABLE II

Egocentric action anticipation results on the **Unseen (S2)** test set of the EPIC-KITCHENS Action Anticipation Challenge [5] with anticipation time T = 1 second. All values are reported as percentage (%).

Methods	Verb		No	oun	Action	
	Top-1 Acc	Top-5 Acc	Top-1 Acc	Top-5 Acc	Top-1 Acc	Top-5 Acc
2SCNN [5]	25.23	68.66	09.97	27.38	02.29	09.35
ATSN [5]	25.30	68.32	10.41	29.50	02.39	06.63
ED [31]	22.52	62.65	07.81	21.42	02.65	07.57
MCE [32]	21.27	63.33	09.90	25.50	05.57	15.71
Transitional [3]	28.37	69.96	12.43	32.20	07.24	19.29
RULSTM [4]	27.01	69.55	15.19	34.38	08.16	21.10
Ours $(l_2 \text{ loss})$	27.35	69.78	15.36	35.34	08.54	20.79
Ours (Contrastive)	29.33	70.67	15.50	35.78	09.25	22.19

late fusion of the three models' predictions. In the following experiments, for fair comparisons with RULSTM, our model takes all the three modalities as input if not specified.

TABLE III Action anticipation results on the EPIC-KITCHENS validation set under different anticipation time T. The performance is measured by the top-5 accuracy of action anticipation.

B. Comparison to the state-of-the-art methods

Compared Methods We compare our method with stateof-the-art action anticipation methods: Deep Multimodal Regressor (DMR) [57], Anticipation Temporal Segment Network (ATSN) of [5], Anticipation Temporal Segment Network trained with verb-noun Marginal Cross Entropy Loss (MCE) [32], and the Encoder-Decoder LSTM (ED) introduced in [31]. We also compare with the early action recognition methods to the problem of egocentric action anticipation: Feedback Network LSTM (FN) [25], and an LSTM trained using the Exponential Anticipation Loss [58] (EL). To compare with state-of-the-art action anticipation methods, we reproduced a vanilla version of Feature Mapping RNN [9] without the kernalised RBF. For a fair comparison, we first train models with the three input modalities, i.e., RGB features, optical flow features, and the object features. Then we obtain the final prediction by a late fusion of the three models. Very recently, RULSTM [4] is proposed by combining two LSTM to anticipate actions from egocentric video, where one LSTM is used to summarize the past, and the other is used to predict future actions based on the past future. We compare our method under both the standard anticipation setting (anticipation time T = 1s) and a more general anticipation setting (with variant anticipation time).

Results on the EPIC-KITCHENS test server. We compare our method with the state-of-the-art methods on the test server

Methods	Тор	-5 Actio	on Accu	racy @	differen	nt T
wiedious	1.5	1.25	1.0	0.75	0.5	0.25
DMR [57]	/	/	16.9	/	/	/
ATSN [5]	/	/	16.3	/	/	/
MCE [32]	/	/	26.1	/	/	/
FMRNN [9]	/	/	32.7	/	/	/
ED [31]	23.2	24.8	25.8	26.7	27.7	29.7
FN [25]	24.7	25.7	26.3	26.9	27.9	29.0
EL [58]	26.4	27.4	28.6	30.3	31.5	33.6
RULSTM [4]	32.2	33.4	35.3	36.3	37.3	39.0
Ours	32.5	33.6	35.6	36.7	38.5	39.4

of EPIC-KITCHENS. Table I and Table II report results obtained from the official EPIC-KITCHENS action anticipation challenge submission server. The official test server computes the performances on two test sets, *i.e.*, the "seen" test, which includes the same scenes appearing in the training set (S1) and the "unseen" test set (S2), with kitchens not appearing in the training set. On both test sets, our method outperforms all previously reported results under all metrics. On the S1 (seen) test set (Table I), our method outperforms the previous method RULSTM by 1.25% on the Top-5 Action accuracy. On the S2 (unseen) test set (Table II) where the videos are captured in new environments, our method significantly improves RULSTM in all metrics on Verb, Noun, and Action prediction. Note that we use the same input features with RULSTM, thus the comparison with RULSTM is a fair comparison, and the performance improvements over RULSTM are all from our

TABLE IV COMPARISON OF THE ANTICIPATED ACTION ACCURACIES WITH DIFFERENT MODALITIES ON THE VALIDATION SET. "OBJ" INDICATES THE OBJECT FEATURES. "W/O INTENTION" IS THE MODEL OPTIMIZED WITHOUT THE FUTURE INTENTION (EQN. (5)). "W/O DIFF" INDICATES THE MODEL WITHOUT FORECASTING THE DIFFERENT.

Modality	Method	Top-1 Acc	Top-5 Acc
	RULSTM [4]	13.05	30.83
DCD	Ours w/o intention	13.23	31.39
KUD	Ours w/o diff	12.97	30.61
	Ours	13.68	31.58
	RULSTM [4]	08.77	21.42
Flow	Ours w/o intention	08.81	21.89
FIOW	Ours w/o diff	08.51	21.68
	Ours	09.23	22.06
Ohi	RULSTM [4]	10.04	29.89
	Ours w/o intention	10.76	30.05
Obj	Ours w/o diff	10.62	30.12
	Ours	10.72	30.27
	RULSTM [4]	15.00	35.24
Fusion	Ours w/o intention	15.04	35.17
	Ours w/o diff	14.91	34.98
	Ours	15.23	35.38

algorithm instead of better features. These results demonstrate our method is better at anticipating the future action.

Results with Different Anticipation Time T. Our method can also be used to predict future action under different anticipation time. Since each time step α in our method is 0.25s, we can evaluate the future anticipation every 0.25s. We compare our method with the state-of-the-art methods under different anticipation time $T \in \{2s, 1.75s, 1.5s, 1.25s, 1s,$ 0.75s, 0.5s, 0.25s $\}$. The results are shown in Table III. Note that some methods [5], [32], [57] can anticipate actions only at a fixed anticipation time. We found the proposed method always outperforms the strong competitor RULSTM [4] under all anticipation time T. Note that the results are reported on the validation set, where the models are selected by choosing the best validation performance, as used by RULSTM [4]. As indicated in [4], the results on the test server are more important in evaluating compared to the validation results.

Results on the EGTEA Gaze+ dataset. We also conduct experiments on the EGTEA Gaze+ dataset. Table V reports Top-5 action accuracy scores on EGTEA Gaze+ under different anticipation times. We use the same input modalities as RULSTM. Our method outperforms the compared methods under different anticipation time T. We also found the relative improvement is smaller on the EGTEA Gaze+ dataset compared to that on the EPIC-KITCHENS dataset. It might be because the EGTEA Gaze+ is relatively small in scale. It only consists 106 actions, which is far less than the 2,513 actions in EPIC-KITCHENS. Thus the anticipation on the EPIC-KITCHENS dataset is more challenging.

C. Ablation Studies

We conduct ablation studies to evaluate the effectiveness of the two components of our method.

Effectiveness of ImagineRNN. Without our proposed ImagineRNN, the model is the baseline RULSTM. From Table I and Table II, we can see the results of our baseline model



Fig. 3. Top-5 accuracies over different lengths of observed past for the encoder. The results are produced by our method with the RGB modality input. Note the anticipation time T is 1s for all experiments.

 TABLE V

 ANTICIPATION RESULTS ON THE EGTEA GAZE+ DATASET.

	Top-5 Action Accuracy% @ different T							
	1.0	0.75	0.5	0.25				
DMR [57]	55.70	/	/	/				
ATSN [5]	40.53	/	/	/				
MCE [32]	56.29	/	/	/				
ED [31]	50.22	51.86	49.99	49.17				
FN [25]	60.12	62.03	63.96	66.45				
RL [26]	62.56	64.65	67.35	70.42				
EL [58]	64.62	66.89	69.60	72.38				
RULSTM [4]	66.40	68.41	71.84	74.28				
Ours	66.71	68.54	72.32	74.59				

only achieve 33.73% in Top-5 accuracy on the seen (S1) test set and 21.10% on the unseen (S2) test set. By adding our ImagineRNN to the framework, we observed a clear performance improvement on both test sets.

Effectiveness of Contrastive Learning. The common used optimization for the ImagineRNN is the regression loss functions (l_2 loss). In Table I and Table II, we show the comparison of different optimization methods on the test set of EPIC-KITCHENS. Ours $(l_2 \text{ loss})$ indicates our models optimized by the l_2 loss, while Ours (Contrastive) is the model optimized in the Contrastive Learning way, i.e., picking the correct one from lots of distractors. With l_2 loss, our model achieves 34.66 on Top-5 Accuracy in the seen test set. In contrast, with Contrastive Learning, our method achieves 34.98% on Top-5 action accuracy. The improvement of Contrastive Learning is more clear in the unseen test set. With the proposed Contrastive Learning, the action anticipation result on the unseen set shows a 1.40% (22.19% versus 20.79%) improvement on the Top-5 action accuracy. The significant performance gap shows that contrastive learning is a better way to optimize ImagineRNN. It leads to a better generalisability across the various benchmarks.

Effectiveness of Forecasting the Difference. In Table IV, we show the comparison of results with and without forecasting the difference. We conduct ablation studies on the RGB input, the flow input, and the fused modalities input. The results show a steady improvement by introducing to forecast the difference. Specifically, our method significantly outperforms the one (w/o diff) by 0.9 points on the Top-5 action accuracy



Fig. 4. Qualitative results with anticipation time T = 2s, 1.5s, 1s, 0.5s. From left to right, the observations are getting closer to future action. We list the Top-5 action predictions of our model at each anticipation time. Orange indicates the ground truth, and green means our prediction matches the ground truth.

on the RGB modality. Similarly, we found our approach also suppresses the model (w/o diff) with optical flow data as inputs. These comparison results prove the effectiveness of forecasting the difference instead of directly generating the whole visual feature. We also validate the effectiveness of forecasting the difference with other architectures. We replace the basic architectures of our ImagineRNN and the encoder-decoder by Gated Recurrent Unit (GRU), instead of the previously used LSTM. The results are shown in Table VII. It can be seen that our predicting the feature difference of adjacent frames still performs better with the GRU-based architecture.

Effectiveness of Future Intention. We also show the comparison of results with and without the future intention optimization Eqn. (5) in Table IV. The ablation studies show a small improvement brought by future intention. Specifically, our final model outperforms the one without future intension on the RGB and flow modalities by about 0.4% in Top-1 accuracy and 0.2% in Top-5 accuracy.

Ablation studies over different lengths of the past. We show the results over different lengths of observed past in Figure 3. Note the anticipation time T is 1s for all experiments. It can be seen from the figure that the performance is relatively low if the encoder period is too short (i.e., less than 2.25 seconds). As the encoding period gets longer, we found the performance gets steady. Inputting more observed frames did not lead to further performance improvement if the encoder period is longer than 2.5 seconds. The reason might be that actions usually change quickly in egocentric videos. Too early frames do not have strong correlations with the future action. Discussion on different optimization methods. We evaluate the models with different optimization methods on the test set of the EPIC-KITCHENS Action Anticipation Challenge. The results are shown in Table VI, where "Con." indicates the contrastive learning loss, and "Adv." is the adversarial loss used in GAN [10]. It can be seen that a combination of contrastive loss and l_2 loss does not outperform the one with the contrastive learning only. Besides, we add the adversarial loss

TABLE VI Comparison of different optimizations on the Unseen (S2) test set of the EPIC-KITCHENS Action Anticipation Challenge. "Con." indicates the contrastive learning loss. "Adv." indicates the adversarial loss used in GAN [10].

Methods	Verb		No	oun	Action	
Wiethous	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
l_2	27.4	69.8	15.4	35.3	8.5	20.8
$Con. + l_2 + Adv.$	27.9	70.3	14.3	34.7	8.5	20.7
Con. + l_2	28.4	70.0	15.1	34.9	9.0	21.1
Con.	29.3	70.7	15.5	35.8	9.3	22.2

TABLE VII Ablation studies of predicting the feature difference between adjacent frames with GRU-based architecture on the EPIC-KITCHENS Action Anticipation validation set.

Methods	Verb		Noun		Action	
Wiethous	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
Ours (GRU w/o diff)	32.9	78.7	22.0	49.2	13.3	32.3
Ours (GRU with diff)	33.7	79.7	22.7	50.2	14.0	33.2

in the model training, where the discriminator is a three-layer MLP. According to the validation results, we set the weight of the adversarial loss to be 0.01 in the overall loss function and the discriminator's learning rate to be 2×10^{-6} . As shown in Table VI, the model trained with the combination of the three loss functions performs worst among all candidates. Our model trained with contrastive learning performs best among all candidates. The reason might be that contrastive learning helps to learn the change of future features essentially, since it needs to distinguish the positive target from lots of distractors. (frame features at other times).

D. Qualitative Results

We show some qualitative results of our method in Fig 4. From left to right, the observations are getting closer to future action. The orange box and "GT" indicates the ground truth of the future action. We list the Top-5 action predictions of our results at the anticipation time $T \in \{2s, 1.5s, 1s, 0.5s\}$. Green indicates the prediction matches the ground truth. Taking the first row as an example, the anticipations become more and more accurate as time flows. It is consistent with our motivation that long-time modeling might involve lots of noise. It is interesting to see the model always predicts "Open fridges" when T is less than 2 seconds probably because the fridge shows up in the observations at T = 1.5s. The other action candidates, including "Take milk" and "Open drawer", are also likely to take place in the near future.

V. CONCLUSIONS AND FUTURE WORKS

In this work, we decompose the action anticipation task into a series of future frame feature predictions. We first imagine how the future feature changes and then predict future action based on these imagined representations. We found that ImagineRNN optimized with contrastive learning is superior than the typical anticipation models. In addition, we further propose to improve ImagineRNN by predicting the feature difference of adjacent frames instead of the whole frame content. It helps promote the model to focus on the change of future states and avoid the noise accumulation during the autoregressive future feature generation. Extensive experimental results on different architectures validate the effectiveness of the proposed method.

In conclusion, we found it useful by decomposing action anticipation into lots of intermediate predictions. Focusing on the future state transition by contrastive learning and predicting future frames' difference improves the quality of intermediate predictions, leading to better results on the final action anticipation task.

In future works, we might further explore the uncertainty of future in the egocentric action anticipation task, which is a limitation of our current work. It might also be helpful in generating better intermediate future features by considering the hand tracking on the observed frames.

VI. ACKNOWLEDGEMENT

Yu Wu is the recipient of the Google PhD Fellowship 2020. We acknowledge the Google PhD Fellowship Programme and ARC Discovery Project DP200100938 for funding this research.

REFERENCES

- L. Chen, J. Lu, Z. Song, and J. Zhou, "Part-activated deep reinforcement learning for action prediction," in *ECCV*, 2018, pp. 421–436. 1, 2
- [2] Y. Li, "Which way are you going? imitative decision learning for path forecasting in dynamic scenes," in CVPR, 2019, pp. 294–303. 1
- [3] A. Miech, I. Laptev, J. Sivic, H. Wang, L. Torresani, and D. Tran, "Leveraging the present to anticipate the future in videos," in *CVPR-W*, 2019. 1, 2, 6
- [4] A. Furnari and G. M. Farinella, "What would you expect? anticipating egocentric actions with rolling-unrolling lstms and modality attention." in *ICCV*, 2019. 1, 2, 3, 5, 6, 7
- [5] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray, "Scaling egocentric vision: The epic-kitchens dataset," in *ECCV*, 2018. 1, 2, 3, 5, 6, 7
- [6] Y. Tang, Z. Wang, J. Lu, J. Feng, and J. Zhou, "Multi-stream deep neural networks for rgb-d egocentric action recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 29, no. 10, pp. 3001– 3015, Oct 2019. 1

- [7] W. Liu, W. Luo, D. Lian, and S. Gao, "Future frame prediction for anomaly detection-a new baseline," in CVPR, 2018, pp. 6536–6545. 1
- [8] P. Luc, C. Couprie, Y. Lecun, and J. Verbeek, "Predicting future instance segmentation by forecasting convolutional features," in *ECCV*, 2018, pp. 584–599. 1
- [9] Y. Shi, B. Fernando, and R. Hartley, "Action anticipation with rbf kernelized feature mapping rnn," in ECCV, 2018, pp. 301–317. 1, 5, 6
- [10] H. Gammulle, S. Denman, S. Sridharan, and C. Fookes, "Predicting the future: A jointly learnt model for action anticipation," in *ICCV*, 2019, pp. 5562–5571. 1, 2, 3, 5, 8
- [11] C. Rodriguez, B. Fernando, and H. Li, "Action anticipation by predicting future dynamic images," in ECCV-W, 2018. 1, 2, 3, 5
- [12] Y. Li, M. Liu, and J. M. Rehg, "In the eye of beholder: Joint learning of gaze and actions in first person video," in ECCV, 2018. 2, 5
- [13] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *NIPS*, 2014. 2
- [14] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Val Gool, "Temporal segment networks: Towards good practices for deep action recognition," in ECCV, 2016. 2
- [15] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*, 2015.
- [16] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, "Youtube-8m: A large-scale video classification benchmark," arXiv preprint arXiv:1609.08675, 2016. 2
- [17] L. Zhu, Z. Xu, and Y. Yang, "Bidirectional multirate reconstruction for temporal modeling in videos," in CVPR, 2017. 2
- [18] E. Kazakos, A. Nagrani, A. Zisserman, and D. Damen, "Epic-fusion: Audio-visual temporal binding for egocentric action recognition," in *ICCV*, 2019. 2
- [19] S. Sudhakaran, S. Escalera, and O. Lanz, "Lsta: Long short-term attention for egocentric action recognition," in CVPR, 2019. 2
- [20] S. Sudhakaran and O. Lanz, "Attention is all we need: Nailing down object-centric attention for egocentric activity recognition," *BMVC*, 2018. 2
- [21] X. Wang, L. Zhu, Y. Wu, and Y. Yang, "Symbiotic attention for egocentric action recognition with object-centric alignment," *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 2020. 2
- [22] M. Sadegh Aliakbarian, F. Sadat Saleh, M. Salzmann, B. Fernando, L. Petersson, and L. Andersson, "Encouraging lstms to anticipate actions very early," in *ICCV*, 2017. 2
- [23] F. Becattini, T. Uricchio, L. Seidenari, A. Del Bimbo, and L. Ballan, "Am I done? predicting action progress in videos," *BMVC*, 2017. 2
- [24] R. De Geest, E. Gavves, A. Ghodrati, Z. Li, C. Snoek, and T. Tuytelaars, "Online action detection," in ECCV, 2016. 2
- [25] R. De Geest and T. Tuytelaars, "Modeling temporal structure with lstm for online action detection," in WACV, 2018. 2, 6, 7
- [26] S. Ma, L. Sigal, and S. Sclaroff, "Learning activity progression in lstms for activity detection and early detection," in CVPR, 2016. 2, 7
- [27] Y. Li, "A deep spatiotemporal perspective for understanding crowd behavior," *IEEE Transactions on Multimedia*, vol. 20, no. 12, pp. 3289– 3297, 2018. 2
- [28] J. Liang, L. Jiang, J. C. Niebles, A. G. Hauptmann, and L. Fei-Fei, "Peeking into the future: Predicting future person activities and locations in videos," in CVPR, 2019, pp. 5725–5734. 2
- [29] Y. Wang, L. Jiang, M.-H. Yang, L.-J. Li, M. Long, and L. Fei-Fei, "Eidetic 3d lstm: A model for video prediction and beyond," in *ICLR*, 2018. 2
- [30] B. F. Yan Bin Ng, "Forecasting future action sequences with attention: a new approach to weakly supervised action forecasting," *IEEE Transactions on Image Processing*, 2020. 2
- [31] J. Gao, Z. Yang, and R. Nevatia, "Red: Reinforced encoder-decoder networks for action anticipation," *BMVC*, 2017. 2, 6, 7
- [32] A. Furnari, S. Battiato, and G. M. Farinella, "Leveraging uncertainty to rethink loss functions and evaluation measures for egocentric action anticipation," in *ECCV-W*, 2018. 2, 3, 5, 6, 7
- [33] Q. Ke, M. Bennamoun, H. Rahmani, S. An, F. Sohel, and F. Boussaid, "Learning latent global network for skeleton-based action prediction," *IEEE Transactions on Image Processing*, vol. 29, pp. 959–970, 2020. 2
- [34] Y. Abu Farha, A. Richard, and J. Gall, "When will you do what?anticipating temporal occurrences of activities," in CVPR, 2018. 2, 3
- [35] P. Felsen, P. Agrawal, and J. Malik, "What will happen next? forecasting player moves in sports videos," in *ICCV*, 2017. 2
- [36] A. Furnari, S. Battiato, K. Grauman, and G. M. Farinella, "Nextactive-object prediction from egocentric videos," *Journal of Visual Communication and Image Representation*, vol. 49, pp. 401–411, 2017.

- [37] T. Mahmud, M. Hasan, and A. K. Roy-Chowdhury, "Joint prediction of activity labels and starting times in untrimmed videos," in *ICCV*, 2017.
- [38] K.-H. Zeng, W. B. Shen, D.-A. Huang, M. Sun, and J. Carlos Niebles, "Visual forecasting by imitating dynamics in natural sequences," in *ICCV*, 2017. 2
- [39] N. Rhinehart and K. M. Kitani, "First-person activity forecasting with online inverse reinforcement learning," in *ICCV*, 2017. 2
- [40] S. Huang, X. Li, Z. Zhang, Z. He, F. Wu, W. Liu, J. Tang, and Y. Zhuang, "Deep learning driven visual path prediction from a single image," *IEEE Transactions on Image Processing*, vol. 25, no. 12, pp. 5892–5904, Dec 2016. 2
- [41] F. Sener, D. Singhania, and A. Yao, "Temporal aggregate representations for long-range video understanding," in ECCV, 2020. 2
- [42] T. Nagarajan, Y. Li, C. Feichtenhofer, and K. Grauman, "Ego-topo: Environment affordances from egocentric video," in CVPR, 2020, pp. 163–172. 3
- [43] E. Dessalene, M. Maynord, C. Devaraj, C. Fermuller, and Y. Aloimonos, "Egocentric object manipulation graphs," arXiv preprint arXiv:2006.03201, 2020. 3
- [44] C. C. d. Santos, P. Moreno, J. L. A. Samatelo, R. F. Vassallo, and J. Santos-Victor, "Action anticipation for collaborative environments: The impact of contextual information and uncertainty-based prediction," *Neurocomputing*, 2019. 3
- [45] Y. Abu Farha and J. Gall, "Uncertainty-aware anticipation of activities," in *ICCV-W*, 2019. 3
- [46] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICLR*, 2020. 3
- [47] T. Han, W. Xie, and A. Zisserman, "Memory-augmented dense predictive coding for video representation learning," in ECCV, 2020. 3
- [48] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," arXiv preprint arXiv:1807.03748, 2018.
- [49] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *CVPR*, 2020, pp. 9729–9738. 3
- [50] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997. 3
- [51] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *AISTATS*, 2010, pp. 297–304. 4
- [52] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *CVPR*, 2018, pp. 3733– 3742. 4
- [53] Y. Wu, Y. Lin, X. Dong, Y. Yan, W. Bian, and Y. Yang, "Progressive learning for person re-identification with one example," *IEEE Transactions on Image Processing*, vol. 28, no. 6, pp. 2872–2881, 2019. 4
- [54] H. S. Koppula and A. Saxena, "Anticipating human activities using object affordances for reactive robotic response," *T-PAMI*, vol. 38, no. 1, pp. 14–29, 2015. 5
- [55] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017. 5
- [56] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *ICML*, 2015.
- [57] C. Vondrick, H. Pirsiavash, and A. Torralba, "Anticipating visual representations from unlabeled video," in CVPR, 2016. 6, 7
- [58] A. Jain, A. Singh, H. S. Koppula, S. Soh, and A. Saxena, "Recurrent neural networks for driver activity anticipation via sensory-fusion architecture," in *ICRA*, 2016. 6, 7