ActionSpotter: Deep Reinforcement Learning Framework for Temporal Action Spotting in Videos

Guillaume Vaudaux-Ruth
ONERA, Sorbonne Université
Palaiseau, France
guillaume.vaudaux-ruth@onera.fr

Adrien Chan-Hon-Tong
ONERA, Université Paris Saclay
Palaiseau, France
adrien.chan_hon_tong@onera.fr

Catherine Achard
Sorbonne Université, CNRS UMR 7222, ISIR
F-75005, Paris, France
catherine.achard@sorbonne-universite.fr

Abstract—Action spotting has recently been proposed as an alternative to action detection and key frame extraction. However, the current state-of-the-art method of action spotting requires an expensive ground truth composed of the search sequences employed by human annotators spotting actions - a critical limitation.

In this article, we propose to use a reinforcement learning algorithm to perform efficient action spotting using only the temporal segments from the action detection annotations, thus opening an interesting solution for video understanding.

Experiments performed on THUMOS14 and ActivityNet datasets show that the proposed method, named ActionSpotter, leads to good results and outperforms state-of-the-art detection outputs redrawn for this application. In particular, the spotting mean Average Precision on THUMOS14 is significantly improved from 59.7% to 65.6% while skipping 23% of video.

I. INTRODUCTION

Many works are interested in the analysis of actions in videos as it leads to a lot of applications. For example, it can be used to index Youtube videos [1] or to control robots by gestures [2]. The evaluation of the quality of actions [3] can be used to improve sports performances and the segmentation of video streams is well suited to monitor elderly people in their homes [4]. Action detection [5] is also an interesting subject since it aims at detecting action realizations in untrimmed videos. Thus, in addition to the action class, action detection returns the precise start and end times of each action instance. It can be used for video surveillance by detecting accidents, thefts or fights, but also for people protection (e.g. detection of falls).

The visual summary of videos has also been widely studied in particular to provide a smart scroll bar when streaming videos. The two most common frameworks are key frame selection [6], [7], [8], [9] and key sub-shot selection [10], [11]. Both frameworks are classically unsupervised and aim at finding clusters that best describe all frames.

However, to summarize the actions in videos, these techniques are not optimal. Indeed, it can be very difficult to find precise action boundaries, as it is done in action detection. Furthermore these boundaries are not necessary to know the existence of an action. Moreover, finding a visual summarizing does not necessarily reflect the number of actions. In fact, the visual summarizing is intended to extract salient or key images from a video stream. Thus, an action composed of several

Fig. 1: **Overview of the proposed spot frame extraction framework.** It extracts one spot frame per action occurrence (in green in the figure) by sparsely browsing the video.

visual parts will be described using several key frames, which does not allow to count the instances of actions. This issue also prohibits reliance on single-image action classification [12], [13] for such task, as this could lead to discontinuous segments and subsequent overestimation of action occurrences.

To overcome these limitations, action spotting has recently been proposed in [14]: it consists in producing an ordered list of the action instances contained in the video, or in other words “the process of finding any temporal occurrence of an action” as defined in [14]. Typically, it can be done by selecting one frame per action instance, as shown in figure 1.

To the best of our knowledge, the spotting task has only recently appeared in the literature: Alwassel *et al.* [14] use it as pre-processing for detection, Bhardwaj *et al.* [15] or Wu *et al.* [16] use it to produce accurate classification.

In this work, we consider action spotting as an individual task (and not to speed up classification or improve detection). The resulting spot frames are not expected to be key frames, and are not annotated in the ground truth (only composed of action segments): they are automatically selected by the method to optimize the spotting results. Thus, for a given action segment, the selection of any frame belonging to this time interval as a spot frame is correct. Consequently, two spotting outputs can be different but equally accurate. The prediction is considered perfect if and only if there is exactly one spot frame per ground truth segment, with a correct label

(allowing for example a perfect count). Thus, the problem is asymmetric: while the ground truth is composed of a set of segments, the prediction is a set of spot frames. We therefore introduce a metric to directly measure the quality of a set of spot frames. It is based on the mean Average Precision (mAP) value used in the AVA dataset toolbox [1] adapted to action spotting. The above-mentioned measure, which estimates spotting performance, is available in supplementary material and is described in more detail in section III-A. In addition to defining this new metric, we have also redrawn the predictions of state-of-the-art action detectors to fit the proposed metric, as will be detailed in section IV-C.

Independently, in order to carry out fully supervised action spotting, [14] introduced expensive annotations on the behavior of human annotators when performing action spotting. We believe that supervised learning is not straightforward as the problem is asymmetric, and we propose the use of reinforcement learning to produce spot frames from action segments only. This algorithm, called ActionSpotter, can handle videos containing multiple classes of action with multiple occurrences and is described in section III-B.

To summarize, our contributions are as follows:

- We release a metric for the spotting task (see attachment).
- We propose to use deep reinforcement learning algorithm to avoid the time-consuming annotations of human behavior and we propose to spot frames in a semi-supervised way (only action segments are available during training).
- We propose ActionSpotter, a reinforcement learning architecture extracting spot frames while observing as few frames as possible. It is based on the state-of-the-art actor-critic architecture.
- We show that this architecture is more relevant for action spotting than a post-processing of state-of-the-art detectors and/or spotting baselines.

Our framework is presented in Section III after related works in Section II. Then, experiments on THUMOS14 and ActivityNet are presented in Section IV before conclusion.

II. RELATED WORKS

Algorithms based on deep reinforcement learning have been proposed in literature to quickly browse videos, especially for fast video classification [16], [17], early video detection [18], [19] or action detection [5], [14].

For example, [16] offers an impressive result producing a state-of-the-art classifier by exploring only (on average) 8 frames per video. However, this is only possible thanks to the prior that videos contain a single action (or even object): in this context, a single good image could be enough to decide the class of the whole video.

In detection or spotting, as this assumption does not hold, performance decreases quickly when frames are skipped. This is particularly evident for short actions in long videos that can be completely missed by a greedy browsing. Typically, Giancola *et al.* [20] introduced a benchmark for locating very sparse events in soccer videos and propose a method based on a classification of video slices. However, as actions are very

short, they process all the frames. Only [5] tackles detection using few images.

In this paper, we rely on reinforcement learning, but for another reason: spot images are not provided by the ground truth, which is only composed of temporal segments. Thus, even though [16] also uses an architecture based on reinforcement learning, the goal of our work is different, our objective being to produce a good spotting. Such an idea can be found for object spotting [21] in 2D images, but we believe we are the first to apply it to temporal action spotting. Our framework also does not require video-retargeting [22], [23].

Moreover, our framework balances the trade-off between accuracy and frames skipping and, in opposition to [5], [16], [17], [14], ActionSpotter processes videos online. Indeed, we show that using the current frame and a memory of previous ones is sufficient to take correct decisions.

These considerations lead to a practical spotting algorithm that does not rely on detection or segmentation. This property is interesting as detectors have to tackle a much more difficult task as they have to find the starting and ending points of actions that are often ambiguously defined.

The article closest to our work is [14] where Alwassel *et al.* propose an algorithm that produces one spot per video. Their process is fully supervised and is done by learning the trajectories made by a human during the exploration of the videos. The authors argue that the use of supervised trajectories during training is much more direct and simple than reinforcement learning. However, this is an important limitation of this method as it requires a lot of human acquisition to obtain the browsing strategies of the videos. Moreover, [14] does not consider the spot frame extraction for itself but only as a first step for action detection and thus does not present spotting results.

III. METHOD

A. Action spotting and proposed evaluation metric

In this section, we present the proposed methodology for action spotting. As presented in introduction, the goal is to browse a video in order to select *spot frames* summarizing human activity in videos (see Figure 1). We are thus interested in optimizing the quality of spot frames (described below), but also the proportion of skipped frames, called the *skip ratio*.

Spot Frames: If C is a predefined set of action classes, and $V = \{v_t\}_{t=1}^T$ a video sequence of T frames (or frame chunks), containing a set of K temporally ordered actions proposals of interest $Q = \{q_k \mid q_k = (a_k, [I_k, J_k])\}_{k=1}^K$, with $a_k \in C$ the action class, and with I_k and J_k the start and end time of a proposal, then our goal is to produce a set of spot frames with likelihoods and labels $\mathcal{V} = \{(\tau_k, l_k, \alpha_k)\}_{k=1}^K$ such that $\forall k \in \{1, \dots, K\}$, $\tau_k \in [I_k, J_k]$ and $\alpha_k = a_k$.

Evaluation Metric: In order to propose an unbiased metric reflecting the quality of the extracted spot frames, we build it on the basis of the state-of-the-art metric used in object detection [24] or temporal action localization [25]: the mean Average Precision (mAP). This metric is for example used on Pascal, THUMOS or AVA challenge. Thus, we adapt this

Browser Agent: In parallel, this memory h_n is also forwarded to the Browser Agent $BROW$ which decides the next frame to visit, i.e. $\tau_{n+1} = \tau_n + BROW(h_n)$.

Skip ratio: Importantly, $\tau_n \neq t$ as our pipeline does not go through all the images: let N such that $\tau_N = T$ (i.e. the number of steps to process the video), then, the ratio of skipped frames is defined as $1 - \frac{N}{T}$.

Global dynamic: Then, at step n , ActionSpotter (AS), has the following dynamic:

$$AS(v_{\tau(n)}, h_{n-1}) : \begin{cases} f_n = BB(v_{\tau(n)}) \\ h_n = GRU(f_n, h_{n-1}) \\ l_n = SF(h_n) \\ p_n = CL(h_n) \\ \alpha_n = \arg \max_c (p_{n,c}) \\ \mathcal{V}_{n+1} = \mathcal{V}_n \cup \{(\tau_n, l_n, \alpha_n)\} \\ \tau_{n+1} = \tau_n + BROW(h_n) \end{cases} \quad (1)$$

with $h_{n-1} = GRU(BB(\{v_{\tau(i)}\}_{i=1}^{n-1}))$, the memory of the past viewed frames (or frame chunks).

Figure 2 gives an illustration of the overall framework.

C. Training and objectives

Knowing that we use a pretrained model for BB (see section IV), the training goal is to optimize the parameters of GRU , $BROW$, SF and CL so that by processing the video V , the policy $\pi = BROW + SF$ provides an accurate set of spot frames $\mathcal{V}$, while skipping a large number of frames (i.e. with $BROW(h_n)$ as high as possible). It is important to note that the browser $BROW$ and spot frame selection network SF are not learned in a supervised way, removing the need for specific annotations. Instead, a state-of-the-art reinforcement learning algorithm is used with a specific reward function quantifying the relevance of each step under policy π (both for accuracy and browsing ratio).

Reward. As we use mAP to evaluate model performance, we choose to introduce a reward directly linked with this metric: the global policy has to maximize the final mAP of the video being processed (plus an entropy term $\rho \mathcal{H}(\pi(n))$ which will be explained later). Moreover, we can easily find a trade-off between the efficiency of video browsing and the accuracy of spot frame selection by discounting the final mAP by γ^N where $\gamma \in (0, 1]$ is the discount factor and N the explored frames.

According to reward shaping theory [27], this final reward can be hybridized with a potential-based shaping to help reinforcement convergence without changing optimal policy. Straightforwardly, using the mAP after each step is an interesting shaping potential. Thus, our local reward $r_{\pi,n}$ at step n is just the difference of mAP between step n and $n-1$ under policy π (plus entropy term):

$$r_{\pi,n} = \gamma \text{mAP}(\mathcal{V}_n) - \text{mAP}(\mathcal{V}_{n-1}) + \rho \mathcal{H}(\pi(n)) \quad (2)$$

Then, $R_{\pi,n}$ the cumulative discounted rewards is:

$$R_{\pi,n} = \sum_{k=1}^n \gamma^k r_{\pi,k} \quad (3)$$

By omitting entropy, we can directly verify that:

$$\begin{aligned} R_{\pi,N} &= \sum_{k=1}^N \gamma^k r_{\pi,k} = \sum_{k=1}^N \gamma^k \gamma \text{mAP}(\mathcal{V}_k) - \gamma^k \text{mAP}(\mathcal{V}_{k-1}) \\ &= \gamma^{N+1} \text{mAP}(\mathcal{V}_N) - \text{mAP}(\mathcal{V}_0) \end{aligned}$$

exactly as desired as $\text{mAP}(\mathcal{V}_0) = 0$. This invariance in shaping is true even with entropy as shown by [27].

Actor-Critic optimization. The Policy Gradient method is based on the total expected reward, and therefore requires a long sequence of actions to update the policy. Good and bad actions are then averaged, which can introduce convergence issues. Actor-Critic approach [28] is known to be a way to avoid this issue by evaluating each action independently. It uses two models named actor and critic.

The actor is straightforwardly trained to find the policy that maximizes the expected return:

$$J_{actor} = \mathbb{E}[R_{\pi,N}] \quad (4)$$

The critic measures how good the policy is (value-based) and produces an estimation of the value function which is the expected discounted reward $\text{crit}(h_n) \approx \mathbb{E}[R_{\pi,N} - R_{\pi,n} | h_n]$. Thus, the loss function linked to the critic is defined as:

$$\mathcal{L}_{critic} = \frac{1}{2} \|\text{crit}(h_n) - \mathbb{E}[R_{\pi,N} - R_{\pi,n} | h_n]\|_2 \quad (5)$$

In reinforcement learning, it is crucial to balance exploration and exploitation. Following the work of Haarnoja *et al.* [28], it is possible to integrate this balance directly into the reward, by adding an entropy penalty which forces the actor to uniformly explore states with equal rewards. Thus, a penalty $\rho \mathcal{H}(\pi(n))$ is added in equation 2. ρ is the temperature parameter which balances exploration and exploitation and is automatically adjusted following [28]. $\mathcal{H}()$ is the entropy function. Let us denote that without shaping or temperature, reinforcement has difficulty to converge.

In our algorithm, the actor is the combination of the Browser Agent $BROW$ and the Spot Frame Selector Agent SF . Thus, $\mathcal{H}(\pi(n))$ is the entropy applied to the distribution $\pi(n)$ over choices to update the current state i.e. $BROW(h_n)$ and $SF(h_n)$.

On the other side, the classification network CL is trained, in a supervised way, using Cross-Entropy (CE). Thus,

$$L_{cls} = CE(p_n, a_{\tau_n}) \quad (6)$$

Final loss: Combining previous losses, our final objective is to minimize:

$$\mathcal{L}_{global} = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{critic} - \lambda_2 J_{actor} \quad (7)$$

As J_{actor} objective is non differentiable we use REINFORCE [29] to derive the expected gradient:

$$\nabla J_{actor} = \nabla \mathbb{E} \left[\sum_{n=1}^N \log(\pi(n)) (R_{\pi,n} - \mathbb{E}[R_{\pi,n} | h_n]) \right] \quad (8)$$

We can then approximate this equation by using Monte Carlo sampling and finally use stochastic gradient descent to minimize our final objective.

THUMOS'14						
Approach	Detection mAP@					Spotting mAP
	0.1	0.2	0.3	0.4	0.5	
Glimpses [5]	48.9	44.0	36.0	26.4	17.1	-
SMS [30]	51.0	45.2	36.5	27.8	17.8	-
M-CNN [31]	47.7	43.5	36.3	28.7	19.0	41.2
CDC [32]	-	-	41.3	30.7	24.7	31.5
TURN [33]	54.0	50.9	44.1	34.9	25.6	44.8
R-C3D [34]	54.5	51.5	44.8	35.6	28.9	52.2
SSN [35]	66.0	59.4	51.9	41.0	29.8	-
A-Search [14]	-	-	51.8	42.4	30.8	-
CBR [36]	60.1	56.7	50.1	41.3	31.0	50.1
BSN + UNet [37]	-	-	53.5	45.0	36.9	-
Re-thinking F-RCNN [38]	59.8	57.1	53.2	48.5	42.5	-
D-SSAD [39]	-	-	60.2	54.1	44.2	59.7
Ours (TSN backbone)	-	-	-	-	-	62.4
Ours (I3D backbone)	-	-	-	-	-	65.6

TABLE I: **Results on THUMOS14 validation set.** Second column: state-of-the-art detector results according to detection metric, computed with different IOU thresholds ranking from 0.1 to 0.5. Last column: mAP results for the spotting task.

IV. EXPERIMENTS

Currently, we are the first to tackle action spotting as an individual task. It is therefore not easy to position our algorithm regarding the start-of-the-art. So, we offer two experiments to highlight the capabilities of this algorithm. First, we show that ActionSpotter produces better sets of spot frames than post-processed state-of-the-art detectors. Then, we compare ActionSpotter to several baselines to highlight that action spotting is a challenging task and that reinforcement learning is the relevant way to deal with this asymmetrical problem where the ideal output is not defined in the ground truth. Finally, some additional experiments reveal that ActionSpotter is able to balance accuracy and skip ratio, simply by modifying the discount factor associated with reinforcement learning.

A. Datasets

We evaluate our approach on the well-known THUMOS14 [40] and ActivityNet [41] datasets.

THUMOS14 dataset is composed of 101 activity classes for action recognition and a subset of 20 classes for action detection. Validation and testing sets contain respectively 200 and 212 untrimmed videos temporally annotated with action instances. We adopt the classical train/test setting of THUMOS14 protocol: training is done on 20 classes validation set and evaluation is done on testing set - original training set being not suited for detection.

ActivityNetv1.2 dataset contains 9,682 videos in 100 classes collected from YouTube. The dataset is divided into three subsets: 4,819 videos for training, 2,383 for validation and 2,480 for testing. Action spotting results on ActivityNet dataset are reported on validation set as evaluation server does not compute our spotting metric.

B. Implementation details

As previously mentioned, any type of backbone network can be used to encode local information from images. In order to have the same backbones as the state-of-the-art action

detectors, we rely on classical TSN [42] and I3D [43] feature extractors. This setting allows good reproducibility as such features are provided by [44] and [37]. These techniques operates on both RGB frames and optical-flow field to capture appearance feature and motion information. TSN operates on individual video frames while I3D features are extracted from non-overlapping 16-frame video slices. For this second technique, as feature represents non-overlapping frame slices, the Browser Agent BROW does not process individual frames but frame slices. This does not change the skip ratio because it is equivalent to considering one slice in K instead of one frame in K .

For the memory, we use a one-layer GRU with respectively 2,048 and 400 hidden units for THUMOS14 and ActivityNet (the size of the respected features). BROW, SF, CL and *crit* have 3 linear layers and ReLu activation function. The first two layers have the same number of hidden units as the GRU layer and the last one has the size of the network output.

In our main setting, BROW can choose to move to the next frame, skip one frame or skip three frames. At training time, in order to approximate Eq.8, the actions performed by BROW and SF are sampled from a categorical distribution parameterized by their respective logits. At testing time, the action performed is the one with the highest likelihood for BROW (and for CL). SF directly outputs a likelihood.

We use PyTorch for implementation and Adam for optimization with an initial learning rate of 10^{-4} and a batch size of 32. Convergence is much faster when *CL* and *SF* are trained alone for few epochs before starting the whole reinforcement learning process with $\lambda_1 = \lambda_2 = 1$.

C. Comparison between ActionSpotter and detectors

We evaluate ActionSpotter performances on **THUMOS14** and **ActivityNet** datasets using our proposed spotting metric. Results are presented in Table I and II. These tables also present results obtained by state-of-the-art detectors for the spotting task. We use published results or available codes to obtain detection results, so there is no re-implementation

ActivityNet v1.2					
Approach	Detection mAP@				Spotting mAP
	0.5	0.75	0.95	Avg	
W-TALC [45]	37.0	14.6	-	18.0	-
SSN-SW [35]	-	-	-	18.1	-
3C-Net [46]	37.2	23.7	9.2	21.7	-
FPTADC [47]	37.6	21.8	2.4	21.9	-
SSN-TAG [35]	39.2	25.3	5.4	25.9	55.4
BSN [48]	46.5	30.0	8.0	30.0	49.6
BMN [49]	50.1	34.8	8.3	33.85	55.3
Ours (TSN backbone)	-	-	-	-	58.1
Ours (I3D backbone)	-	-	-	-	60.2

TABLE II: **Results on ActivityNet v1.2 validation set.** The column AVG indicates the average mAP for IoU thresholds: 0.5:0.05:0.95.

issues. Detection results are redrawn into spotting results by extracting the centers of the predicted segments.

These experiments show that ActionSpotter significantly outperforms state-of-the-art detectors for the spotting task: the mAP of the latest action detector D-SSAD [50] is improved from 59% to 65% on THUMOS14 database.

It can be pointed out that FrameGlimpses [5] offers very low performance but has a skip ratio of 98%. Currently, ActionSpotter performance also decreases from 62.4% to 50.9% when using TSN backbone and increasing the skip ratio to 98%. This shows that it is difficult to do accurate spotting (and even more detection) with only 2% of the frames, which is not the case for action classification [16].

One may wonder why spotting performance are not similar to detection performance with a low Intersection over Union (IoU) value. In fact, this is due to a difference in the matching mechanism. In the case of detection, predicted segments match with ground-truth segments according to the best IoU while, in the case of spotting, the predicted spots match the ground-truth according to their scores (which make sense since IoU no longer exists). Thus, for detection task, it is better to predict segments with good localization and random confidence than to produce one segment per action with good confidence but coarse localization. Conversely, only the confidence score matters in spotting. This difference between the matching processes induces some changes in the ranking of predictions between detection and spotting: while CDC [32] is more efficient in detection than M-CNN[31], it is the opposite for spotting.

Thus, spotting results obtained with detectors cannot be easily compared with those of spotting algorithms as their primary goal is not the same. However, the large performance gap between our method and detection methods shows that it is not sufficient to post-process detector results to have an optimal spotting, and that specific algorithms such as ActionSpotter are required.

D. Ablation study and comparison with baselines

Previous experiments show that action spotting requires a specific approach. Moreover, action spotting is an asymmetric

problem since any frame of a ground truth segment can be used as a spot frame and, therefore, optimal spot frames are not defined by these segments. Based on this, Reinforcement Learning appears to be suited to tackle this problem and we show: (i) comparative results between supervised and reinforced spotting algorithm, (ii) an ablated version of ActionSpotter. Results are presented in table III.

Method	mAP (%)
Naive Segmentation	32
Multi-Task Segmentation	43
Supervised ActionSpotter	52
No Memory ActionSpotter	45
ActionSpotter (memory + reinforcement)	65

TABLE III: **Comparison between ActionSpotter and other spotting algorithms on THUMOS14 validation set.**

Naive segmentation is a simple semantic CNN-based segmentation (based on I3D + 2 layers) where the ground truth is constructed as follows: the center of each action segment is labeled by its action class and all other frames are labeled as background. This spotting-oriented baseline, only reaches 32% of mAP.

A two-tasks learning process is then used as a multi-task segmentation: the first task learns the center of segments and the second predicts the action class. This second baseline deals with a balanced problem helping gradient stabilization and leads to 43% of mAP.

Then, we train the ActionSpotter architecture to the supervised segmentation of action centers (instead of reinforcement learning). In this setting, performance reaches 52.3% mAP, which is higher than other baselines but well below ActionSpotter.

Finally, we consider ActionSpotter (with reinforcement learning) by removing its memory. Performance drops to 45%: as expected, without memory, reinforcement and supervision are practically the same since the decision is frame based.

These results highlight the fact that it is not easy to supervise action spotting. Actually, it is preferable to let the network select the easiest spot frame rather than imposing them. That's exactly what reinforcement does as the reward is only based on the final output.

Figure 3 shows qualitative results on three videos: in most cases, the spot frames (third row in red) are not extracted in the middle of the actions (first row in green). ActionSpotter adapts itself to the difficulty of the processed videos, as shown by the fact that the browser explores less frames (second row in blue) in areas without actions. These results are presented in video in supplementary material.

In addition, ActionSpotter allows direct optimization of the mAP score using state-of-the-art shaping technique. Indeed, this shaping technique is an important component as performance drops to 47% without it. ActionSpotter is therefore based on 3 key ideas: end-to-end learning, reinforcement learning to tackle asymmetry and shaping to allow better convergence.

Fig. 3: **Example of ActionSpotter outputs.** The figure displays outputs of ActionSpotter on THUMOS14 validation videos: columns represent the frames and, the 3 rows represent respectively ground truth segments, explored frames and selected spot frames. The same results are presented in video in supplementary material.

Our	γ				
	1.0	0.99	0.98	0.96	0.95
mAP (%)	65.6	64.3	63.4	61.4	61.3
skip ratio (%)	23	29	53	51	51

TABLE IV: **Performance on THUMOS14 validation set according to the discount factor γ .**

E. Trade-off between accuracy and skip ratio

ActionSpotter manages the lack of spot frame ground truth and tackles the spotting problem by optimizing the mAP. Moreover, it allows to balance the mAP and the browser skip ratio. We report in Table IV the mAP and the skip ratio for different values of γ to highlight the impact of the discount factor on the policy learning.

As expected, for small values, the lower the γ discount factor, the higher the skip ratio (convergence issues appear quickly when the γ decreases). But, as a trade-off exists between speed and accuracy, accuracy decreases too. Nevertheless, when the skip ratio increases from 23% to 53%, the mAP decreases only by 2%. For $\gamma = 0.98$, our algorithm is still better than the best detector for spotting while using only 47% of the frames. Currently, with the presented action space setting that puts a lot of emphasis on mAP, it is difficult to go below a 53% skip ratio. But, by adding more actions, it is possible to skip more frames even if mAP decreases significantly. For example, a mAP of 50.9% is obtained using only 2% of the frames (same skip ratio as [5]). But there is no point in skipping so many frames if it leads to such low performance level. As mentioned before, it seems difficult to perform accurate detection or spotting with 2% of frames, unlike classification [16] which assumes only one action.

Thus, while skipping many frames degrades performance, skipping some frames improves it. Indeed, by removing the browser and processing all video frames, the mAP decreases from 65.6% to 64.0%. Using the browser agent leads to a 1.6% improvement in mAP while using 23% fewer frames.

It is interesting to note that we also train our pipeline without the Browser but with uniform subsampling. Precisely, we subsample videos with different sampling rates and observe the mAP. Figure 4 shows that ActionSpotter produces better results regardless of the skip ratio.

Fig. 4: **Benefit of reinforcement browsing:** mAP according to the percentage of viewed frames for ActionSpotter algorithm and uniform sub-sampling.

V. CONCLUSION

We propose an algorithm called ActionSpotter able to tackle action spotting: collecting one frame per action instance. Our algorithm is based on state-of-the-art reinforcement algorithm and is able to tackle action spotting in a streaming context (frames are not stored) while skipping some video frames.

To evaluate the proposed algorithm, we introduce a metric that quantifies the accuracy of the action spotting. It is based on the mean Average Precision (mAP) conventionally used in detection. The main result of this paper is that the adaptation of state-of-the-art action detectors into spotting algorithms is less efficient than learning an end-to-end spotting: ActionSpotter reaches 65% of mAP (while skipping 23% of video frames), when state-of-the-art detectors only reach 59% of mAP (with a dense exploration).

Indeed, unlike action detectors that have to deal with ambiguous temporal boundaries, ActionSpotter focuses on extracting one frame per instance thanks to reinforcement learning.

ACKNOWLEDGMENTS

This work was performed using HPC resources from GENCI-IDRIS (2019-AD011011269)

REFERENCES

- [1] C. Gu, C. Sun, D. A. Ross, C. Vondrick, C. Pantofaru, Y. Li, S. Vijayanarasimhan, G. Toderici, S. Ricco, R. Sukthankar *et al.*, "Ava: A video dataset of spatio-temporally localized atomic visual actions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [2] S. S. Rautaray and A. Agrawal, "Vision based hand gesture recognition for human computer interaction: a survey," *Artificial intelligence review*, vol. 43, no. 1, 2015.
- [3] M. Morel, C. Achard, R. Kulpa, and S. Dubuisson, "Automatic evaluation of sports motion: A generic computation of spatial and temporal errors," *Image and Vision Computing*, vol. 64, 2017.
- [4] A. Chan-Hon-Tong, C. Achard, and L. Lucat, "Simultaneous segmentation and classification of human actions in video streams using deeply optimized hough transform," *Pattern Recognition*, vol. 47, no. 12, 2014.
- [5] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei, "End-to-end learning of action detection from frame glimpses in videos," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [6] T. Liu, H.-J. Zhang, and F. Qi, "A novel video key-frame-extraction algorithm based on perceived motion energy model," *IEEE transactions on circuits and systems for video technology*, 2003.
- [7] S. K. Kuanar, R. Panda, and A. S. Chowdhury, "Video key frame extraction through dynamic delaunay clustering with a structural constraint," *Journal of Visual Communication and Image Representation*, 2013.
- [8] N. Ejaz, I. Mehmood, and S. W. Baik, "Efficient visual attention based framework for extracting key frames from videos," *Signal Processing: Image Communication*, vol. 28, no. 1, 2013.
- [9] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman, "Video summarization with long short-term memory," in *European conference on computer vision*. Springer, 2016.
- [10] I. Mademlis, A. Tefas, N. Nikolaidis, and I. Pitas, "Movie shot selection preserving narrative properties," in *2016 IEEE 18th International Workshop on Multimedia Signal Processing (MMSP)*. IEEE, 2016.
- [11] I. Mademlis, A. Tefas, and I. Pitas, "A salient dictionary learning framework for activity video summarization via key-frame extraction," *Information Sciences*, vol. 432, 2018.
- [12] Z. Zhao, H. Ma, and S. You, "Single image action recognition using semantic body part actions," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [13] G. Guo and A. Lai, "A survey on still image based human action recognition," *Pattern Recognition*, vol. 47, no. 10, 2014.
- [14] H. Alwassel, F. Caba Heilbron, and B. Ghanem, "Action search: Spotting actions in videos and its application to temporal action localization," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [15] S. Bhardwaj, M. Srinivasan, and M. M. Khapra, "Efficient video classification using fewer frames," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [16] Z. Wu, C. Xiong, C.-Y. Ma, R. Socher, and L. S. Davis, "Adaframe: Adaptive frame selection for fast video recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [17] W. Wu, D. He, X. Tan, S. Chen, and S. Wen, "Multi-agent reinforcement learning based frame sampling for effective untrimmed video recognition," *arXiv preprint arXiv:1907.13369*, 2019.
- [18] S. Ma, L. Sigal, and S. Sclaroff, "Learning activity progression in lstms for activity detection and early detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [19] M. Gao, M. Xu, L. S. Davis, R. Socher, and C. Xiong, "Startnet: Online detection of action start in untrimmed videos," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- [20] S. Giancola, M. Amine, T. Dghaily, and B. Ghanem, "Soccernet: A scalable dataset for action spotting in soccer videos," *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018.
- [21] H. Perreault, G.-A. Bilodeau, N. Saunier, and M. Hritier, "Spotnet: Self-attention multi-task network for object detection," *arXiv*, 2020.
- [22] A. Bansal, S. Ma, D. Ramanan, and Y. Sheikh, "Recycle-gan: Unsupervised video retargeting," in *Proceedings of the European conference on computer vision (ECCV)*, 2018.
- [23] L. Wolf, M. Guttmann, and D. Cohen-Or, "Non-homogeneous content-driven video-retargeting," in *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 2007.
- [24] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International Journal of Computer Vision*, vol. 88, no. 2, Jun. 2010.
- [25] B. G. Fabian Caba Heilbron, Victor Escorcia and J. C. Niebles, "Activitynet: A large-scale video benchmark for human activity understanding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [26] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," *arXiv preprint arXiv:1412.3555*, 2014.
- [27] A. D. Laud, "Theory and application of reward shaping in reinforcement learning," 2004.
- [28] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," *arXiv preprint arXiv:1801.01290*, 2018.
- [29] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st ed. Cambridge, MA, USA: MIT Press, 1998.
- [30] Z. Yuan, J. Stroud, T. Lu, and J. Deng, "Temporal action localization by structured maximal sums," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [31] Z. Shou, D. Wang, and S.-F. Chang, "Temporal action localization in untrimmed videos via multi-stage cnns," in *CVPR*, 2016.
- [32] Z. Shou, J. Chan, A. Zareian, K. Miyazawa, and S.-F. Chang, "Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos," in *CVPR*, 2017.
- [33] J. Gao, Z. Yang, K. Chen, C. Sun, and R. Nevatia, "Turn tap: Temporal unit regression network for temporal action proposals," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [34] H. Xu, A. Das, and K. Saenko, "R-C3D: region convolutional 3d network for temporal activity detection," *CoRR*, 2017.
- [35] Y. Zhao, Y. Xiong, L. Wang, Z. Wu, X. Tang, and D. Lin, "Temporal action detection with structured segment networks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [36] J. Gao, Z. Yang, and R. Nevatia, "Cascaded boundary regression for temporal action detection," in *BMVC*, 2017.
- [37] T. Lin, X. Zhao, H. Su, C. Wang, and M. Yang, "BSN: boundary sensitive network for temporal action proposal generation," *CoRR*, 2018.
- [38] B. Seybold, D. Ross, J. Deng, R. Sukthankar, S. Vijayanarasimhan, and Y.-W. Chao, "Rethinking the faster r-cnn architecture for temporal action localization," in *CVPR 2018*, 2018.
- [39] Y. Huang, Q. Dai, and Y. Lu, "Decoupling localization and classification in single shot temporal action detection," *CoRR*, 2019.
- [40] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar, "THUMOS challenge: Action recognition with a large number of classes," <http://crcv.ucf.edu/THUMOS14/>, 2014.
- [41] F. C. Heilbron, V. Escorcia, B. Ghanem, and J. C. Niebles, "Activitynet: A large-scale video benchmark for human activity understanding," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 961–970.
- [42] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014.
- [43] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," *CoRR*, 2017.
- [44] S. Paul, S. Roy, and A. K. Roy-Chowdhury, "W-TALC: weakly-supervised temporal activity localization and classification," *CoRR*, 2018.
- [45] S. Paul, S. Roy, and A. K. Roy-Chowdhury, "W-talc: Weakly-supervised temporal activity localization and classification," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [46] S. Narayan, H. Cholakkal, F. S. Khan, and L. Shao, "3c-net: Category count and center loss for weakly-supervised action localization," 2019.
- [47] J. He, Y. Song, and H. Jiang, "Bi-direction feature pyramid temporal action detection network," 02 2020.
- [48] T. Lin, X. Zhao, H. Su, C. Wang, and M. Yang, "Bsn: Boundary sensitive network for temporal action proposal generation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [49] T. Lin, X. Liu, X. Li, E. Ding, and S. Wen, "Bmn: Boundary-matching network for temporal action proposal generation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- [50] Y. Huang, Q. Dai, and Y. Lu, "Decoupling localization and classification in single shot temporal action detection," in *2019 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2019.