

Multi-Agent Reinforcement Learning Based Frame Sampling for Effective Untrimmed Video Recognition

Wenhao Wu^{1,2*} Dongliang He³ Xiao Tan³ Shifeng Chen^{1†} Shilei Wen³
¹ Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China
² University of Chinese Academy of Sciences, China
³ Department of Computer Vision Technology (VIS), Baidu Inc., China

{wh.wu, shifeng.chen}@siat.ac.cn

{hedongliang01, tanxiao01, wenshilei}@baidu.com

Abstract

Video Recognition has drawn great research interest and great progress has been made. A suitable frame sampling strategy can improve the accuracy and efficiency of recognition. However, mainstream solutions generally adopt hand-crafted frame sampling strategies for recognition. It could degrade the performance, especially in untrimmed videos, due to the variation of frame-level saliency. To this end, we concentrate on improving untrimmed video classification via developing a learning-based frame sampling strategy. We intuitively formulate the frame sampling procedure as multiple parallel Markov decision processes, each of which aims at picking out a frame/clip by gradually adjusting an initial sampling. Then we propose to solve the problems with multi-agent reinforcement learning (MARL). Our MARL framework is composed of a novel RNN-based context-aware observation network which jointly models context information among nearby agents and historical states of a specific agent, a policy network which generates the probability distribution over a predefined action space at each step and a classification network for reward calculation as well as final recognition. Extensive experimental results show that our MARL-based scheme remarkably outperforms hand-crafted strategies with various 2D and 3D baseline methods. Our single RGB model achieves a comparable performance of ActivityNet v1.3 champion submission with multi-modal multi-model fusion and new state-of-the-art results on YouTube Birds and YouTube Cars.

1. Introduction

Recently, video recognition has attracted great research interest in the computer vision community, due to its importance in real-world applications such as video surveillance,

*This work was done when Wenhao Wu was a research intern at Baidu.

†Corresponding author (e-mail: shifeng.chen@siat.ac.cn).

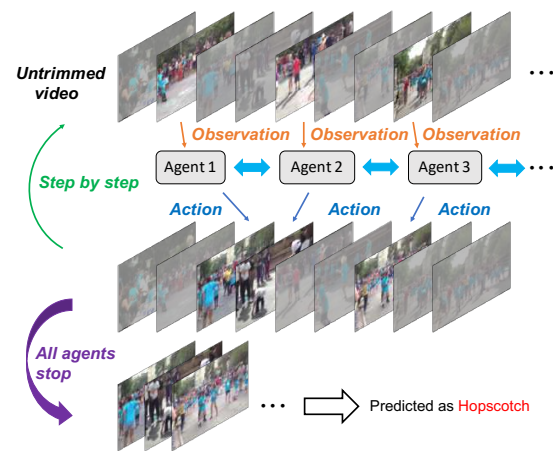


Figure 1. A demonstration of our approach. At each time step, each agent relies on context information among nearby agents to take action for adjusting its sampling location. When all agents stop, the prediction is emitted by the classification decision.

video search, and video recommendation. A video contains a sequence of frames, both spatial information, and temporal relation are important for accurate recognition. Compared to recognizing well-trimmed clips, untrimmed videos pose a more critical challenge since not all the frames consistently respond to the specified ground-truth label. Picking out the most informative frames can be an effective method for recognition.

Existing research efforts [45, 36, 12, 10, 13, 7, 11, 41, 3, 52, 33, 47, 8, 1, 16, 29, 35, 42, 44, 56, 5, 20] mainly focus on building effective and efficient video modeling networks. Generally speaking, they can be divided into two directions, namely, (1) two-stage solutions [8, 1, 16, 29, 35] which extract spatial feature vectors from video frames and then integrate the obtained local feature sequence into one compact video descriptor for recognition; (2) the 2D [45, 36, 10, 7] or 3D convolution based [41, 11, 3, 52, 33, 47, 42, 44, 56, 5]

end-to-end video classification methods. Though great progress has been achieved by these methods, limited attention is paid to the aforementioned variation of frame-level saliency among different frames. The mainstream methods all propose to sample input frames by hand-crafted strategies, for instance, evenly sampling N frames/clips from the original video is done in [45, 20, 29], sampling N successive frames from a video is adopted by [41, 33, 47], and directly feeding all video frames in the test phase is chosen in [3, 52, 1, 35]. Sampling N frames/clips either evenly or successively from an untrimmed video cannot guarantee optimum. Meanwhile, feeding all the frames is a sort of brute-force and introduces a much unnecessary computation burden. Therefore, in this paper, we concentrate on how to pick out the most discriminative frames from untrimmed videos to achieve better classification performance.

Intuitively, an effective algorithm for humans to pick out N representative frames from an untrimmed video can be as follows: we first observe N scenes of the whole video, according to the initial observation, then we infer where to check next time until we find out the satisfying frames round by round. Motivated by this, we proposed to formulate frame selection as N parallel Markov decision processes. As is known, a Markov decision process can be naturally modeled by the reinforcement learning framework [38]. Inspired by the success of reinforcement learning in solving sequential decision-making problems, we propose a multi-agent reinforcement learning (MARL) framework to select multiple discriminative frames or video clips from an untrimmed video to improve the recognition performance. The workflow of our system is illustrated in Figure 1.

Specifically, there are N agents in our framework and each of them is responsible for selecting one informative frame/clip from an untrimmed video. They initially sample N frames/clips evenly from the entire video and iteratively decide that each of their samples should come from temporally preceding or later location until encountering a STOP action at T_{stop} -th step or the maximum step number T_{max} is reached. To both combine nearby context and track previous status information for better decision making, we design a shared RNN-based *context-aware observation network* to model the local environment and its nearby context information as well as historical states to generate a status vector for each agent. Conditioned on the status vector, the *policy network* estimates the probability distribution over the predefined action space, according to which each agent takes action to adjust its sampling location. A carefully designed reward function is proposed and the MARL framework is optimized following REINFORCE [49] by maximizing the expected reward.

To verify the effectiveness of our proposed multi-agent reinforcement learning framework for frame sampling, extensive experiments are conducted on several popular

untrimmed video classification datasets, including ActivityNet v1.2 and v1.3 [2], YouTube Birds and YouTube Cars [55]. Results show that the proposed scheme achieves remarkable improvement over 2D/3D CNN baseline solutions which are equipped with different hand-crafted sampling strategies. In more detail, a new state-of-the-art on YouTube Birds and YouTube Cars is achieved and our single RGB model reaches a comparable performance of ActivityNet v1.3 champion submission with multi-modal multi-model fusion. To sum up, we make the following contributions.

- We focus on a previously overlooked point, i.e., the frame sampling strategy, in improving untrimmed video classification performance and intuitively formulate it as a Markov decision process.
- Multi-agent reinforcement learning is adopted to solve the formulated sequential decision problems. A novel framework that takes both context information and historical environment states into consideration for decision making is designed.
- The proposed method can be effectively applied to various existing untrimmed video recognition models to improve the performance, which is well witnessed by the excellent experimental results.

2. Related Work

2.1. Action Recognition

Our paper is closely related to works on deep-learning based action recognition, including end-to-end convolutional classification networks and two-stage recognition solutions. Karpathy *et al.* firstly introduces CNN for video classification in [24]. Then two-stream ConvNet [36] is proposed to merge the predicted scores from RGB and optical flow modalities, and the performance is improved by a large margin. ST-ResNet [10] further introduces residual connections between the two streams of [36] and shows great advantages in results. Currently, TSN [45] and C3D [41] are two well-known baseline methods for video recognition. The former is a 2D CNN based approach while the latter is based on 3D CNN. There are numerous follow-up studies to improve the aforementioned two baselines. For example, TLE [7], ShuttleNet [35], AttentionClusters [29] and NetVlad [1, 16] are proposed for better local feature integration instead of directly AVG-Pooling as used in TSN. OFF [37] and motion feature network [26] are proposed for integrating motion information modeling into a spatial CNN network, instead of using two streams. I3D [3] inflates deeper networks than C3D for spatial-temporal modeling. Given the heavy computation overhead of I3D, a series of works [52, 33, 5, 56, 42, 44, 20] are done to strike good effectiveness-efficiency trade-off. Meanwhile,

nonlocal network [47], compact generalized nonlocal network [54] and Nonlocal + GCN [48] bring great performance gain by leveraging extra network modules to capture feature-point-wise or ROI-wise spatial-temporal relations. The frame sampling strategies of all the aforementioned state-of-the-arts are hand-crafted. Instead, we propose to use the learning-based strategy to improve recognition performance for untrimmed videos.

2.2. Reinforcement Learning

Our work is largely inspired by those on leveraging reinforcement learning to solve vision problems. [31] is a classical work that uses RL for the spatial attention policy in image recognition, as well as in [28]. [27] proposes an aesthetic aware reinforcement learning framework to learn the policy for image cropping and achieves promising results. RL is also adopted for object detection in [32], for video object segmentation in [18] and object tracking in [17, 34].

As for semantic level video understanding, RL has also played an important role. For instance, it is utilized for action detection in [53], for action anticipation in [14], for natural language grounding in [19] and for video description in [46]. Our work is most closely related to the ones that use reinforcement learning for action recognition. In [4], the part activation policy of human body parts is learned with RL for action prediction. [40] focuses on skeleton-based action recognition and RL is adopted to distinguish discriminative poses. Following [53], the Fast-Forward algorithm is proposed in [9] to reduce the computation burden for untrimmed video classification. In this work, RL is utilized for both frames skipping planning and early stop decision making. In [51], the authors also focus on fast video prediction by adaptively selecting relevant frames on a per-video bias using reinforcement learning. Closely related work mainly focuses on improving prediction efficiency and a single agent is used for decision making. In contrast, we emphasize improving untrimmed video classification baselines and our model adopts multiple agents to cooperatively select multiple frames.

3. Approach

In this paper, we formulate frames selection as multiple sequential decision-making problems. Therefore, it naturally fits into the reinforcement learning framework. Figure 2 illustrates the multi-agent reinforcement learning framework of our proposed model. The model can be seen as multiple agents that interact with a video sequence of F frames/clips over time. Each agent picks a specific frame out of F frames and employs a *context-aware observation network* to encode the explored environment into a vector, h_t^a which is then fed into the following *policy network* to generate a proper action u_t^a from the action space \mathcal{A} . This action adjusts the frame that the agent to pick at the next

time step. N agents are identified by $a \in A \equiv \{1, 2, \dots, N\}$. All agents are initialized to be evenly distributed over the temporal space. When all agents decide to pick the current frame (to take the STOP action), a *classification network* emits the prediction based on the selected frames.

3.1. Architecture

Context-aware observation network The context-aware observation network is composed of a basic observation network f_o parameterized by θ_o , followed by a context network. The basic observation network is used to encode the video information at the frame/clip selected by the agent (namely, the local environment) into a feature vector o_t^a providing it as input to the context network. CNN based networks, both 2D CNN and 3D CNN are good at capturing features in frames or frame clips, and hence either of them can be adopted as the basic observation network. Unlike the single-agent system, for the multi-agent system, the action selected by each agent not only depends on its local environment state but is also impacted by its context information. Hence, we design a context-aware module f_h upon the basic observation network, parameterized by θ_h , to maintain a joint internal state of agents which summarizes history context information by a recurrent neural network. To make it work effectively, each agent only accesses context information from its $2M$ neighboring agents but not from all agents. More formally, at the time step t , the agent a observes a concatenated state $s_t^a = [o_t^{a-M}, \dots, o_t^a, o_t^{a+M}]$ and its previous hidden states h_{t-1}^a as inputs of the context module, then produces its current hidden states h_t^a :

$$h_t^a = f_h(h_{t-1}^a, s_t^a; \theta_h). \quad (1)$$

Especially, if M is set to 1, then the first agent loses the rear information while the last agent loses the front information. We use the beginning and ending frame/clip of the video to compensate for the information needed for the two cases respectively. For the multi-agent model, all of these agents share the context-aware observation network.

Policy network We adopt a fully connected layer followed by a softmax function as the policy network parameterized by θ_u . In every time-step t , each agent a selects an action $u_t^a \in U$ to execute, according to the probability distribution $\pi(u_t^a | h_t^a; \theta_u)$ generated by the policy network. The action space consists of three predefined actions, namely, *moving ahead*, *moving back* and *staying*. The moving stride is set to be δ frames/clips. When all agents choose *staying*, it means a STOP action is encountered. Note that we still make parameters of the policy network shared among all agents so that our model is readily applicable for testing at arbitrary N .

Classification network The classification network f_p is parameterized by θ_p . At each time step, these selected frames go through the classification network to produce its

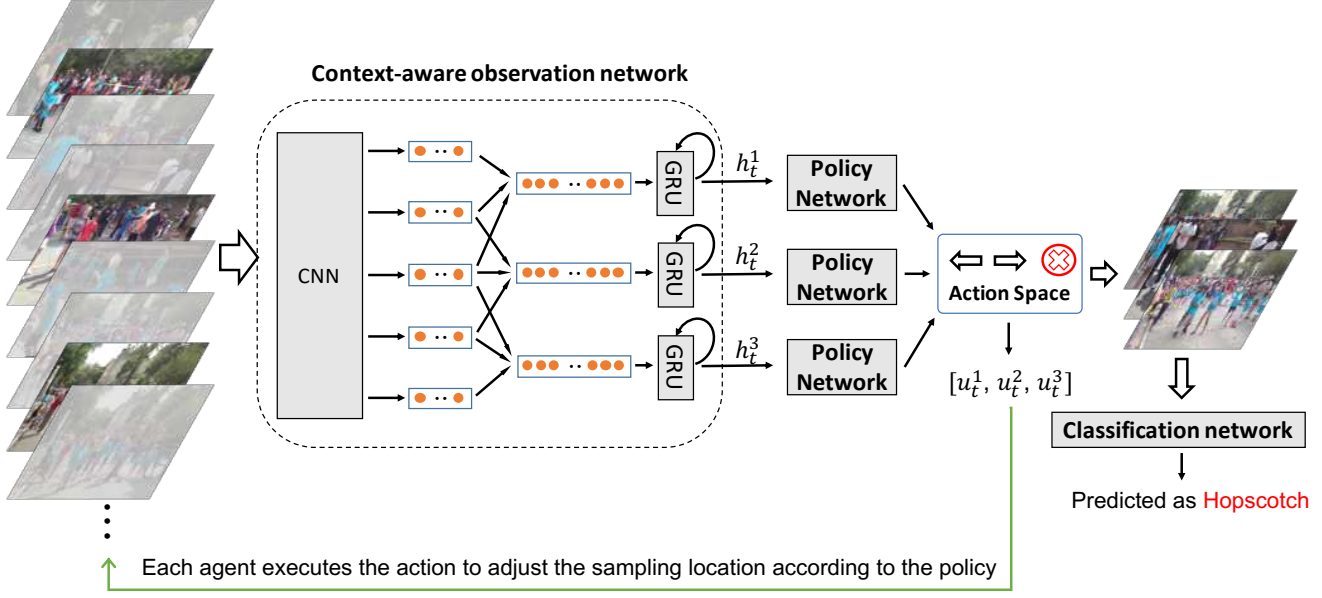


Figure 2. The proposed multi-agent reinforcement learning framework is composed of a context-aware observation network for capturing environment state h_t^n , a policy network for estimating probabilistic distribution over action space and a classification network for video-level prediction. There are N (which is 3 for illustration) agents and the environment is the F sampled frames/clips in our system. The agent interacts with the environment by taking action u_t^n to adjust sampling locations iteratively. The context-aware observation network is designed to allow context information communication among nearby agents. GRU, which integrates preceding and current states, is used to model the sequential decision making property of our system.

corresponding prediction logit $l_t^a \in \mathbb{R}^C$, where C is the number of classes. At the last time step, these prediction scores $l_{T_{stop}} = \{l_{T_{stop}}^1, l_{T_{stop}}^2, \dots, l_{T_{stop}}^N\}$ before Softmax normalization are aggregated with average pooling to yield the final video-level prediction. In our proposed method, the classification network can be easily replaced with any video classification module, such as Two Stream [36], TSN [45] and 3D CNNs [41]. For the simplicity in design, the classification network shares the parameters of layers before the last classifier layer with the observation network.

3.2. Objectives

The overall objective of the MARL-framework is to simultaneously maximize the expected reward of the frame sampling network and minimizing the classification loss. We use standard back-propagation to train the classification network parameterized by θ_p , and REINFORCE [49] to optimize the parameters of the basic observation network, the context module and the policy network $\theta_\pi = \{\theta_o, \theta_h, \theta_u\}$. Hence, our loss function consists of the MARL loss $\mathcal{L}_{MARL}(\theta_\pi)$ and the classification loss $\mathcal{L}_{Cls}(\theta_p)$.

3.2.1 MARL Objective

Reward function The reward function reflects how good the actions taken by the agents. When all the agents take

actions, each agent gets its own reward r_t^a based on the classification probability $p_t^a = \text{Softmax}(l_t^a)$ of the t -th time step. The reward is given to encourage the agent to find a more informative frame which can improve the probability of correct prediction step by step. Hence, we design a simple reward function that encourages the model to increase its confidence. Specifically, for the t -th ($t > 1$) time step, agent a receives a reward follows:

$$r_t^a = p_{t,gt}^a - p_{t-1,gt}^a, \quad (2)$$

where $p_{t,c}^a$ represents the probability of predicting the video as class c at the t -th time step for agent a , and gt is the ground-truth label of the video. All agents share the same reward function. Considering the sequential decision-making scenario, it is more appropriate to consider a cumulative discounted rewards, where rewards obtained in the more distant future contribute less to the current step. Specifically, at the time step t , the discounted return for agent a is

$$R_t^a = \sum_{i=0}^{T_{stop}-t} \gamma^i r_{i+t}^a, \quad (3)$$

where $\gamma \in (0, 1]$ is a constant discount factor that controls the importance of future rewards.

Policy gradient Given \mathcal{U} , a space of action sequences. Following REINFORCE [49], our objective can be ex-

pressed as

$$J(\theta_\pi) = \sum_{a=1}^N \sum_{u \in \mathcal{U}} \pi(u|s; \theta_\pi) R^a. \quad (4)$$

In our case we wish to learn network parameters θ_π that maximize the equation (4). The gradient of $J(\theta_\pi)$ is

$$\nabla_{\theta_\pi} J(\theta_\pi) = \sum_{a=1}^N \sum_{u \in \mathcal{U}} \pi(u|s; \theta_\pi) \nabla_{\theta_\pi} \log \pi(a|s; \theta_\pi) R^a. \quad (5)$$

This leads to a non-trivial optimization problem due to the high dimension of the action sequence space. REINFORCE addresses this by using Monte Carlo sampling to obtain K interaction sequences to approximate the policy gradients:

$$\nabla_{\theta_\pi} J(\theta_\pi) \approx \frac{1}{K} \sum_{k=1}^K \sum_{a=1}^N \sum_{t=0}^{T_{stop}} \nabla_{\theta_\pi} \log \pi(u_{t,k}^a | s_{t,k}^a; \theta_\pi) R_t^a. \quad (6)$$

Then we can use stochastic gradient descent to minimize the loss function:

$$\mathcal{L}_J(\theta_\pi) = -\frac{1}{K} \sum_{k=1}^K \sum_{a=1}^N \sum_{t=0}^{T_{stop}} \log \pi(u_{t,k}^a | s_{t,k}^a; \theta_\pi) R_t^a. \quad (7)$$

Maximum entropy To prevent policies from becoming deterministic too quickly, researchers use entropy regularization [50, 30]. Its success has sometimes been attributed to the fact that it “encourages exploration” [30]. The greater the entropy, the more ability of exploration an agent will have. Therefore, we follow the practice of using the entropy of policy to increase the ability to explore by:

$$\mathcal{L}_H(\theta_\pi) = -\sum_{a=1}^N \sum_{t=0}^{T_{stop}} \sum_{u \in \mathcal{A}} \pi(u_t^a | s_t^a; \theta_\pi) \log \pi(u_t^a | s_t^a; \theta_\pi). \quad (8)$$

Hence, the overall loss for MARL Objective is a combination of the two losses:

$$\mathcal{L}_{MARL}(\theta_\pi) = \mathcal{L}_J(\theta_\pi) + \lambda_1 \mathcal{L}_H(\theta_\pi), \quad (9)$$

where λ_1 is a constant scaling factor.

3.2.2 Classification Objective

A cross-entropy loss is applied to minimize the KL divergence between the ground truth distribution y and prediction p :

$$\mathcal{L}_{Cls}(\theta_p) = -\sum_{c=1}^C y_c \log p_c. \quad (10)$$

Finally, we minimize a hybrid loss that combines all the losses:

$$Loss = \mathcal{L}_{Cls}(\theta_p) + \lambda_2 \mathcal{L}_{MARL}(\theta_\pi), \quad (11)$$

where λ_2 is a constant scaling factor.

4. Experiments

4.1. Datasets and Evaluation Metrics

ActivityNet is a large-scale video benchmark for human activity understanding [2]. The first version of this dataset (termed as ActivityNet v1.2) has 100 classes of human activities. and its second version (termed as ActivityNet v1.3) contains 19,994 videos from 200 activity categories. Moreover, the ActivityNet dataset has temporal annotations of action instances for training data. It is also worth noting that the labels of the test set are not publicly available and thus the performances on ActivityNet dataset are all reported on the validation set. We decode videos at 1fps and use RGB frames in our experiments. Following the official evaluation script, the **evaluation metric** is based on mean average precision (mAP) for action recognition on ActivityNet.

YouTube Birds and YouTube Cars are two challenging video datasets for fine-grained video classification which consist of 200 different bird species and 196 different car models respectively [55]. We experimented with the RGB frames of the two datasets. Videos in YouTube Birds and YouTube Cars were downsampled to 2fps and 4fps respectively. We employ top-1 precision as the **evaluation metric** for the two datasets.

4.2. Implementation Details

One of the major differences in current video architectures is whether the convolutional operators use 2D (image-based) or 3D (video-based) kernels. Hence, we choose two successful video classification methods for feature extraction in our method, namely temporal segment network [45] and C3D [41]. The temporal segment network is equipped with segmental modeling (5 segments) to capture long-range temporal information and C3D uses 3D ConvNets to extract the temporal and spatial information of a video clip.

Training When using 2D ConvNets as our backbone, we first pre-train classification network f_p with the method introduced in [45]. For 3D ConvNets, we use the C3D [41] features provided by ActivityNet’s [2] website which are extracted every 8 frames with a temporal resolution of 16 frames. We use the pre-trained weights to initialize f_p , then jointly train the network with extra components. Adam with the initial learning rate of 0.0001 is adopted. The parameters used in the experiments are set as follows. We set F to 120, 100 and 100 for ActivityNet, YouTube Birds and YouTube Cars respectively. For the videos shorter than F frames, we cyclically repeat their frames to derive a video of F frames. In the policy network, we use a gated recurrent unit (GRU) cell with the hidden size of 1024 to model the sequential decision process and T_{max} is set to 10. γ is empirically set to 0.9 and λ_1, λ_2 are set to 1. We set N to 5 during training.

Testing Many state-of-the-art methods rely on some so-

Architecture	ActivityNet v1.2				ActivityNet v1.3			
	R25	U25	All	Ours	R25	U25	All	Ours
C3D	62.06	62.89	63.00	64.13	59.73	60.68	60.83	62.00
BN-Inception	78.76	80.02	80.50	81.99	75.08	76.48	77.33	78.32
ResNet-101	80.73	81.94	82.26	83.76	78.69	79.96	80.64	81.54
Inception-V3	81.90	82.66	83.25	85.01	79.27	80.33	80.86	82.34
ResNet-152	82.72	83.71	84.07	85.70	80.69	82.08	82.53	83.81

Table 1. Performance comparison of different ConvNet architectures on the ActivityNet dataset. For different architectures, randomly sampling 25 frames (R25), uniformly sampling 25 frames (U25), using all frames (All) and using our method to sample 25 frames (Ours) are evaluated. All architectures are based on ImageNet [6] pre-trained model, except C3D.

phisticated testing strategy or post-processing techniques, such as 10-crop testing, to boost performance. However, our framework does not demand these testing strategies. We simply sample 25 frames (N is set to 25) uniformly per video as the initial temporal location, and single-center cropping is applied to the selected frames from MARL to make the final prediction directly. During evaluation, we use maximum a posterior estimation to choose the action for each agent according to $\pi(u_t^a | s_t^a; \theta_\pi)$.

4.3. Improvements over 2D/3D CNN Baselines

MARL improves various backbones Here we examine our framework with several recent network architectures on the ActivityNet dataset using the RGB modality. Specifically, we equip MARL to five deep architectures: BN-Inception [23], Inception-V3 [39], ResNet-101 [21], ResNet-152 [21], and C3D [41]. For each architecture, we evaluate three hand-crafted frame sampling strategies as baselines: randomly sampling 25 frames (R25), uniformly sampling (U25) and using all frames (All). For random sampling, we evaluate three times and report the average results. Since the average duration of videos in ActivityNet is 117s and we decode videos at one fps, we set F to be 120, which indicates that using all frames means using 120 frames. For 2D ConvNets, we follow the TSN [45] framework and train the network on annotated action instances, all these 2D ConvNets are initialized with ImageNet pre-trained weights [6]. For C3D, we also carry out experiments with TSN-style training strategy by sampling 5 clips and predicting labels based on consensus. An average pooling on logits predicted from each sampled frame/clip is followed to produce video-level predictions for different sampling strategies in these experiments.

The results are summarized in Table 1. We observe that deeper models with higher accuracy (on ImageNet dataset [6]) result in better performance in the video classification task, and our method consistently obtains robust improvements over various 2D/3D architectures. For each architecture, our method achieves the best performance and improved over them by a large margin on both datasets comparing with randomly sampling, uniformly sampling and

Strategies	Instance	Video
R25/U25/All	80.69/82.08/82.53	80.17/81.23/81.73
Ours	83.81	82.98

Table 2. Performance comparison between different supervision on ActivityNet v1.3 val set.

using all frames.

We can see that even with a very powerful ResNet-152 backbone, MARL can largely boost recognition performance, specifically, the gain over U25 and All is 1.99% and 1.63% on ActivityNet v1.2, and as is 1.73% and 1.28% on ActivityNet v1.3. This verifies the effectiveness of our method, regardless of shallower or deeper baseline models.

Instance level v.s. video level supervision Annotations of action instances in untrimmed videos are expensive and hardly available under most circumstances. Experiments are carried out to show that our method consistently improves over the baseline method. Besides using these annotated activity instances when training the network, we train our model when only the video-level labels are used. In this experiment, the powerful ResNet-152 is utilized as our backbone network. We present the experimental results of our proposed MARL with annotated instances information (Instance) and video-level information (Video) in Table 2. From the results of ActivityNet v1.3 val set, we observe that our method achieves the best performance in comparison with the three hand-crafted baseline strategies, whether instance-level supervision is available or not. It is also worthy of noting that, our method with video-level information achieves even better performance than the model trained with instance-level supervision and tested with all frames. It further confirms that MARL is effective in picking discriminant frames in untrimmed videos.

4.4. Comparison with State-of-the-arts

Comparison with other frame selection methods We make comparisons under the scenario of observing the same number of total frames and results demonstrate our method outperforms other methods as shown in Figure 3. The re-




Figure 3. **Mean average precision vs. number of observed frames.** Comparisons of our method with AdaFrame [51], FrameGlimpse [53], Fast-Forward [9], and uniformly sampling.

Method	Backbone	Pre-trained	top-1	mAP
IDT [43]	-	ImageNet	64.70	68.69
C3D [33]	-	Sports1M	65.80	67.68
TSN [45]*	BN-Inception	ImageNet	72.97	76.56
P3D [33]	ResNet-152	ImageNet	75.12	78.86
RRA [55]	ResNet-152	ImageNet	78.81	83.42
Ours	ResNet-152	ImageNet	79.82	83.81
TSN [45]*	BN-Inception	Kinetics	78.98	81.80
Ours	BN-Inception	Kinetics	80.22	83.52
C16	Ensemble	-	-	90.9
Ours	SEResNeXt152	Kinetics	85.72	90.05

Table 3. Comparing with methods on the ActivityNet v1.3 validation dataset using RGB modality. * indicates the results of our implementation. C16 denotes the champion submission of the ActivityNet 2016 challenge, it fuses multiple powerful models and multi-modal (RGB, optical flow and audio) results.

sults of these methods are directly cited from Adaframe [51] whose authors reproduce FrameGlimpse [53] and Fast-Forward [9]. To be fair, we use the same preprocessing method, backbone network (*i.e.* ResNet-101) and training strategy of backbone as Adaframe [51]. Note that the performances of our solution in Figure 3 are lower than those reported in Table 1, this is because the backbone is not trained as what TSN [45] does for a fair comparison. In this experiment, the number of observed frames is calculated by averaging the total number of frames observed for each video at all steps of MARL.

Comparison with other SOTAs We also compare our method with other state-of-the-art methods on these challenging datasets. Table 3 shows results on ActivityNet

Method	YouTube Birds	YouTube Cars
BN-Inception*	60.13	61.96
I3D [3]*	40.68	40.92
TSN [45]*	72.361	74.340
RRA [55]*	73.205	77.625
U25/All/Ours	76.56/76.77/ 79.01	76.49/76.99/ 79.77

Table 4. Comparing with methods on YouTube Birds and YouTube Cars. * indicates the results of the method come from the latest project page of these datasets.

v1.3, where the results of state-of-the-art-methods all come from published paper or tech reports. We compare with several well known video recognition methods, which once achieved the state-of-the-art performance, including improved trajectories (iDT) [43], 3D convolutional networks (C3D) [41], temporal segment networks (TSN) [45], Pseudo-3d residual network (P3D) [33], and Redundancy Reduction Attention (RRA) [55]. We see that our method with ResNet-152 outperforms all these previous methods on the ActivityNet v1.3. Moreover, as shown in Table 3, it is beneficial to pre-train a model using Kinetics dataset[25] and then transfer it to the video classification task on ActivityNet dataset. “SEResNeXt152” pre-trained on Kinetics and finetuned on ActivityNet1.3 is the current SOTA model which won the champion of Action Proposal task in ActivityNet2018[15]. We conducted experiments with RGB modality based on “SEResNeXt152”, it achieves SOTA mAP of 87.95% (U25), 88.15% (All) and 90.05% (MARL). For readers’ reference, here we list the results of the 2016 champion submission as well as ours with SE-ResNeXt-152 [22] pre-trained on Kinetics dataset. To be noted, our results on ActivityNet v1.3 are obtained by only using 1fps RGB frames during the training phase and a single-center cropping strategy in validation. Even with this simple train and test setup, our single RGB model still achieves comparable results with the 2016 champion submission which fuses multiple powerful models and integrates RGB, optical flow and audio modalities.

Table 4 shows results on YouTube Birds and YouTube Cars. To make a fair comparison, the Inception-V3 [39] with the same architecture as RRA [55] is used as the backbone. Our method surpasses RRA on all these two datasets since categories in fine-grained tasks often share a similar appearance in general and hence require to focus more on the informative frames to distinguish from each other. Our best performance is 5.805% above that of other methods on the YouTube Birds and 2.145% on the YouTube Cars.

4.5. Ablation Study

Number of agents Our model is trained with N of 5 and we evaluate how the testing performance of ActivityNet 1.3 val set varies when N changes. Results can be found in




Figure 4. Visualization of the selected frame with different strategies on ActivityNet. The first row show frames from uniformly sampling, the second row depicts frames from our method without context-aware observation while the last row contains frames from our method.

	N=5	N=15	N=25	N=120	All
ResNet-152	80.19	82.99	83.81	83.72	82.53

Table 5. Impact of N on ActivityNet v1.3 val set in terms of mAP.

Table 5. It can be observed when N increases, the performance gradually improves at first. There is a very interesting phenomenon. When N reaches the total number of frames (which is 120), the performance slightly drops to 83.72% but is still better than feeding all frames directly (82.53%). It is because that some agents might select the same frame to avoid less informative frames, but not all the irrelevant frames are skipped by the agents. In this case, though performance drops, it is still better than 82.53%. Such results evident the variation of frame-level saliency in the untrimmed video and support our motivation.

Context range We explore the impact of parameter M in context-aware observation network by ablation study. We carry out experiments with MARL frameworks with various settings of M . The experimental results are listed in Table 6. Basically, our method still works for untrimmed videos when $M = 0$. It outperforms uniformly sampling by 0.91% and is even better than using all frames. The use of the context-aware observation network improves the non-context model with clear margins. However, no obvious gain is obtained by increasing M . Larger M means larger network size and cost, so we empirically set M to 1.

Policy network transferring We show that the learned policy network can still be effective when it is transferred directly. (1) With Resnet-152, performances of directly applying sampling networks trained on ActivityNet1.2/Youtube-Cars(Birds) to ActivityNet1.3 are shown in Table 7, see Cars’S, Birds’S and ANet1.2’S. (2) Our policy network trained for ResNet-101 classifier still works for ResNet-152 classifier on ActivityNet1.3, which is indicated as “cross-cls” in Table 7.

Qualitative results We also visualize some examples of

	U25	All	M=0	M=1	M=2	M=4
mAP	82.08	82.53	82.99	83.81	83.80	83.72

Table 6. Evaluation of context range on ActivityNet v1.3 val set using ResNet-152. U25 and All are hand-crafted strategies.

U25	All	MARL	Birds’S	Cars’S	ANet1.2’S	cross-cls
82.08	82.53	83.81	82.70	82.66	83.41	83.43

Table 7. The performance of different settings on ActivityNet v1.3 val set.

selected frames with different strategies on the validation data of ActivityNet in Figure 4. From top to bottom, frames picked by uniformly sampling, $M = 0$ and $M = 1$ are depicted. We see that our method is able to automatically select important frames according to surroundings and to avoid irrelevant frames.

5. Conclusion

In this paper, we presented a multi-agent reinforcement learning method for untrimmed video recognition, which can be effectively applied to existing video recognition frameworks to select the informative frames/clips from the untrimmed video. Experiments demonstrate that our method outperforms state-of-the-art baseline methods by a substantial margin, which verifies the effectiveness of our method. In the future, we plan to distillate our MARL based recurrent frame sampling network into a smaller feed-forward CNN to achieve more efficient untrimmed video classification.

Acknowledgment

The work was funded by National Natural Science Foundation of China (U1713203) and Shenzhen Science and Technology Innovation Commission (Project KQJSCX20180330170238897).

References

- [1] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pfister, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *CVPR*, 2016. 1, 2
- [2] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2, 5
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 1, 2, 7
- [4] Lei Chen, Jiwen Lu, Zhanjie Song, and Jie Zhou. Part-activated deep reinforcement learning for action prediction. In *ECCV*, 2018. 3
- [5] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. Multi-fiber networks for video recognition. In *ECCV*, 2018. 1, 2
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [7] Ali Diba, Vivek Sharma, and Luc Van Gool. Deep temporal linear encoding networks. In *CVPR*, 2017. 1, 2
- [8] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 1
- [9] Hehe Fan, Zhongwen Xu, Linchao Zhu, Chenggang Yan, Jianjun Ge, and Yi Yang. Watching a small portion could be as good as watching all: Towards efficient video classification. In *IJCAI*, 2018. 3, 7
- [10] Christoph Feichtenhofer, Axel Pinz, and Richard Wildes. Spatiotemporal residual networks for video action recognition. In *NIPS*, 2016. 1, 2
- [11] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Temporal residual networks for dynamic scene recognition. In *CVPR*, 2017. 1
- [12] Chuang Gan, Chen Sun, Lixin Duan, and Boqing Gong. Webly-supervised video recognition by mutually voting for relevant web images and web video frames. In *ECCV*, 2016. 1
- [13] Chuang Gan, Ting Yao, Kuiyuan Yang, Yi Yang, and Tao Mei. You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images. In *CVPR*, 2016. 1
- [14] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. Red: Reinforced encoder-decoder networks for action anticipation. *arXiv preprint arXiv:1707.04818*, 2017. 3
- [15] Bernard Ghanem, Juan Carlos Niebles, Cees Snoek, Fabian Caba Heilbron, Humam Alwassel, Victor Escorcia, Ranjay Khristna, Shyamal Buch, and Cuong Duc Dao. The activitynet large-scale activity recognition challenge 2018 summary. *arXiv preprint arXiv:1808.03766*, 2018. 7
- [16] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *CVPR*, 2017. 1, 2
- [17] Minghao Guo, Jiwen Lu, and Jie Zhou. Dual-agent deep reinforcement learning for deformable face tracking. In *ECCV*, 2018. 3
- [18] Junwei Han, Le Yang, Dingwen Zhang, Xiaojun Chang, and Xiaodan Liang. Reinforcement cutting-agent learning for video object segmentation. In *CVPR*, 2018. 3
- [19] Dongliang He, Xiang Zhao, Jizhou Huang, Fu Li, Xiao Liu, and Shilei Wen. Read, watch, and move: Reinforcement learning for temporally grounding natural language descriptions in videos. In *AAAI*, 2019. 3
- [20] Dongliang He, Zhichao Zhou, Chuang Gan, Fu Li, Xiao Liu, Yandong Li, Limin Wang, and Shilei Wen. Stnet: Local and global spatial-temporal modeling for action recognition. In *AAAI*, 2019. 1, 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [22] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 7
- [23] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 6
- [24] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 2
- [25] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 7
- [26] Myunggi Lee, Seungeui Lee, Sungjoon Son, Gyutae Park, and Nojun Kwak. Motion feature network: Fixed motion filter for action recognition. In *ECCV*, 2018. 2
- [27] Debang Li, Huikai Wu, Junge Zhang, and Kaiqi Huang. A2-RL: Aesthetics aware reinforcement learning for image cropping. In *CVPR*, 2018. 3
- [28] Xiao Liu, Jiang Wang, Shilei Wen, Errui Ding, and Yuanqing Lin. Localizing by describing: Attribute-guided attention localization for fine-grained recognition. In *AAAI*, 2017. 3
- [29] Xiang Long, Chuang Gan, Gerard de Melo, Jiajun Wu, Xiao Liu, and Shilei Wen. Attention clusters: Purely attention based local feature integration for video classification. In *CVPR*, 2018. 1, 2
- [30] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *ICML*, 2016. 5
- [31] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *NIPS*, 2014. 3
- [32] Aleksis Pirinen and Cristian Sminchisescu. Deep reinforcement learning of region proposal networks for object detection. In *CVPR*, 2018. 3
- [33] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, 2017. 1, 2, 7

- [34] Liangliang Ren, Jiwen Lu, Zifeng Wang, Qi Tian, and Jie Zhou. Collaborative deep reinforcement learning for multi-object tracking. In *ECCV*, 2018. 3
- [35] Yemin Shi, Yonghong Tian, Yaowei Wang, Wei Zeng, and Tiejun Huang. Learning long-term dependencies for action recognition with a biologically-inspired deep network. In *ICCV*, 2017. 1, 2
- [36] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014. 1, 2, 4
- [37] Shuyang Sun, Zhanghui Kuang, Lu Sheng, Wanli Ouyang, and Wei Zhang. Optical flow guided feature: a fast and robust motion representation for video action recognition. In *CVPR*, 2018. 2
- [38] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press, 2011. 2
- [39] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 6, 7
- [40] Yansong Tang, Yi Tian, Jiwen Lu, Peiyang Li, and Jie Zhou. Deep progressive reinforcement learning for skeleton-based action recognition. In *CVPR*, 2018. 3
- [41] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 1, 2, 4, 5, 6, 7
- [42] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 1, 2
- [43] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, 2013. 7
- [44] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification. In *CVPR*, 2018. 1, 2
- [45] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 1, 2, 4, 5, 6, 7
- [46] Xin Wang, Wenhua Chen, Jiawei Wu, Yuan-Fang Wang, and William Yang Wang. Video captioning via hierarchical reinforcement learning. In *CVPR*, 2018. 3
- [47] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 1, 2, 3
- [48] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018. 3
- [49] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 1992. 2, 4
- [50] Ronald J Williams and Jing Peng. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 1991. 5
- [51] Zuxuan Wu, Caiming Xiong, Chih-Yao Ma, Richard Socher, and Larry S Davis. Adaframe: Adaptive frame selection for fast video recognition. In *CVPR*, 2019. 3, 7
- [52] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018. 1, 2
- [53] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, 2016. 3, 7
- [54] Kaiyu Yue, Ming Sun, Yuchen Yuan, Feng Zhou, Errui Ding, and Fuxin Xu. Compact generalized non-local network. In *NIPS*, 2018. 3
- [55] Chen Zhu, Xiao Tan, Feng Zhou, Xiao Liu, Kaiyu Yue, Errui Ding, and Yi Ma. Fine-grained video categorization with redundancy reduction attention. In *ECCV*, 2018. 2, 5, 7
- [56] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. Eco: Efficient convolutional network for online video understanding. In *ECCV*, 2018. 1, 2