



ECO: Efficient Convolutional Network for Online Video Understanding

Mohammadreza Zolfaghari^(✉) , Kamaljeet Singh, and Thomas Brox

University of Freiburg, Freiburg im Breisgau, Germany
{zolfagha,singhk,brox}@cs.uni-freiburg.de

Abstract. The state of the art in video understanding suffers from two problems: (1) The major part of reasoning is performed locally in the video, therefore, it misses important relationships within actions that span several seconds. (2) While there are local methods with fast per-frame processing, the processing of the whole video is not efficient and hampers fast video retrieval or online classification of long-term activities. In this paper, we introduce a network architecture (<https://github.com/mzolfaghari/ECO-efficient-video-understanding>) that takes long-term content into account and enables fast per-video processing at the same time. The architecture is based on merging long-term content already in the network rather than in a post-hoc fusion. Together with a sampling strategy, which exploits that neighboring frames are largely redundant, this yields high-quality action classification and video captioning at up to 230 videos per second, where each video can consist of a few hundred frames. The approach achieves competitive performance across all datasets while being 10× to 80× faster than state-of-the-art methods.

Keywords: Online video understanding · Real-time
Action recognition · Video captioning

1 Introduction

Video understanding and, specifically, action classification have benefited a lot from deep learning and the larger datasets that have been created in recent years. The new datasets, such as Kinetics [20], ActivityNet [13], and SomethingSomething [10] have contributed more diversity and realism to the field. Deep learning provides powerful classifiers at interactive frame rates, enabling applications like real-time action detection [30].

While action detection, which quickly decides on the present action within a short time window, is fast enough to run in real-time, activity understanding, which is concerned with longer-term activities that can span several seconds,

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01216-8_43) contains supplementary material, which is available to authorized users.

requires the integration of the long-term context to achieve full accuracy. Several 3D CNN architectures have been proposed to capture temporal relations between frames, but they are computationally expensive and, thus, can cover only comparatively small windows rather than the entire video. Existing methods typically use some post-hoc integration of window-based scores, which is suboptimal for exploiting the temporal relationships between the windows.

In this paper, we introduce a straightforward, end-to-end trainable architecture that exploits two important principles to avoid the above-mentioned dilemma. Firstly, a good initial classification of an action can already be obtained from just a single frame. The temporal neighborhood of this frame comprises mostly redundant information and is almost useless for improving the belief about the present action¹. Therefore, we process only a single frame of a temporal neighborhood efficiently with a 2D convolutional architecture in order to capture appearance features of such frame. Secondly, to capture the contextual relationships between distant frames, a simple aggregation of scores is insufficient. Therefore, we feed the feature representations of distant frames into a 3D network that learns the temporal context between these frames and so can improve significantly over the belief obtained from a single frame – especially for complex long-term activities. This principle is much related to the so-called early or late fusion used for combining the RGB stream and the optical flow stream in two-stream architectures [8]. However, this principle has been mostly ignored so far for aggregation over time and is not part of the state-of-the-art approaches.

Consequent implementation of these two principles together leads to a high classification accuracy without bells and whistles. The long temporal context of complex actions can be fully captured, whereas the fact that the method only looks at a very small fraction of all frames in the video leads to extremely fast processing of entire videos. This is very beneficial especially in video retrieval applications.

Additionally, this approach opens the possibility for online video understanding. In this paper, we also present a way to use our architecture in an online setting, where we provide a fast first guess on the action and refine it using the longer term context as a more complex activity establishes. In contrast to online action detection, which has been enabled recently [30], the approach provides not only fast reaction times, but also takes the longer term context into account.

We conducted experiments on various video understanding problems including action recognition and video captioning. Although we just use RGB images as input, we obtain on-par or favorable performance compared to state-of-the-art approaches on most datasets. The runtime-accuracy trade-off is superior on all datasets.

¹ An exception is the use of two frames for capturing motion, which could be achieved by optionally feeding optical flow together with the RGB image. In this paper, we only provide RGB images, but an extension with optical flow, e.g., a fast variant of FlowNet [16] would be straightforward.

2 Related Work

Video Classification with Deep Learning. Most recent works on video classification are based on deep learning [6, 19, 29, 34, 48]. To explore the temporal context of a video, 3D convolutional networks are an obvious option. Tran et al. [33] introduced a 3D architecture with 3D kernels to learn spatio-temporal features from a sequence of frames. In a later work, they studied the use of a Resnet architecture with 3D convolutions and showed the improvements over their earlier c3d architecture [34]. An alternative way to model the temporal relation between frames is by using recurrent networks [6, 24, 25]. Donahue et al. [6] employed a LSTM to integrate features from a CNN over time. However, the performance of recurrent networks on action recognition currently lags behind that of recent CNN-based methods, which may indicate that they do not sufficiently model long-term dynamics [24, 25]. Recently, several works utilized 3D architectures for action recognition [5, 35, 39, 48]. These approaches model the short-term temporal context of the input video based on a sliding window. At inference time, these methods must compute the average score over multiple windows, which is quite time consuming. For example, ARTNet [39] requires on average 250 samples to classify one video.

All these approaches do not sufficiently use the comprehensive information from the entire video during training and inference. Partial observation not only causes confusion in action prediction, but also requires an extra post-processing step to fuse scores. Extra feature/score aggregation reduces the speed of video processing and disables the method to work in a real-time setting.

Long-Term Representation Learning. To cope with the problem of partial observation, some methods increased the temporal resolution of the sliding window [4, 36]. However, expanding the temporal length of the input has two major drawbacks. (1) It is computationally expensive, and (2) still fails to cover the visual information of the entire video, especially for longer videos.

Some works proposed encoding methods [26, 28, 42] to learn a video representation from samples. In these approaches, features are usually calculated for each frame independently and are aggregated across time to make a video-level representation. This ignores the relationship between the frames.

To capture long-term information, recent works [2, 3, 7, 40] employed a sparse and global temporal sampling method to choose frames from the entire video during training. In the TSN model [40], as in the aggregation methods above, frames are processed independently at inference time and their scores are aggregated only in the end. Consequently, the performance in their experiments stays the same when they change the number of samples, which indicates that their model does not really benefit from the long-range temporal information.

Our work is different from these previous approaches in three main aspects: (1) Similar to TSN, we sample a fixed number of frames from the entire video to cover long-range temporal structure for understanding of video. In this way, the sampled frames span the entire video independent of the length of the video. (2) In contrast to TSN, we use a 3D-network to learn the relationship between

the frames throughout the video. The network is trained end-to-end to learn this relationship. (3) The network directly provides video-level scores without post-hoc feature aggregation. Therefore, it can be run in online mode and in real-time even on small computing devices.

Video Captioning. Video captioning is a widely studied problem in computer vision [9, 14, 43, 45]. Most approaches use a CNN pre-trained on image classification or action recognition to generate features [9, 43, 45]. These methods, like the video understanding methods described above, utilize a frame-based feature aggregation (e.g. Resnet or TSN) or a sliding window over the whole video (e.g. C3D) to generate video-level features. The features are then passed to a recurrent neural network (e.g. LSTM) to generate the video captions via a learned language model. The extracted visual features should represent both the temporal structure of the video and the static semantics of the scene. However, most approaches suffer from the problem that the temporal context is not properly extracted. With the network model in this work, we address this problem, and can consequently improve video captioning results.

Real-Time and Online Video Understanding. Deep learning accelerated image classification, but video classification remains challenging in terms of speed. A few works dealt with real-time video understanding [18, 30, 31, 44]. EMV [44] introduced an approach for fast calculation of motion vectors. Despite this improvement, video processing is still slow. Kantorov [18] introduced a fast dense trajectory method. The other works used frame-based hand-crafted features for online action recognition [15, 22]. Both accuracy and speed of feature extraction in these methods are far from that of deep learning methods. Soomro et al. [31] proposed an online action localization approach. Their model utilizes an expensive segmentation method which, therefore, cannot work in real-time. More recently, Singh et al. [30] proposed an online detection approach based on frame-level detections at 40 fps. We compare to the last two approaches in Sect. 5.

3 Long-Term Spatio-Temporal Architecture

The network architecture is shown in Fig. 1. A whole video with a variable number of frames is provided as input to the network. The video is split into N subsections S_i , $i = 1, \dots, N$ of equal size, and in each subsection, exactly one frame is sampled randomly. Each of these frames is processed by a single 2D convolutional network (weight sharing), which yields a feature representation encoding the frame’s appearance. By jointly processing frames from time segments that cover the whole video, we make sure that we capture the most relevant parts of an action over time and the relationship among these parts.

Randomly sampling the position of the frame is advantageous over always using the same position, because it leads to more diversity during training and makes the network adapt to variations in the instantiation of an action. Note that this kind of processing exploits all frames of the video during training to

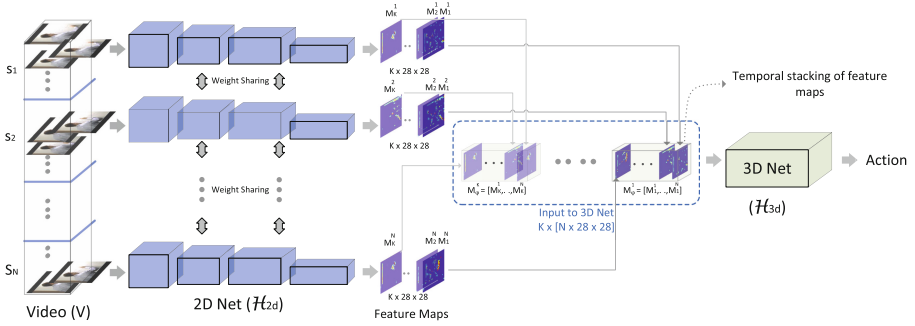


Fig. 1. Architecture overview of ECO Lite. Each video is split into N subsections of equal size. From each subsection a single frame is randomly sampled. The samples are processed by a regular 2D convolutional network to yield a representation for each sampled frame. These representations are stacked and fed into a 3D convolutional network, which classifies the action, taking into account the temporal relationship.

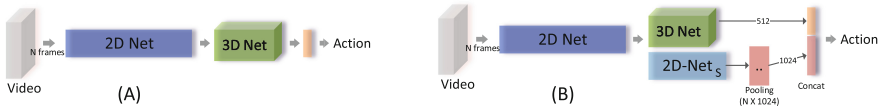


Fig. 2. (A) ECO Lite architecture as shown in more detail in Fig. 1. (B) Full ECO architecture with a parallel 2D and 3D stream.

model the variation. At the same time, the network must only process N frames at runtime, which makes the approach very fast. We also considered more clever partitioning strategies that take the content of the subsections into account. However, this comes with the drawback that each frame of the video must be processed at runtime to obtain the partitioning, and the actual improvement by such smarter partitioning is limited, since most of the variation is already captured by the random sampling during training.

Up to this point, the different frames in the video are processed independently. In order to learn how actions are made up of the different appearances of the scene over time, we stack the representations of all frames and feed them into a 3D convolutional network. This network yields the final action class label.

The architecture is very straightforward, and it is obvious that it can be trained efficiently end-to-end directly on the action class label and on large datasets. It is also an architecture that can be easily adapted to other video understanding tasks, as we show later in the video captioning Sect. 5.4.

3.1 ECO Lite and ECO Full

The 3D architecture in ECO Lite is optimized for learning relationships between the frames, but it tends to waste capacity in case of simple short-term actions that can be recognized just from the static image content. Therefore, we suggest

an extension of the architecture by using a 2D network in parallel; see Fig. 2(B). For the simple actions, this 2D network architecture can simplify processing and ensure that the static image features receive the necessary importance, whereas the 3D network architecture takes care of the more complex actions that depend on the relationship between frames.

The 2D network receives feature maps of all samples and produces N feature representations. Afterwards, we apply average pooling to get a feature vector that is a representative for static scene semantics. We call the full architecture ECO and the simpler architecture in Fig. 2(A) ECO Lite.

3.2 Network Details

2D-Net: For the 2D network (\mathcal{H}_{2D}) that analyzes the single frames, we use the first part of the BN-Inception architecture (until inception-3c layer) [17]. Details are given in the supplemental material. It has 2D filters and pooling kernels with batch normalization. We chose this architecture due to its efficiency. The output of \mathcal{H}_{2D} for each single frame consist of 96 feature maps with size of 28×28 .

3D-Net: For the 3D network \mathcal{H}_{3D} we adopt several layers of 3D-Resnet18 [34], which is an efficient architecture used in many video classification works [34, 39]. Details on the architecture are provided in the supplemental material. The output of \mathcal{H}_{3D} is a one-hot vector for the different class labels.

2D-Net_s: In the ECO full design, we use 2D-Net_s in parallel with 3D-net to directly providing static visual semantics of video. For this network, we use the BN-Inception architecture from inception-4a layer until last pooling layer [17]. The last pooling layer will produce 1024 dimensional feature vector for each frame. We apply average pooling to generate video-level feature and then concatenate with features obtained from 3D-net.

3.3 Training Details

We train our networks using mini-batch SGD with Nesterov momentum and utilize dropout in each fully connected layer. We split each video into N segments and randomly select one frame from each segment. This sampling provides robustness to variations and enables the network to fully exploit all frames. In addition, we apply the data augmentation techniques introduced in [41]: we resize the input frames to 240×320 and employ fixed-corner cropping and scale jittering with horizontal flipping (temporal jittering provided by sampling). Afterwards, we run per-pixel mean subtraction and resize the cropped regions to 224×224 .

The initial learning rate is 0.001 and decreases by a factor of 10 when validation error saturates for 4 epochs. We train the network with a momentum of 0.9, a weight decay of 0.0005, and mini-batches of size 32.

We initialize the weights of the 2D-Net weights with the BN-Inception architecture [17] pre-trained on Kinetics, as provided by [41]. In the same way, we use

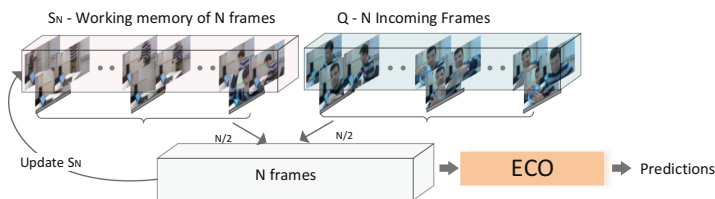


Fig. 3. Scheme of our sampling strategy for online video understanding. Half of the frames are sampled uniformly from the working memory in the previous time step, the other half from the queue (Q) of incoming frames.

the pre-trained model of 3D-Resnet-18, as provided by [39] for initializing the weights of our 3D-Net. Afterwards, we train ECO and ECO Lite on the Kinetics dataset for 10 epochs.

For other datasets, we finetune the above ECO/ECO Lite models on the new datasets. Due to the complexity of the Something-Something dataset, we finetune the network for 25 epochs reducing the learning rate every 10 epochs by a factor of 10. For the rest, we finetune for 4k iterations and the learning rate drops by a factor of 10 as soon as the validation loss saturates. The whole training process on UCF101 and HMDB51 takes around 3 h on one Tesla P100 GPU for the ECO architecture. We adjusted the dropout rate and the number of iterations based on the dataset size.

3.4 Test Time Inference

Most state-of-the-art methods run some post-processing on the network result. For instance, TSN and ARTNet [39, 41], collect 25 independent frames/volumes per video, and for each frame/volume sample 10 regions by corner and center cropping, and their horizontal flipping. The final prediction is obtained by averaging the scores of all 250 samples. This kind of inference at test time is computationally expensive and thus unsuitable for real-time setting.

In contrast, our network produces action labels for the whole video directly without any additional aggregation. We sample N frames from the video, apply only center cropping then feed them directly to the network, which provides the prediction for the whole video with a single pass.

4 Online Video Understanding

Most works on video understanding process in batch mode, i.e., they assume that the whole video is available when processing starts. However, in several application scenarios, the video will be provided as a stream and the current belief is supposed to be available at any time. Such online processing is possible with a sliding window approach, yet this comes with restrictions regarding the size of the window, i.e., long-term context is missing, or with a very long delay.

Algorithm 1. Online video understanding

Input : Live video stream (V), ECO pretrained model (ECO_{NF}), Number of Samples = Sampling window (N)

Output: Predictions

Initialize an empty queue Q to queue N incoming frames;

Initialize working memory S_N ;

Initialize average predictions P_A ;

while new frames available from V **do**

 Add frame f_i from V to queue Q ;

if $i \% N$ **then**

$S_N :=$ Sample 50% frames Q and 50% from S_N ;

 Empty queue Q ;

 Feed S_N to model ECO_{NF} to get output probabilities P ;

$P_A :=$ Average P and P_A ;

 Output average predictions P_A ;

end

end

In this section, we show how ECO can be adapted to run very efficiently in online mode, too. The modification only affects the sampling part and keeps the network architecture unchanged. To this end, we partition the incoming video content into segments of N frames, where N is also the number of frames that go into the network. We use a working memory S_N , which always comprises the N samples that are fed to the network together with a time stamp. When a video starts, i.e., only N frames are available, all N frames are sampled densely and are stored in the working memory S_N . With each new time segment, N additional frames come in, and we replace half of the samples in S_N by samples from this time segment and update the prediction of the network; see Fig. 3. When we replace samples from S_N , we uniformly replace samples from previous time segments. This ensures that changes can be anticipated in real time, while the temporal context is taken into account and slowly fades out via the working memory. Details on the update of S_N are shown in Algorithm 1.

The online approach with ECO runs at 675 fps (and at 970 fps with ECO Lite) on a Tesla P100 GPU. In addition, the model is memory efficient by just keeping exactly N frames. This enables the implementation also on much smaller hardware, such as mobile devices. The video in the supplemental material shows the recorded performance of the online version of ECO in real-time.

5 Experiments

We evaluate our approach on different video understanding problems to show the generalization ability of approach. We evaluated the network architecture on the most common action classification datasets in order to compare its performance against the state-of-the-art approaches. This includes the older but still very popular datasets UCF101 [32] and HMDB51 [21], but also the more recent

Table 1. Comparison to the state-of-the-art on UCF101 and HMDB51 datasets (over all three splits), using just RGB modality.

Method	Pre-training	Dataset	
		UCF101 (%)	HMDB51 (%)
I3D [4]	ImageNet	84.5	49.8
TSN [41]	ImageNet	86.4	53.7
DTPP [47]	ImageNet	89.7	61.1
Res3D [34]	Sports-1M	85.8	54.9
TSN [41]	ImageNet + Kinetics	91.1	-
I3D [4]	ImageNet + Kinetics	95.6	74.8
ResNeXt-101 [12]	Kinetics	94.5	70.2
ARTNet [39]	Kinetics	93.5	67.6
T3D [5]	Kinetics	91.7	61.1
ECO _{E_n}	Kinetics	94.8	72.4

Table 2. Comparing performance of ECO with state-of-the-art methods on the Kinetics dataset.

Methods	Val (%)		Test (%)
	Top-1	Avg	Avg
ResNeXt-101 [12]	65.1	75.4	78.4
Res3D [34]	65.6	75.7	74.4
I3D-RGB [4]	–	–	78.2
ARTNet [39]	69.2	78.7	77.3
T3D [5]	62.2	–	71.5
ECO _{E_n}	70.0	79.7	76.3

Table 3. Comparison with state-of-the-arts on something-something dataset. Last row shows the results using both Flow and RGB.

Methods	Val (%)	Test (%)
I3D by [10]	-	27.23
M-TRN [46]	34.44	33.60
ECO _{E_n} Lite	46.4	42.3
ECO _{E_n} Lite $\left\{ \begin{array}{l} RGB \\ Flow \end{array} \right.$	49.5	43.9

datasets Kinetics [20] and Something-Something [10]. Moreover, we applied the architecture to video captioning and tested it on the widely used Youtube2text dataset [11]. For all of these datasets, we use the standard evaluation protocol provided by the authors.

The comparison is restricted to approaches that take the raw RGB videos as input without further pre-processing, for instance, by providing optical flow or human pose. The term ECO_{NF} describes a network that gets N sampled frames as input. The term ECO_{E_n} refers to average scores obtained from an ensemble of networks with $\{16, 20, 24, 32\}$ number of frames.

5.1 Benchmark Comparison on Action Classification

The results obtained with ECO on the different datasets are shown in Tables 1, 2, and 3 and compare them to the state of the art. For UCF-101, HMDB-51, and Kinetics, ECO outperforms all existing methods except I3D, which uses a much heavier network. On Something-Something, it outperforms the other methods with a large margin. This shows the strong performance of the comparatively simple and small ECO architecture.

Table 4. Runtime comparison with state-of-the-art approaches using Tesla P100 GPU on UCF101 and HMDB51 datasets (over all splits). For other approaches, we just consider one crop per sample to calculate the runtime. We reported runtime without considering I/O.

Method	Inference speed (VPS)	UCF101 (%)	HMDB51 (%)
Res3D [34]	<2	85.8	54.9
TSN [41]	21	87.7	51
EMV [44]	15.6	86.4	-
I3D [4]	0.9	95.6	74.8
ARTNet [39]	2.9	93.5	67.6
ECO _{Lite-4F}	237.3	87.4	58.1
ECO _{4F}	163.4	90.3	61.7
ECO _{12F}	52.6	92.4	68.3
ECO _{20F}	32.9	93.0	69.0
ECO _{24F}	28.2	93.6	68.4

Table 5. Accuracy and runtime of ECO and ECO Lite for different numbers of sampled frames. The reported runtime is without considering I/O.

Model	Sampled frames	Speed (VPS)		Accuracy (%)			
		Titan X	Tesla P100	UCF101	HMDB51	Kinetics	Someth
ECO	4	99.2	163.4	90.3	61.7	66.2	—
	8	49.5	81.5	91.7	65.6	67.8	39.6
	16	24.5	41.7	92.8	68.5	69.0	41.4
	32	12.3	20.8	93.3	68.7	67.8	—
ECO Lite	4	142.9	237.3	87.4	58.1	57.9	—
	8	71.1	115.9	90.2	63.3	—	38.7
	16	35.3	61.0	91.6	68.2	64.4	42.2
	32	18.2	30.2	93.1	68.3	—	41.3

5.2 Accuracy-Runtime Comparison

The advantages of the ECO architectures becomes even more prominent as we look at the accuracy-runtime trade-off shown in Table 4 and Fig. 4. The ECO architectures yield the same accuracy as other approaches at much faster rates.

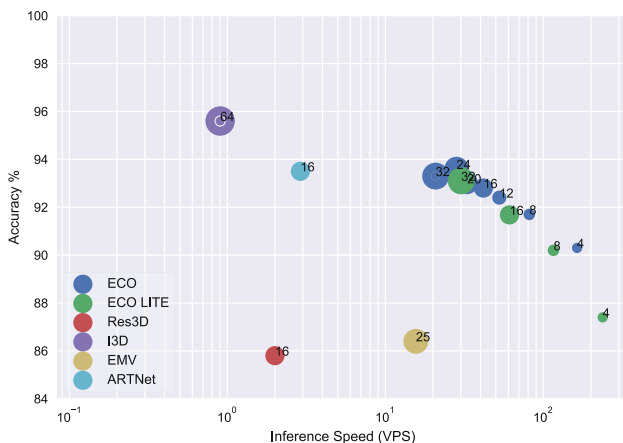


Fig. 4. Accuracy-runtime trade-off on UCF101 for various versions of ECO and other state-of-the-art approaches. ECO is much closer to the top right corner.




Fig. 5. Examples from the Something-Something dataset. In this dataset, the temporal context plays an even bigger role than on other datasets, since the same action is done with different objects, i.e., the appearance of the object or background gives almost no cues about the action.

Previous works typically measure the speed of an approach in frames per second (fps). Our model with ECO runs at 675 fps (and at 970 fps with ECO Lite) on a Tesla P100 GPU. However, this does not reflect the time needed to process a whole video. This becomes relevant for methods like TSN and ours, which do not look at every frame of the video, and motivates us to report *videos per second (vps)* to compare the speed of video understanding methods.

ECO can process videos at least an order of magnitude faster than the other approaches, making it an excellent architecture for video retrieval applications.

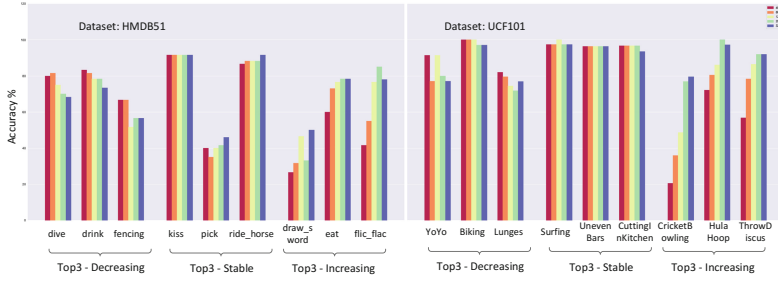


Fig. 6. Effect of the complexity of an action on the need for denser sampling. While simple short-term actions (leftmost group) even suffer from more samples, complex actions (rightmost group) clearly benefit from a denser sampling.

Number of Sampled Frames. Table 5 compares the two architecture variants ECO and ECO Lite and evaluates the influence on the number of sampled frames N . As expected, the accuracy drops when sampling fewer frames, as the subsections get longer and important parts of the action can be missed. This is especially true for fast actions, such as “throw discus”. However, even with just 4 samples the accuracy of ECO is still much better than most approaches in literature, since ECO takes into account the relationship between these 4 instants in the video, even if they are far apart. Figure 6 even shows that for simple short-term actions, the performance decreases when using more samples. This is surprising on first glance, but could be explained by the better use of the network’s capacity for simple actions when there are fewer channels being fed to the 3D network.

ECO Vs. ECO Lite. The full ECO architecture yields slightly better results than the plain ECO Lite architecture, but is also a little slower. The differences in accuracy and runtime between the two architectures can usually be compensated by using more or fewer samples. On the Something-Something dataset, where the temporal context plays a much bigger role than on other datasets (see Fig. 5), ECO Lite performs equally well as the full ECO architecture even with the same number of input samples, since the raw processing of single image cues has little relevance on this dataset.

5.3 Early Action Recognition in Online Mode

Figure 7 evaluates our approach in online mode and shows how many frames the method needs to achieve its full accuracy. We ran this experiment on the J-HMDB dataset due to the availability of results from other online methods on this dataset. Compared to these existing methods, ECO reaches a good accuracy faster and also saturates at a higher absolute accuracy.

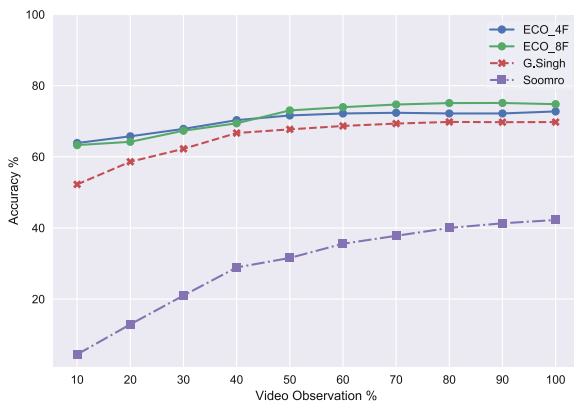


Fig. 7. Early action classification results of ECO in comparison to existing online methods [30, 31] on the J-HMDB dataset. The online version of ECO yields a high accuracy already after seeing a short part of the video. Singh et al. [30] uses both RGB and optical flow.

Table 6. Captioning results on Youtube2Text (MSVD) dataset.

Methods	Metrics			
	B-3	B-4	METEOR	CIDEr
S2VT [38]	-	-	0.292	-
GRU-RCN [1]	-	0.479	0.311	0.678
h-RNN [43]	-	0.499	0.326	0.658
TDDF [45]	-	0.458	0.333	0.730
AF [14]	-	0.524	0.320	0.688
SCN-c3d [9]	0.587	0.482	0.330	0.692
SCN-resnet [9]	0.602	0.506	0.336	0.755
SCN-ensemble of 5 [9]	-	0.511	0.335	0.777
ECO _{Lite-16F}	0.601	0.504	0.339	0.833
ECO _{32F}	0.616	0.521	0.345	0.857
ECO _{32F} + resnet	0.626	0.535	0.350	0.858

5.4 Video Captioning

To show the wide applicability of the ECO architecture, we also combine it with a video captioning network. To this end, we use ECO pre-trained on Kinetics to analyze the video content and train the state-of-the-art Semantic Compositional Network [9] for captioning. We evaluated on the Youtube2Text (MSVD) dataset [11], which consists of 1,970 video clips with an average duration of 9 s and covers various types of videos, such as sports, landscapes, animals, cooking, and human activities. The dataset contains 80,839 sentences and each video is annotated with around 40 sentences.

Table 7. Qualitative results on MSVD. First row corresponds to the examples where ECO improved over SCN and the second row shows the examples where ECO decreased the quality compared to SCN. ECO_L refers to $ECO_{Lite-16F}$, ECO to ECO_{32F} , and ECO_R to $ECO_{32F+resnet}$.





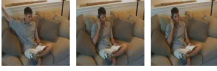

 <p>SCN: a woman is cooking ECO_L: the woman is seasoning the meat ECO: a woman is seasoning some meat ECO_R: a woman is seasoning some meat</p>	 <p>SCN: a man is playing a flute ECO: a man is playing a violin ECO_R: a man is playing a violin</p>	 <p>SCN: a man is cooking ECO_L: a man is pouring water into a container ECO: a man is putting a lid on a plastic container ECO_R: a man is draining pasta</p>
 <p>SCN: a man is riding a horse ECO_L: a woman is riding a motorcycle ECO: a man is riding a horse ECO_R: a man is riding a boat</p>	 <p>SCN: a girl is sitting on a couch ECO_L: a baby is sitting on the bed ECO: a woman is playing with a toy ECO_R: a woman is sleeping on a bed</p>	 <p>SCN: two elephants are walking ECO_L: a rhino is walking ECO: a group of elephants are walking ECO_R: a penguin is walking</p>

Table 6 shows that ECO compares favorably to previous approaches across all popular evaluation metrics (BLEU [27], METEOR [23], CIDEr [37]). Even ECO Lite is already on-par with a ResNet architecture pre-trained on ImageNet. Concatenating the features from ECO with those of ResNet improves results further. Qualitative examples that correspond to the improved numbers are shown in Table 7.

6 Conclusions

In this paper, we have presented a simple and very efficient network architecture that looks only at a small subset of frames from a video and learns to exploit the temporal context between these frames. This principle can be used in various video understanding tasks. We demonstrate excellent results on action classification, online action classification, and video captioning. The computational load and the memory footprint makes an implementation on mobile devices a viable future option. The approaches runs $10\times$ to $80\times$ faster than state-of-the-art methods.

Acknowledgements. We thank Facebook for providing us a GPU server with Tesla P100 processors for this research work.

References

1. Ballas, N., Yao, L., Pal, C., Courville, A.C.: Delving deeper into convolutional networks for learning video representations. In: ICLR (2016)
2. Bilen, H., Fernando, B., Gavves, E., Vedaldi, A.: Action recognition with dynamic image networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1 (2017). <https://doi.org/10.1109/TPAMI.2017.2769085>
3. Bilen, H., Fernando, B., Gavves, E., Vedaldi, A., Gould, S.: Dynamic image networks for action recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3034–3042, June 2016. <https://doi.org/10.1109/CVPR.2016.331>
4. Carreira, J., Zisserman, A.: Quo vadis, action recognition? A new model and the kinetics dataset. *CoRR abs/1705.07750* (2017). <http://arxiv.org/abs/1705.07750>
5. Diba, A., et al.: Temporal 3D ConvNets: new architecture and transfer learning for video classification. *CoRR abs/1711.08200* (2017). <http://arxiv.org/abs/1711.08200>
6. Donahue, J., et al.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR (2015)
7. Feichtenhofer, C., Pinz, A., Wildes, R.P.: Spatiotemporal multiplier networks for video action recognition. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7445–7454, July 2017. <https://doi.org/10.1109/CVPR.2017.787>
8. Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. *CoRR abs/1604.06573* (2016). <http://arxiv.org/abs/1604.06573>
9. Gan, Z., et al.: Semantic compositional networks for visual captioning. In: CVPR (2017)
10. Goyal, R., et al.: The “something something” video database for learning and evaluating visual common sense. *CoRR abs/1706.04261* (2017). <http://arxiv.org/abs/1706.04261>
11. Guadarrama, S., et al.: YouTube2Text: recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition. In: 2013 IEEE International Conference on Computer Vision, pp. 2712–2719, December 2013. <https://doi.org/10.1109/ICCV.2013.337>
12. Hara, K., Kataoka, H., Satoh, Y.: Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? *CoRR abs/1711.09577* (2017). <http://arxiv.org/abs/1711.09577>
13. Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C.: ActivityNet: a large-scale video benchmark for human activity understanding. In: CVPR, pp. 961–970. IEEE Computer Society (2015). <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2015.html#HeilbronEGN15>
14. Hori, C., et al.: Attention-based multimodal fusion for video description. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 4203–4212, October 2017. <https://doi.org/10.1109/ICCV.2017.450>
15. Hu, B., Yuan, J., Wu, Y.: Discriminative action states discovery for online action recognition. *IEEE Signal Process. Lett.* **23**(10), 1374–1378 (2016). <https://doi.org/10.1109/LSP.2016.2598878>
16. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: evolution of optical flow estimation with deep networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017. <http://lmb.informatik.uni-freiburg.de//Publications/2017/IMKDB17>

17. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML 2015, pp. 448–456. JMLR.org (2015). <http://dl.acm.org/citation.cfm?id=3045118.3045167>
18. Kantorov, V., Laptev, I.: Efficient feature extraction, encoding, and classification for action recognition. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition, pp. 2593–2600, June 2014. <https://doi.org/10.1109/CVPR.2014.332>
19. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, pp. 1725–1732. IEEE Computer Society, Washington (2014). <https://doi.org/10.1109/CVPR.2014.223>
20. Kay, W., et al.: The kinetics human action video dataset. CoRR abs/1705.06950 (2017). <http://arxiv.org/abs/1705.06950>
21. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: a large video database for human motion recognition. In: Proceedings of the International Conference on Computer Vision (ICCV) (2011)
22. Kviatkovsky, I., Rivlin, E., Shimshoni, I.: Online action recognition using covariance of shape and motion. *Comput. Vis. Image Underst.* **129**(C), 15–26 (2014). <https://doi.org/10.1016/j.cviu.2014.08.001>
23. Lavie, A., Agarwal, A.: METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments. In: Proceedings of the Second Workshop on Statistical Machine Translation, StatMT 2007, pp. 228–231. Association for Computational Linguistics, Stroudsburg (2007). <http://dl.acm.org/citation.cfm?id=1626355.1626389>
24. Lev, G., Sadeh, G., Klein, B., Wolf, L.: RNN fisher vectors for action recognition and image annotation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9910, pp. 833–850. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46466-4_50
25. Li, Z., Gavriluk, K., Gavves, E., Jain, M., Snoek, C.G.: VideoLSTM convolves, attends and flows for action recognition. *Comput. Vis. Image Underst.* **166**(C), 41–50 (2018). <https://doi.org/10.1016/j.cviu.2017.10.011>
26. Ng, J.Y.H., Hausknecht, M.J., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G.: Beyond short snippets: deep networks for video classification. In: CVPR, pp. 4694–4702. IEEE Computer Society (2015). <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2015.html#NgHVVM15>
27. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL 2002, pp. 311–318. Association for Computational Linguistics, Stroudsburg (2002). <https://doi.org/10.3115/1073083.1073135>
28. Qiu, Z., Yao, T., Mei, T.: Deep quantization: encoding convolutional activations with deep generative model. In: CVPR (2017)
29. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS 2014, pp. 568–576. MIT Press, Cambridge (2014). <http://dl.acm.org/citation.cfm?id=2968826.2968890>

30. Singh, G., Saha, S., Sapienza, M., Torr, P.H.S., Cuzzolin, F.: Online real-time multiple spatiotemporal action localisation and prediction. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, 22–29 October 2017, pp. 3657–3666 (2017). <https://doi.org/10.1109/ICCV.2017.393>
31. Soomro, K., Idrees, H., Shah, M.: Predicting the where and what of actors and actions through online action localization. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2016
32. Soomro, K., Zamir, A.R., Shah, M.: UCF101: a dataset of 101 human actions classes from videos in the wild. CoRR abs/1212.0402 (2012). <http://arxiv.org/abs/1212.0402>
33. Tran, D., Bourdev, L.D., Fergus, R., Torresani, L., Paluri, M.: C3D: generic features for video analysis. CoRR abs/1412.0767 (2014). <http://arxiv.org/abs/1412.0767>
34. Tran, D., Ray, J., Shou, Z., Chang, S., Paluri, M.: ConvNet architecture search for spatiotemporal feature learning. CoRR abs/1708.05038 (2017). <http://arxiv.org/abs/1708.05038>
35. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. CoRR abs/1711.11248 (2017). <http://arxiv.org/abs/1711.11248>
36. Varol, G., Laptev, I., Schmid, C.: Long-term temporal convolutions for action recognition. CoRR abs/1604.04494 (2016). <http://arxiv.org/abs/1604.04494>
37. Vedantam, R., Zitnick, C.L., Parikh, D.: CIDEr: consensus-based image description evaluation. In: CVPR, pp. 4566–4575. IEEE Computer Society (2015). <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2015.html#VedantamZP15>
38. Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R., Darrell, T., Saenko, K.: Sequence to sequence - video to text. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2015)
39. Wang, L., Li, W., Li, W., Gool, L.V.: Appearance-and-relation networks for video classification. CoRR abs/1711.09125 (2017). <http://arxiv.org/abs/1711.09125>
40. Wang, L., et al.: Temporal segment networks for action recognition in videos. CoRR abs/1705.02953 (2017). <http://arxiv.org/abs/1705.02953>
41. Wang, L., et al.: Temporal segment networks: towards good practices for deep action recognition. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9912, pp. 20–36. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_2
42. Xu, Z., Yang, Y., Hauptmann, A.G.: A discriminative CNN video representation for event detection. CoRR abs/1411.4006 (2014). <http://arxiv.org/abs/1411.4006>
43. Yu, H., Wang, J., Huang, Z., Yang, Y., Xu, W.: Video paragraph captioning using hierarchical recurrent neural networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4584–4593 (2016)
44. Zhang, B., Wang, L., Wang, Z., Qiao, Y., Wang, H.: Real-time action recognition with enhanced motion vector CNNs. CoRR abs/1604.07669 (2016). <http://arxiv.org/abs/1604.07669>
45. Zhang, X., Gao, K., Zhang, Y., Zhang, D., Li, J., Tian, Q.: Task-driven dynamic fusion: reducing ambiguity in video description. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 6250–6258, July 2017. <https://doi.org/10.1109/CVPR.2017.662>
46. Zhou, B., Andonian, A., Torralba, A.: Temporal relational reasoning in videos. CoRR abs/1711.08496 (2017). <http://arxiv.org/abs/1711.08496>

47. Zhu, J., Zou, W., Zhu, Z., Li, L.: End-to-end video-level representation learning for action recognition. CoRR abs/1711.04161 (2017). <http://arxiv.org/abs/1711.04161>
48. Zolfaghari, M., Oliveira, G.L., Sedaghat, N., Brox, T.: Chained multi-stream networks exploiting pose, motion, and appearance for action classification and detection. In: IEEE International Conference on Computer Vision (ICCV) (2017). <http://lmb.informatik.uni-freiburg.de/Publications/2017/ZOSB17a>