

TEA: Temporal Excitation and Aggregation for Action Recognition

Yan Li¹ Bin Ji² Xintian Shi¹ Jianguo Zhang^{3*} Bin Kang^{1*} Limin Wang^{2*}

¹ Platform and Content Group (PCG), Tencent

² State Key Laboratory for Novel Software Technology, Nanjing University, China

³ Department of Computer Science and Engineering, Southern University of Science and Technology, China

phoenixyli@tencent.com, binjinju@smail.nju.edu.cn, tinaxtshi@tencent.com

zhangjg@sustech.edu.cn, binkang@tencent.com, 07wanglimin@gmail.com

Abstract

Temporal modeling is key for action recognition in videos. It normally considers both *short-range* motions and *long-range* aggregations. In this paper, we propose a Temporal Excitation and Aggregation (TEA) block, including a motion excitation (ME) module and a multiple temporal aggregation (MTA) module, specifically designed to capture both short- and long-range temporal evolution. In particular, for short-range motion modeling, the ME module calculates the feature-level temporal differences from spatiotemporal features. It then utilizes the differences to excite the motion-sensitive channels of the features. The long-range temporal aggregations in previous works are typically achieved by stacking a large number of local temporal convolutions. Each convolution processes a local temporal window at a time. In contrast, the MTA module proposes to deform the local convolution to a group of sub-convolutions, forming a hierarchical residual architecture. Without introducing additional parameters, the features will be processed with a series of sub-convolutions, and each frame could complete multiple temporal aggregations with neighborhoods. The final equivalent receptive field of temporal dimension is accordingly enlarged, which is capable of modeling the long-range temporal relationship over distant frames. The two components of the TEA block are complementary in temporal modeling. Finally, our approach achieves impressive results at low FLOPs on several action recognition benchmarks, such as Kinetics, Something-Something, HMDB51, and UCF101, which confirms its effectiveness and efficiency.

1. Introduction

Action recognition is a fundamental problem in video-based tasks. It becomes increasingly demanding in video-

based applications, such as intelligent surveillance, autonomous driving, personal recommendation, and entertainment [28]. Though visual appearances (and its context) is important for action recognition, it is rather important to model the temporal structure. Temporal modeling normally presents (or is considered) at different scales: 1) *short-range* motion between adjacent frames and 2) *long-range* temporal aggregation at large scales. There are lines of works considering one or both of those aspects, especially in the current era of deep CNNs [21, 31, 47, 6, 36, 46, 34, 2, 1, 39, 50, 49, 27, 41, 29, 24, 20]. Nevertheless, they still leave some gaps, and the problem is far from being solved, *i.e.*, it remains unclear how to model the temporal structure with significant variations and complexities effectively and efficiently.

For short-range motion encoding, most of the existing methods [31, 42] extract hand-crafted optical flow [48] first, which is then fed into a 2D CNN-based two-stream framework for action recognition. Such a two-stream architecture processes RGB images and optical flow in each stream separately. The computation of optical flow is time-consuming and storage demanding. In particular, the learning of spatial and temporal features is isolated, and the fusion is performed only at the late layers. To address these issues, we propose a *motion excitation* (ME) module. Instead of adopting the pixel-level optical flow as an additional input modality and separating the training of temporal stream with the spatial stream, our module could integrate the motion modeling into the whole spatiotemporal feature learning approach. Concretely, the feature-level motion representations are firstly calculated between adjacent frames. These motion features are then utilized to produce modulation weights. Finally, the motion-sensitive information in the original features of frames can be excited with the weights. In this way, the networks are forced to discover and enhance the informative temporal features that capture differentiated information.

For long-range temporal aggregation, existing methods

*Corresponding authors.

either 1) adopt 2D CNN backbones to extract frame-wise features and then utilize a simple temporal max/average pooling to obtain the whole video representation [42, 11]. Such a simple summarization strategy, however, results in temporal information loss/confusion; or 2) adopt local 3D/(2+1)D convolutional operations to process local temporal window [36, 3]. The long-range temporal relationship is indirectly modeled by repeatedly stacking local convolutions in deep networks. However, repeating a large number of local operations will lead to optimization difficulty [14], as the message needs to be propagated through the long path between distant frames. To tackle this problem, we introduce a *multiple temporal aggregation* (MTA) module. The MTA module also adopts (2+1)D convolutions, but a group of sub-convolutions replaces the 1D temporal convolution in MTA. The sub-convolutions formulate a hierarchical structure with residual connections between adjacent subsets. When the spatiotemporal features go through the module, the features realize multiple information exchanges with neighboring frames, and the equivalent temporal receptive field is thus increased multiple times to model long-range temporal dynamics.

The proposed ME module and MTA module are inserted into a standard ResNet block [14, 15] to build the Temporal Excitation and Aggregation (TEA) block, and the entire network is constructed by stacking multiple blocks. The obtained model is efficient: benefiting from the lightweight configurations, the FLOPs of the TEA network are controlled at a low level (only $1.06\times$ as many as 2D ResNet). The proposed model is also effective: the two components of TEA are complementary and cooperate in endowing the network with both short- and long-range temporal modeling abilities. To summarize, the main contributions of our method are three-fold:

1. The motion excitation (ME) module to integrate the short-range motion modeling with the whole spatiotemporal feature learning approach.
2. The multiple temporal aggregation (MTA) module to efficiently enlarge the temporal receptive field for long-range temporal modeling.
3. The two proposed modules are both simple, lightweight, and can be easily integrated into standard ResNet block to cooperate for effective and efficient temporal modeling.

2. Related Works

With the tremendous success of deep learning methods on image-based recognition tasks [22, 32, 35, 14, 15], some researchers started to explore the application of deep networks on video action recognition task [21, 31, 36, 47, 6, 46]. Among them, Karpathy *et al.* [21] proposed to apply a single 2D CNN model on each frame of videos independently and explored several strategies to

fuse temporal information. However, the method does not consider the motion change between frames, and the final performance is inferior to the hand-crafted feature-based algorithms. Donahue *et al.* [6] used LSTM [16] to model the temporal relation by aggregating 2D CNN features. In this approach, the feature extraction of each frame is isolated, and only high-level 2D CNN features are considered for temporal relation learning.

The existing methods usually follow two approaches to improve temporal modeling ability. The first one was based on two-stream architecture proposed by Simonyan and Zisserman [31]. The architecture contained a spatial 2D CNN that learns still feature from frames and a temporal 2D CNN that models motion information in the form of optical flow [48]. The training of the two streams is separated, and the final predictions for videos are averaged over two streams. Many following works had extended such a framework. [9, 8] explored different mid-level combination strategies to fuse the features of two streams. TSN [42] proposed the sparse sampling strategy to capture long-range video clips. All these methods require additional computation and storage costs to deal with optical flow. Moreover, the interactions between different frames and the two modalities are limited, which usually occur at late layers only. In contrast, our proposed method discards optical flow extraction and learns approximate feature-level motion representations by calculating temporal differences. The motion encoding can be integrated with the learning of spatiotemporal features and utilized to discover and enhance their motion-sensitive ingredients.

The most recent work STM [20] also attempted to model feature-level motion features and inserts motion modeling into spatiotemporal feature learning. Our method differs from STM in that STM directly *adds* the spatiotemporal features and motion encoding together. In contrast, our method utilizes motion features to *recalibrate* the features to enhance the motion pattern.

Another typical video action recognition approach is based on 3D CNNs and its (2+1)D CNN variants [36, 34, 3, 38, 44]. The first work in this line was C3D [36], which performed 3D convolutions on adjacent frames to jointly model the spatial and temporal features in a unified approach. To utilize pre-trained 2D CNNs, Carreira and Zisserman [3] proposed I3D to inflate the pre-trained 2D convolutions to 3D ones. To reduce the heavy computations of 3D CNNs, some works proposed to decompose the 3D convolution into a 2D spatial convolution and a 1D temporal convolution [34, 5, 25, 13, 29, 37] or utilize a mixup of 2D CNN and 3D CNN [38, 45, 52]. In these methods, the long-range temporal connection can be theoretically established by stacking multiple local temporal convolutions. However, after a large number of local convolution operations, the useful features from distant frames have already been weak-

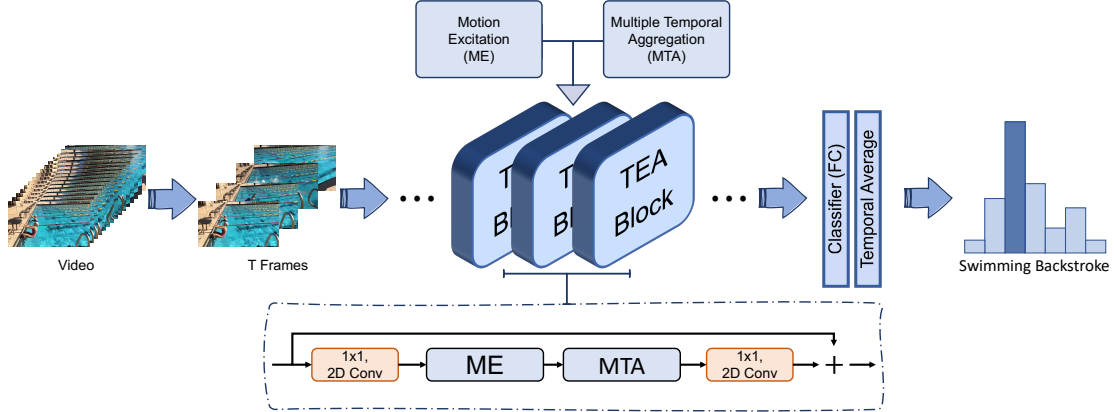


Figure 1. The framework of the proposed method for action recognition. The sparse sampling strategy [42] is adopted to sample T frames from videos. The 2D ResNet [14] is utilized as the backbone, and the ME and MTA modules are inserted into each ResNet block to form the TEA block. The simple temporal pooling is applied to average action predictions for the entire video.

ened and cannot be captured well. To address this issue, T3D [5] proposed to adopt densely connected structure [19] and combined different temporal windows [35]. Non-local module [43] and stnet [13] applied self-attention mechanism to model long-range temporal relationship. Either additional parameters or time-consuming operations accompany these attempts. Different from these works, our proposed multiple temporal aggregation module is simple and efficient without introducing extra operators.

3. Our Method

The framework of the proposed method is illustrated in Figure 1. The input videos with variable lengths are sampled using the sparse temporal sampling strategy proposed by TSN [42]. Firstly, the videos are evenly divided into T segments. Then one frame is randomly selected from each segment to form the input sequence with T frames. For spatiotemporal modeling, our model is based on 2D CNN ResNet [14] and constructed by stacking multiple Temporal Excitation and Aggregation (TEA) blocks. The TEA block contains a motion excitation (ME) module to excite motion patterns and a multiple temporal aggregation (MTA) module to establish a long-range temporal relationship. Following previous methods [42, 25], the simple temporal average pooling is utilized at the end of the model to average the predictions of all frames.

3.1. Motion Excitation (ME) Module

Motion measures the content displacements of the two successive frames and mainly reflects the actual actions. Many previous works utilize motion representations for action recognition [42, 3]. Still, most of them only consider *pixel-level* motion pattern in the form of optical flow [48] and separate the learning of motions from spatiotemporal features. Different from this, in the proposed *motion excitation* (ME) module, the motion modeling is extended from the raw *pixel-level* to a largely scoped *feature-level*,

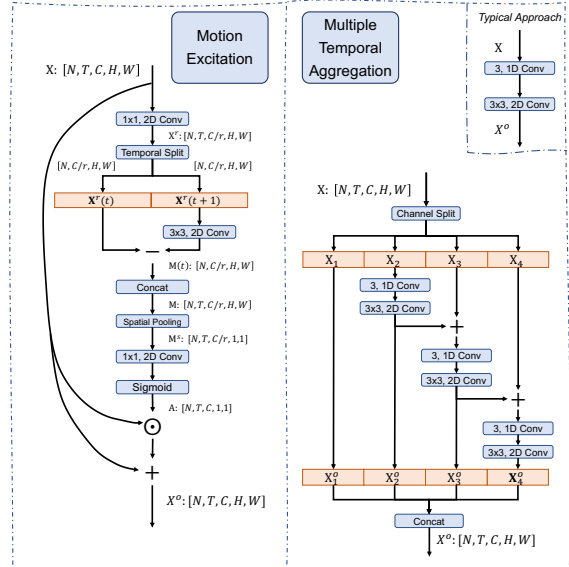


Figure 2. The implementations of the motion excitation (ME) module (left panel) and multiple temporal aggregation (MTA) module (right panel).

such that the motion modeling and spatiotemporal features learning are incorporated into a unified framework.

The architecture of the ME module is shown in the left panel of Figure 2. The shape of input spatiotemporal feature \mathbf{X} is $[N, T, C, H, W]$, where N is the batch size. T and C denote temporal dimension and feature channels, respectively. H and W correspond to spatial shape. The intuition of the proposed ME module is that, among all feature channels, different channels would capture distinct information. A portion of channels tends to model the static information related to background scenes; other channels mainly focus on dynamic motion patterns describing the temporal difference. For action recognition, it is beneficial to enable the model to discover and then enhance these motion-sensitive channels.

Given an input feature \mathbf{X} , a 1×1 2D convolution layer is

firstly adopted to reduce feature channels for efficiency.

$$\mathbf{X}^r = \text{conv}_{red} * \mathbf{X}, \quad \mathbf{X}^r \in \mathbb{R}^{N \times T \times C/r \times H \times W} \quad (1)$$

where \mathbf{X}^r denotes the channel-reduced feature. $*$ indicates the convolution operation. $r = 16$ is the reduction ratio.

The feature-level motion representations at time step t is approximately considered as the difference between the two adjacent frames, $\mathbf{X}^r(t)$ and $\mathbf{X}^r(t+1)$. Instead of directly subtracting the original features, we propose to perform the *channel-wise* transformation on features first and then utilize the transformed feature to calculate motions. Formally,

$$\mathbf{M}(t) = \text{conv}_{trans} * \mathbf{X}^r(t+1) - \mathbf{X}^r(t), 1 \leq t \leq T-1, \quad (2)$$

where $\mathbf{M}(t) \in \mathcal{R}^{N \times C/r \times H \times W}$ is the motion feature at time t . conv_{trans} is a 3×3 2D channel-wise convolution layer performing transformation for each channel.

We denote the motion feature at the end of time steps as zero, *i.e.*, $\mathbf{M}(T) = 0$, and construct the final motion matrix \mathbf{M} by concatenating all the motion features $[\mathbf{M}(1), \dots, \mathbf{M}(T)]$. Then a global average pooling layer is utilized to summarize the spatial information since our goal is to excite the motion-sensitive channels where the detailed spatial layouts are of no great importance:

$$\mathbf{M}^s = \text{Pool}(\mathbf{M}), \quad \mathbf{M}^s \in \mathbb{R}^{N \times T \times C/r \times 1 \times 1}. \quad (3)$$

Another 1×1 2D convolution layer conv_{exp} is utilized to expand the channel dimension of motion features to the original channel dimension C , and the motion-attentive weights \mathbf{A} can be obtained by using the sigmoid function.

$$\mathbf{A} = 2\delta(\text{conv}_{exp} * \mathbf{M}^s) - 1, \quad \mathbf{A} \in \mathbb{R}^{N \times T \times C \times 1 \times 1}, \quad (4)$$

where δ indicates the sigmoid function.

Finally, the goal of the module is to excite the motion-sensitive channels; thus, a simple way is to conduct channel-wise multiplication between the input feature \mathbf{X} and attentive weight \mathbf{A} . However, such an approach will suppress the static background scene information, which is also beneficial for action recognition. To address this issue, in the proposed motion-based excitation module, we propose to adopt a *residual* connection to enhance motion information meanwhile preserve scene information.

$$\mathbf{X}^o = \mathbf{X} + \mathbf{X} \odot \mathbf{A}, \quad \mathbf{X}^o \in \mathbb{R}^{N \times T \times C \times H \times W}, \quad (5)$$

where \mathbf{X}^o is the output of the proposed module, in which the motion pattern has been excited and enhanced. \odot indicates the channel-wise multiplication.

3.1.1 Discussion with SENet

The excitation scheme is firstly proposed by SENet [18, 17] for image recognition tasks. We want to highlight our

differences with SENet. 1) SENet is designed for image-based tasks. When SENet is applied to spatiotemporal features, it processes each frame of videos independently without considering temporal information. 2) SENet is a kind of self-gating mechanism [40], and the obtained modulation weights are utilized to enhance the informative channels of feature \mathbf{X} . While our module aims to enhance the motion-sensitive ingredients of the feature. 3) The useless channels will be completely suppressed in SENet, but the static background information can be preserved in our module by introducing a residual connection.

3.2. Multiple Temporal Aggregation (MTA) Module

Previous action recognition methods [36, 34] typically adopt the local temporal convolution to process neighboring frames at a time, and the long-range temporal structure can be modeled only in deep networks with a large number of stacked local operations. It is an ineffective approach since the optimization message delivered from distant frames has been dramatically weakened and cannot be well handled. To address this issue, we propose the *multiple temporal aggregation* (MTA) module for effective long-range temporal modeling. The MTA module is inspired by Res2Net [10], in which the spatiotemporal features and corresponding local convolution layers are split into a group of subsets. This approach is efficient since it does not introduce additional parameters and time-consuming operations. In the module, the subsets are formulated as a hierarchical residual architecture such that a serial of sub-convolutions are successively applied to the features and could accordingly enlarge the equivalent receptive field of the temporal dimension.

As shown in the upper-right corner of Figure 2, given an input feature \mathbf{X} , a typical approach is to process it with a single local temporal convolution and another spatial convolution. Different from this, we split the feature into four fragments along the channel dimension, and the shape of each fragment thus becomes $[N, T, C/4, H, W]$. The local convolutions are also divided into multiple sub ones. The last three fragments are sequentially processed with one *channel-wise* temporal sub-convolution layer and another spatial sub-convolution layer. Each of them only has 1/4 parameters as original ones. Moreover, the residual connection is added between the two adjacent fragments, which transforms the module from a parallel architecture to a hierarchical cascade one. Formally¹,

$$\begin{aligned} \mathbf{X}_i^o &= \mathbf{X}_i, & i &= 1, \\ \mathbf{X}_i^o &= \text{conv}_{spa} * (\text{conv}_{temp} * \mathbf{X}_i), & i &= 2, \\ \mathbf{X}_i^o &= \text{conv}_{spa} * (\text{conv}_{temp} * (\mathbf{X}_i + \mathbf{X}_{i-1}^o)), & i &= 3, 4, \end{aligned} \quad (6)$$

¹The necessary reshape and permutation operations are ignored for simplicity. In fact, to conduct 1D temporal convolution on input feature \mathbf{X} , it requires to be reshaped from $[N, T, C, H, W]$ to $[NHW, C, T]$.

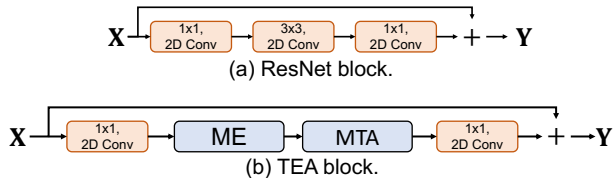


Figure 3. The motion excitation (ME) module is placed after the first 1×1 convolution layer. The multiple temporal aggregation (MTA) module is utilized to replace the 3×3 convolution layer.

where $\mathbf{X}_i^o \in \mathbb{R}^{N \times T \times C/4 \times H \times W}$ is the output of i -th fragment. conv_{temp} denotes the 1D channel-wise temporal sub-convolution whose kernel size is 3 and conv_{spa} indicates the 3×3 2D spatial sub-convolution.

In this module, the different fragments have different receptive fields. For example, the output of the first fragment \mathbf{X}_1^o is the same as input fragment \mathbf{X}_1 ; thus, its receptive field is $1 \times 1 \times 1$. By aggregating information from former fragments in series, the equivalent receptive field of the last fragment \mathbf{X}_4^o has been enlarged three times. Finally, a simple concatenation strategy is adopted to combine multiple outputs.

$$\mathbf{X}^o = [\mathbf{X}_1^o; \mathbf{X}_2^o; \mathbf{X}_3^o; \mathbf{X}_4^o], \quad \mathbf{X}^o \in \mathbb{R}^{N \times T \times C \times H \times W} \quad (7)$$

The obtained output feature \mathbf{X}^o involves spatiotemporal representations capturing different temporal ranges. It is superior to the local temporal representations obtained by using a single local convolution in typical approaches.

3.3. Integration with ResNet Block

Finally, we describe how to integrate the proposed modules into standard ResNet block [14] to construct our temporal excitation and aggregation (TEA) block. The approach is illustrated in Figure 3. For computational efficiency, the motion excitation (ME) module is integrated into the residual path after the bottleneck layer (the first 1×1 Conv layer). The multiple temporal aggregation (MTA) module is utilized to replace the original 3×3 Conv layer in the residual path. The action recognition network can be constructed by stacking the TEA blocks.

4. Experiments

4.1. Datasets

The proposed approach is evaluated on two large-scale action recognition datasets, Something-Something V1 [12] and Kinetic400 [3], and other two small-scale datasets, HMDB51 [23] and UCF101 [33]. As pointed in [45, 51], most of the categories in Kinetics, HMDB, and UCF can be recognized by considering the background scene information only, and the temporal understanding is not very important in most cases. While the categories of Something-Something focus on human interactions with

daily life objects, for example, “pull something” and “push something”. Classifying these interactions requires more considerations of temporal information. Thus the proposed method is mainly evaluated on Something-Something since our goal is to improve the temporal modeling ability.

Kinetics contains 400 categories and provides download URL links for $\sim 240k$ training videos and $\sim 20k$ validation videos. In our experiments, we successfully collect 223,127 training videos and 18,153 validation videos, because a small fraction of the URLs (around 10%) is no longer valid. For the Kinetics dataset, the methods are learned on the training set and evaluated on the validation set. HMDB contains 51 classes and 6,766 videos, while UCF includes 101 categories with 13,320 videos. For these two datasets, we follow TSN [42] to utilize three different training/testing splits for evaluation, and the average results are reported.

Something-Something V1 includes 174 categories with 86,017 training videos, 11,522 validation videos, and 10,960 test videos. All of them have been split into individual frames at the same rate, and the extracted frames are also publicly available. The methods are learned on the training set and measured on the validation set and test set.

4.2. Implementation Details

We utilize 2D ResNet-50 as the backbone and replace each ResNet block with the TEA block from conv2 to conv5. The sparse sampling strategy [42] is utilized to extract T frames from the video clips ($T = 8$ or 16 in our experiments). During training, random scaling and corner cropping are utilized for data augmentation, and the cropped region is resized to 224×224 for each frame².

During the test, two evaluation protocols are considered to trade-off accuracy and speed. 1) *efficient protocol* (center crop $\times 1$ clip), in which 1 clip with T frames is sampled from the video. Each frame is resized to 256×256 , and a central region of size 224×224 is cropped for action prediction. 2) *accuracy protocol* (full resolution $\times 10$ clips), in which 10 different clips are randomly sampled from the video, and the final prediction is obtained by averaging all clips’ scores. For each frame in a video clip, we follow the strategy proposed by [43] and resize the shorter size to 256 with maintaining the aspect ratio. Then 3 crops of 256×256 that cover the full-frame are sampled for action prediction.

4.3. Experimental Results

4.3.1 Ablation Study

In this section, we first conduct several ablation experiments to testify the effectiveness of different components in our proposed TEA block. Without loss of generality, the models are trained with 8 frames on the Something-Something V1 training set and evaluated on the validation set. Six

²More training details can be found in supplementary materials.

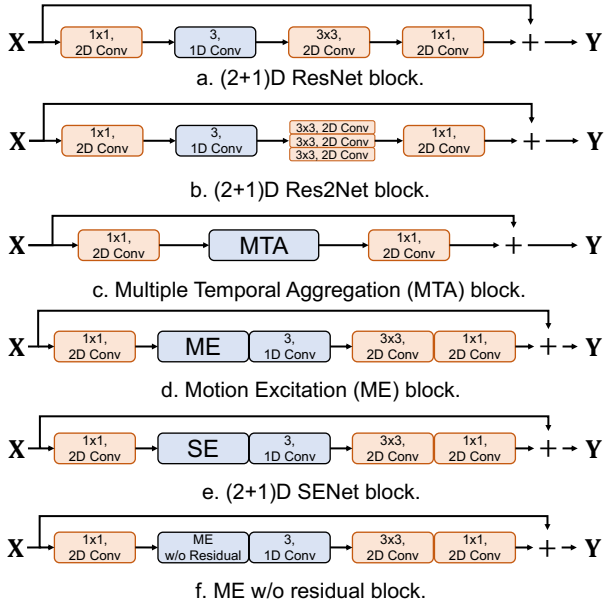


Figure 4. The altered blocks of different baselines based on standard ResNet block [14].

baseline networks are considered for comparison, and their corresponding blocks are illustrated in Figure 4. The comparison results, including the classification accuracies and inference protocols, are shown in Table 1.

- **(2+1)D ResNet.** In the residual branch of the standard ResNet block, a 1D channel-wise temporal convolution is inserted after the first 2D spatial convolution.
- **(2+1)D Res2Net.** The channel-wise temporal convolution is integrated into Res2Net block [10]. In Res2Net, the 3×3 spatial convolution of ResNet block is deformed to be a group of sub-convolutions.
- **Multiple Temporal Aggregation (MTA).** The motion excitation module is removed from the proposed TEA network.
- **Motion Excitation (ME).** Compared with the (2+1)D ResNet baseline, the proposed motion excitation module is added to the residual path.
- **(2+1)D SENet.** The SE block [18, 17] replaces the motion excitation module in the ME baseline. The SE block utilizes two fully connected layers to produce modulation weights from original features, and then apply the obtained weights to rescale the features.
- **ME w/o Residual.** The residual connection is removed from the ME baseline. Thus the output feature is obtained by directly multiplying the input feature with the motion-sensitive weights, *i.e.*, $\mathbf{X}^o = \mathbf{X} \odot \mathbf{A}$.

Table 1. Comparison results on Something-Something.

Method	Frames \times Crops \times Clips	Val Top-1 (%)	Val Top-5 (%)
(2+1)D ResNet (a) ¹	$8 \times 1 \times 1$	46.0	75.3
(2+1)D Res2Net (b) ¹	$8 \times 1 \times 1$	46.2	75.5
MTA (c) ¹	$8 \times 1 \times 1$	47.5	76.4
TEA	$8 \times 1 \times 1$	48.9	78.1
(2+1)D ResNet (a) ¹	$8 \times 1 \times 1$	46.0	75.3
(2+1)D SENet (e) ¹	$8 \times 1 \times 1$	46.5	75.6
ME w/o Residual (f) ¹	$8 \times 1 \times 1$	47.2	76.1
STM [20] ²	$8 \times 1 \times 1$	47.5	-
ME (d) ¹	$8 \times 1 \times 1$	48.4	77.5
TEA	$8 \times 1 \times 1$	48.9	78.1

1. **XX (y).** **XX** indicates the XX baseline, and y represents that the architecture of the corresponding block is the y-th one in Figure 4.

2. The result of STM using efficient inference protocol is cited from Table 9 in [20].

Effect of Multiple Temporal Aggregation. Firstly, it can be seen from the first compartment of Table 1 that the MTA baseline outperforms the (2+1)D ResNet baseline by a large margin (47.5% vs. 46.0%). Compared with the (2+1)D ResNet baseline, the capable long-range temporal aggregation can be constructed in the MTA module by utilizing the hierarchical structure to enlarge the equivalent receptive field of the temporal dimension in each block, which results in the performance improvements.

Moreover, considering the proposed MTA module enlarges both spatial and temporal receptive fields, it is thus necessary to ascertain the independent impact of the two aspects. To this end, we then compare the (2+1)D ResNet baseline with the (2+1)D Res2Net baseline. In (2+1)D Res2Net, the group of sub-convolutions is applied to spatial dimension only, and the equivalent receptive field of temporal dimension is unchanged in this model. We can see that the accuracies of the two baselines are similar and both inferior to that of MTA (46.0%/46.2% vs. 47.5%). It proves that exploring complicated spatial structures and sophisticated spatial representations have, to some extent, limit impacts on the action recognition task. The key to improving the performance of action recognition is capable and reliable temporal modeling ability.

Effect of Motion Modeling. To testify the effectiveness of the motion modeling for action recognition, we compare the ME baseline with the (2+1)D ResNet baseline. In the second compartment of Table 1, we can see that the action recognition performance is significantly increased by considering the motion encoding (48.1% vs. 46.0%). The discovery of motion-sensitive features will force the networks to focus on dynamic information that reflects the actual actions.

To prove that such improvement is not brought by introducing extra parameters and soft attention mechanisms,

Table 2. Comparison results of TEA with other state-of-the-art methods on Something-Something V1.

Method	Backbone	Frames×Crops×Clips	FLOPs	Pre-train	Val Top-1 (%)	Val Top-5 (%)	Test Top-1 (%)
3D/(2D+3D) CNNs:							
I3D-RGB [44]	3D ResNet50		153G×3×2	ImgNet	41.6	72.2	-
NL I3D-RGB [44]	3D ResNet50	32×3×2	168G×3×2	+	44.4	76.0	-
NL I3D+GCN-RGB [44]	3D ResNet50+GCN		303G×3×2	K400	46.1	76.8	45.0
2D/(2+1)D CNNs:							
TSN-RGB [42]	BNInception	8×1×1	16G×1×1	ImgNet	19.5	-	-
TSN-RGB [42]	ResNet50		33G×1×1		19.7	-	-
STM-RGB [20]	ResNet50	8×3×10	33G×3×10	ImgNet	49.2	79.3	-
STM-RGB [20]		16×3×10	67G×3×10		50.7	80.4	43.1
TSM-RGB [25]		8×1×1	33G×1×1	ImgNet	43.4	73.2	-
TSM-RGB [25]		16×1×1	65G×1×1		44.8	74.5	-
TSM _{en} -RGB [25]	ResNet50	8 + 16	33G + 65G	+	46.8	76.1	-
TSM-(RGB+Flow) [25]		16 + 16	N/A ²	K400	50.2	79.5	47.0
TEA (Ours)		8×1×1	35G×1×1		48.9	78.1	-
TEA (Ours)	ResNet50	8×3×10	35G×3×10	ImgNet	51.7	80.5	45.3
TEA (Ours)		16×1×1	70G×1×1		51.9	80.3	-
TEA (Ours)		16×3×10	70G×3×10		52.3	81.9	46.6

1. “ImgNet” denotes ImageNet dataset [4, 30] and “K400” indicates Kinetics400 datasets [3].

2. “N/A” represents that the FLOPs cannot be accurately measured because of extracting optical flow.

we then compare the (2+1)D SENet baseline with the (2+1)D ResNet baseline. (2+1)D SENet adds the SE block at the start of the trunk path, aiming to excite the informative feature channels. However, the SE block is applied to each frame of videos independently, and the temporal information is not considered in this approach. Thus, the performance of the (2+1)D SENet baseline is similar to the (2+1)D ResNet baseline (46.5% vs. 46.0%). The improvement is quite limited.

Finally, we explore several designs for motion modeling. We first compare the ME baseline with the ME w/o Residual baseline. It can be seen that the performance decreases from 48.1% to 47.2% without residual connections since the static information related background scenes will be eliminated in ME w/o Residual. It proves that the scene information is also beneficial for action recognition, and the residual connection is necessary for the motion excitation module. Then we compare the ME baseline with STM [20]. We can see that ME attains higher accuracy than STM (48.4% vs. 47.5%), which verifies the excitation mechanism utilized in the proposed method is superior to the simple *add* approach used in STM. When additionally considering the long-range temporal relationship by introducing the MTA module, the accuracy of our method (TEA) can be further improved to 48.9%.

4.3.2 Comparisons with the State-of-the-arts

In this section, we first compare TEA with the existing state-of-the-art action recognition methods on Something-

Something V1 and Kinetics400. The comprehensive statistics, including the classification results, inference protocols, and the corresponding FLOPs, are shown in Table 2 and 3.

In both tables, the first compartment contains the methods based on 3D CNNs or the mixup of 2D and 3D CNNs, and the methods in the second compartment are all based on 2D or (2+1)D CNNs. Due to the high computation costs of 3D CNNs, the FLOPs of methods in the first compartment are typically higher than others. Among all existing methods, the most efficient ones are TSN_{8f} [42] and TSM_{8f} [25] with only 33G FLOPs. Compared with these methods, the FLOPs of our proposed TEA network slightly increases to 35G (1.06×), but the performance is increased by a big margin, a relative improvement of 5.4% (48.8% vs. 43.4%).

The superiority of our TEA on Something-Something is quite impressive. It confirms the remarkable ability of TEA for temporal modeling. Using efficient inference protocol (center crop×1 clip) and 8 input frames, the proposed TEA obtains 48.8%, which significantly outperforms TSN and TSM with similar FLOPs (19.7%/43.4%). This results even exceeds the ensemble result of TSM, which combines the two models using 8 and 16 frames, respectively (TSM_{En}, 46.8%). When utilizing 16 frames as input and applying a more laborious accuracy evaluation protocol (full resolution×10 clips), the FLOPs of our method increase to ~2000G, which is similar to NL I3D+GCN [44]. But the proposed method significantly surpasses NL I3D+GCN and all other existing methods (52.3% vs. 46.1%) on the validation set. Our performance on the test set

Table 3. Comparison results of TEA with other state-of-the-art methods on Kinetics400 validation set.

Method	Backbone	Frames×Crops×Clips	FLOPs	Pre-train	Top-1 (%)	Top-5 (%)
3D/(2D+3D) CNNs:						
I3D-RGB [3]	Inception V1	$64 \times N/A \times N/A^1$	$108G \times N/A \times N/A$	ImgNet	72.1	90.3
I3D-RGB [3]			$108G \times N/A \times N/A$	None	67.5	87.2
ECO-RGB _{En} [52]	BNIncep+3D Res18	$92 \times 1 \times 1$	$267G \times 1 \times 1$	None	70.0	-
NL I3D-RGB [44]	3D ResNet101	$32 \times 6 \times 10$	$359G \times 6 \times 10$	ImgNet	77.7	93.3
NL SlowFast [7]	3D ResNet101	$(16+8) \times 3 \times 10$	$234G \times 3 \times 10$	None	79.8	93.9
2D/(2+1)D CNNs:						
TSN-RGB [42]	BNInception Inception v3	$25 \times 10 \times 1$	$53G \times 10 \times 1$	ImgNet	69.1	88.7
TSN-RGB [42]			$80G \times 10 \times 1$		72.5	90.2
R(2+1)D [38]	ResNet-34	$32 \times 1 \times 10$	$152G \times 1 \times 10$	None	72.0	90.0
STM-RGB [20]	ResNet50	$16 \times 3 \times 10$	$67G \times 3 \times 10$	ImgNet	73.7	91.6
TSM-RGB [25]	ResNet50	$8 \times 3 \times 10$	$33G \times 3 \times 10$	ImgNet	74.1	-
TSM-RGB [25]		$16 \times 3 \times 10$	$65G \times 3 \times 10$		74.7	-
TEA (Ours)	ResNet50	$8 \times 1 \times 1$	$35G \times 1 \times 1$	ImgNet	72.5	90.4
TEA (Ours)		$8 \times 3 \times 10$	$35G \times 3 \times 10$		75.0	91.8
TEA (Ours)		$16 \times 1 \times 1$	$70G \times 1 \times 1$		74.0	91.3
TEA (Ours)		$16 \times 3 \times 10$	$70G \times 3 \times 10$		76.1	92.5

1. “ImgNet” denotes ImageNet dataset [4, 30] and “None” indicates training models from scratch.
2. “N/A” represents that the authors do not report the inference protocol in their paper.

(46.6%) also outperforms most of the existing methods. Moreover, we do not require additional COCO images [26] to pre-train an object detector as in [44]. When compared with the methods utilizing both RGB and optical flow modalities, *i.e.*, ECO_{En}-(RGB+Flow) [52] (49.5%) and TSM-(RGB+Flow) [25] (50.2%), the obtained result (52.3%) also shows substantial improvements.

On Kinetics400, the performance of our method (76.1%) is inferior to that of SlowFast [7] (79.8%). However, the SlowFast networks adopt the deeper networks (ResNet101) based on 3D CNNs and utilize time-consuming non-local [43] operations. When comparing methods with similar efficiency, such as TSM [25] and STM [20], TEA obtains better performance. When adopting 8 frames as input, TEA gains $\sim 1\%$ higher accuracy than TSM (75.0% *vs.* 74.1%). While utilizing 16 input frames, our TEA method outperforms both TSM_{16f} and STM_{16f} with a large margin (76.1% *vs.* 74.7%/73.7%).

Finally, we report comparison results on HMDB51 and UCF101 in Table 4. Our method achieves 73.3% on HMDB51 and 96.9% on UCF101 with the accuracy inference protocol. The performance of our model (TEA_{16f}) outperforms most of the existing methods except for I3D [3]. I3D is based on 3D CNNs and additional input modality; thus, its computational FLOPs will be far more than ours.

5. Conclusion

In this paper, we propose the Temporal Excitation and Aggregation (TEA) block, including the motion excitation

Table 4. Comparison results on HMDB51 and UCF101.

Method	Backbone	HMDB51 MCA (%) ¹	UCF101 MCA (%) ¹
I3D-(RGB+Flow) [3]	3D Inception	80.7	98.0
TSN-(RGB+Flow) [42]	BNInception	68.5	94.0
StNet [13]	ResNet50	-	93.5
TSM ²	ResNet50	70.7	94.5
STM [20]	ResNet50	72.2	96.2
TEA (Ours)	ResNet50	73.3	96.9

1. MCA denotes mean class accuracy.
2. TSM does not report MCA results, and the listed results are cited from STM [20].

(ME) module and the multiple temporal aggregation (MTA) module for both short- and long-range temporal modeling. Specifically, the ME module could insert the motion encoding into the spatiotemporal feature learning approach and enhance the motion pattern in spatiotemporal features. In the MTA module, the reliable long-range temporal relationship can be established by deforming the local convolutions into a group of sub-convolutions to enlarge the equivalent temporal receptive field. The two proposed modules are integrated into the standard ResNet block and cooperate for capable temporal modeling.

6. Acknowledgement

This work is supported by the Video Understanding Middle Platform of the Platform and Content Group (PCG) at Tencent. The authors would like to thank Wei Shen for his helpful suggestions.

References

- [1] Hakan Bilen, Basura Fernando, Efstratios Gavves, and Andrea Vedaldi. Action recognition with dynamic image networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2799–2813, 2017.
- [2] Hakan Bilen, Basura Fernando, Efstratios Gavves, Andrea Vedaldi, and Stephen Gould. Dynamic image networks for action recognition. In *CVPR*, pages 3034–3042, 2016.
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.
- [5] Ali Diba, Mohsen Fayyaz, Vivek Sharma, Amir Hossein Karami, Mohammad Mahdi Arzani, Rahman Yousefzadeh, and Luc Van Gool. Temporal 3d convnets: New architecture and transfer learning for video classification. *arXiv preprint arXiv:1711.08200*, 2017.
- [6] 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*, pages 2625–2634, 2015.
- [7] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6202–6211, 2019.
- [8] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Spatiotemporal multiplier networks for video action recognition. In *CVPR*, pages 4768–4777, 2017.
- [9] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, pages 1933–1941, 2016.
- [10] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *arXiv preprint arXiv:1904.01169*, 2019.
- [11] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *CVPR*, pages 971–980, 2017.
- [12] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, pages 5843–5851, 2017.
- [13] 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*, pages 8401–8408, 2019.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*, pages 630–645, 2016.
- [16] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [17] J Hu, L Shen, S Albanie, G Sun, and E Wu. Squeeze-and-excitation networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [18] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141, 2018.
- [19] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, pages 4700–4708, 2017.
- [20] Boyuan Jiang, MengMeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. Stm: Spatiotemporal and motion encoding for action recognition. In *ICCV*, pages 2000–2009, 2019.
- [21] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, pages 1725–1732, 2014.
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [23] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, pages 2556–2563, 2011.
- [24] Yanghao Li, Sijie Song, Yuqi Li, and Jiaying Liu. Temporal bilinear networks for video action recognition. In *AAAI*, pages 8674–8681, 2019.
- [25] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019.
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014.
- [27] Joe Yue-Hei Ng and Larry S Davis. Temporal difference networks for video action recognition. In *WACV*, pages 1587–1596, 2018.
- [28] Oleksandra Poquet, Lisa Lim, Negin Mirriahi, and Shane Dawson. Video and learning: a systematic review (2007–2017). In *ICLAK*, pages 151–160. ACM, 2018.
- [29] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *CVPR*, pages 12056–12065, 2019.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [31] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, pages 568–576, 2014.
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

- [33] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [34] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *ICCV*, pages 4597–4605, 2015.
- [35] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, pages 1–9, 2015.
- [36] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [37] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. *arXiv preprint arXiv:1904.02811*, 2019.
- [38] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018.
- [39] Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1510–1517, 2017.
- [40] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, pages 3156–3164, 2017.
- [41] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification. In *CVPR*, pages 1430–1439, 2018.
- [42] 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*, pages 20–36, 2016.
- [43] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018.
- [44] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, pages 399–417, 2018.
- [45] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, pages 305–321, 2018.
- [46] Li Yao, Atousa Torabi, Kyunghyun Cho, Nicolas Ballas, Christopher Pal, Hugo Larochelle, and Aaron Courville. Describing videos by exploiting temporal structure. In *ICCV*, pages 4507–4515, 2015.
- [47] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, pages 4694–4702, 2015.
- [48] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Joint Pattern Recognition Symposium*, pages 214–223, 2007.
- [49] Yue Zhao, Yuanjun Xiong, and Dahua Lin. Recognize actions by disentangling components of dynamics. In *CVPR*, pages 6566–6575, 2018.
- [50] Yue Zhao, Yuanjun Xiong, and Dahua Lin. Trajectory convolution for action recognition. In *NIPS*, pages 2204–2215, 2018.
- [51] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, pages 803–818, 2018.
- [52] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. Eco: Efficient convolutional network for online video understanding. In *ECCV*, pages 695–712, 2018.