

TSM: Temporal Shift Module for Efficient Video Understanding

Ji Lin
MIT

jilini@mit.edu

Chuang Gan
MIT-IBM Watson AI Lab

ganchuang@csail.mit.edu

Song Han
MIT

songhan@mit.edu

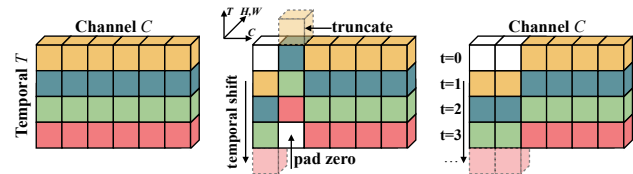
Abstract

The explosive growth in video streaming gives rise to challenges on performing video understanding at high accuracy and low computation cost. Conventional 2D CNNs are computationally cheap but cannot capture temporal relationships; 3D CNN based methods can achieve good performance but are computationally intensive, making it expensive to deploy. In this paper, we propose a generic and effective Temporal Shift Module (TSM) that enjoys both high efficiency and high performance. Specifically, it can achieve the performance of 3D CNN but maintain 2D CNN's complexity. TSM shifts part of the channels along the temporal dimension; thus facilitate information exchanged among neighboring frames. It can be inserted into 2D CNNs to achieve temporal modeling at zero computation and zero parameters. We also extended TSM to online setting, which enables real-time low-latency online video recognition and video object detection. TSM is accurate and efficient: it ranks the first place on the Something-Something leaderboard upon publication; on Jetson Nano and Galaxy Note8, it achieves a low latency of 13ms and 35ms for online video recognition. The code is available at: <https://github.com/mit-han-lab/temporal-shift-module>.

1. Introduction

Hardware-efficient video understanding is an important step towards real-world deployment, both on the cloud and on the edge. For example, there are over 10^5 hours of videos uploaded to YouTube every day to be processed for recommendation and ads ranking; tera-bytes of sensitive videos in hospitals need to be processed locally on edge devices to protect privacy. All these industry applications require both accurate and efficient video understanding.

Deep learning has become the standard for video understanding over the years [45, 48, 4, 49, 61, 53, 58]. One key difference between video recognition and image recognition is the need for *temporal modeling*. For example, to distinguish between opening and closing a box, reversing the order will give opposite results, so temporal modeling is critical.



(a) The original tensor without shift. (b) Offline temporal shift (bi-directional). (c) Online temporal shift (uni-directional).

Figure 1. **Temporal Shift Module (TSM)** performs efficient temporal modeling by moving the feature map along the temporal dimension. It is computationally free on top of a 2D convolution, but achieves strong temporal modeling ability. TSM efficiently supports both **offline** and **online** video recognition. Bi-directional TSM mingles both past and future frames with the current frame, which is suitable for high-throughput offline video recognition. Uni-directional TSM mingles only the past frame with the current frame, which is suitable for low-latency online video recognition.

Existing efficient video understanding approaches directly use 2D CNN [24, 39, 48, 58]. However, 2D CNN on individual frames cannot well model the temporal information. 3D CNNs [45, 4] can jointly learn spatial and temporal features but the computation cost is large, making the deployment on edge devices difficult; it cannot be applied to real-time online video recognition. There are works to trade off between temporal modeling and computation, such as post-hoc fusion [13, 9, 58, 7] and mid-level temporal fusion [61, 53, 46]. Such methods sacrifice the low-level temporal modeling for efficiency, but much of the useful information is lost during the feature extraction before the temporal fusion happens.

In this paper, we propose a new perspective for efficient temporal modeling in video understanding by proposing a novel Temporal Shift Module (TSM). Concretely, an activation in a video model can be represented as $A \in \mathbb{R}^{N \times C \times T \times H \times W}$, where N is the batch size, C is the number of channels, T is the temporal dimension, H and W are the spatial resolutions. Traditional 2D CNNs operate independently over the dimension T ; thus no temporal modeling takes effects (Figure 1a). In contrast, our Temporal Shift Module (TSM) shifts the channels along the temporal dimension, both forward and backward. As shown in Figure 1b, the information from neighboring frames is mingled with the current frame after shifting. Our intuition is: the convo-

lution operation consists of *shift* and *multiply-accumulate*. We *shift* in the time dimension by ± 1 and fold the *multiply-accumulate* from time dimension to channel dimension. For real-time online video understanding, future frames can't get shifted to the present, so we use a uni-directional TSM (Figure 1c) to perform online video understanding.

Despite the zero-computation nature of the shift operation, we empirically find that simply adopting the spatial shift strategy [51] used in image classifications introduces two major issues for video understanding: (1) it is *not efficient*: shift operation is conceptually zero FLOP but incurs data movement. The additional cost of data movement is non-negligible and will result in latency increase. This phenomenon has been exacerbated in the video networks since they usually have a large memory consumption (5D activation). (2) It is *not accurate*: shifting too many channels in a network will significantly hurt the spatial modeling ability and result in performance degradation. To tackle the problems, we make two technical contributions. (1) We use a *temporal partial shift* strategy: instead of shifting all the channels, we shift only a small portion of the channels for efficient temporal fusion. Such strategy significantly cuts down the data movement cost (Figure 2a). (2) We insert TSM inside *residual branch* rather than outside so that the activation of the current frame is preserved, which does not harm the spatial feature learning capability of the 2D CNN backbone.

The contributions of our paper are summarized as follows:

- We provide a new perspective for efficient video model design by temporal shift, which is computationally free but has strong spatio-temporal modeling ability.
- We observed that naive shift cannot achieve high efficiency or high performance. We then proposed two technical modifications *partial shift* and *residual shift* to realize a high efficiency model design.
- We propose *bi-directional TSM* for *offline* video understanding that achieves state-of-the-art performance. It ranks the first on Something-Something leaderboard upon publication.
- We propose *uni-directional TSM* for *online* real-time video recognition with strong temporal modeling capacity at low latency on edge devices.

2. Related Work

2.1. Deep Video Recognition

2D CNN. Using the 2D CNN is a straightforward way to conduct video recognition [24, 39, 48, 11, 8, 9, 2]. For example, Simonyan *et al.* [39] designed a two-stream CNN for RGB input (spatial stream) and optical flow [55] input

(temporal stream) respectively. Temporal Segment Networks (TSN) [48] extracted averaged features from strided sampled frames. Such methods are more efficient compared to 3D counterparts but cannot infer the temporal order or more complicated temporal relationships.

3D CNN. 3D convolutional neural networks can jointly learn spatio-temporal features. Tran *et al.* [45] proposed a 3D CNN based on VGG models, named C3D, to learn spatio-temporal features from a frame sequence. Carreira and Zisserman [4] proposed to inflate all the 2D convolution filters in an Inception V1 model [43] into 3D convolutions. However, 3D CNNs are computationally heavy, making the deployment difficult. They also have more parameters than 2D counterparts, thus are more prone to over-fitting. On the other hand, our TSM has the same spatio-temporal modeling ability as 3D CNN while enjoying the same computation and parameters as the 2D CNNs.

Trade-offs. There have been attempts to trade off expressiveness and computation costs. Lee *et al.* [27] proposed a motion filter to generate spatio-temporal features from 2D CNN. Tran *et al.* [46] and Xie *et al.* [53] proposed to study mixed 2D and 3D networks, either first using 3D and later 2D (bottom-heavy) or first 2D and later 3D (top-heavy) architecture. ECO [61] also uses a similar top-heavy architecture to achieve a very efficient framework. Another way to save computation is to decompose the 3D convolution into a 2D spatial convolution and a 1D temporal convolution [46, 33, 42]. For mixed 2D-3D CNNs, they still need to remove low-level temporal modeling or high-level temporal modeling. Compared to decomposed convolutions, our method completely removes the computation cost of temporal modeling and enjoys better hardware efficiency.

2.2. Temporal Modeling

A direct way for temporal modeling is to use 3D CNN based methods as discussed above. Wang *et al.* [49] proposed a spatial-temporal non-local module to capture long-range dependencies. Wang *et al.* [50] proposed to represent videos as space-time region graphs. An alternative way to model the temporal relationships is to use 2D CNN + post-hoc fusion [13, 9, 58, 7]. Some works use LSTM [19] to aggregate the 2D CNN features [54, 7, 41, 10, 12]. Attention mechanism also proves to be effective for temporal modeling [37, 28, 32]. Zhou *et al.* [58] proposed Temporal Relation Network to learn and reason about temporal dependencies. The former category is computational heavy, while the latter cannot capture the useful low-level information that is lost during feature extraction. Our method offers an efficient solution at the cost of 2D CNNs, while enabling both low-level and high-level temporal modeling, just like 3D-CNN based methods.

2.3. Efficient Neural Networks

The efficiency of 2D CNN has been extensively studied. Some works focused on designing an efficient model [21, 20, 36, 56]. Recently neural architecture search [62, 63, 31] has been introduced to find an efficient architecture automatically [44, 3]. Another way is to prune, quantize and compress an existing model for efficient deployment [16, 15, 29, 59, 18, 47]. Address shift, which is a hardware-friendly primitive, has also been exploited for compact 2D CNN design on image recognition tasks [51, 57]. Nevertheless, we observe that directly adopting the shift operation on video recognition task neither maintains efficiency nor accuracy, due to the complexity of the video data.

3. Temporal Shift Module (TSM)

We first explain the intuition behind TSM: data movement and computation can be separated in a convolution. However, we observe that such naive shift operation neither achieves high efficiency nor high performance. To tackle the problem, we propose two techniques minimizing the data movement and increasing the model capacity, which leads to the efficient TSM module.

3.1. Intuition

Let us first consider a normal convolution operation. For brevity, we used a 1-D convolution with the kernel size of 3 as an example. Suppose the weight of the convolution is $W = (w_1, w_2, w_3)$, and the input X is a 1-D vector with infinite length. The convolution operator $Y = \text{Conv}(W, X)$ can be written as: $Y_i = w_1 X_{i-1} + w_2 X_i + w_3 X_{i+1}$. We can decouple the operation of convolution into two steps: *shift* and *multiply-accumulate*: we shift the input X by $-1, 0, +1$ and multiply by w_1, w_2, w_3 respectively, which sum up to be Y . Formally, the *shift* operation is:

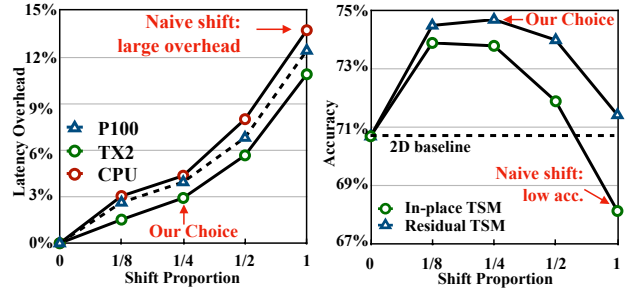
$$X_i^{-1} = X_{i-1}, \quad X_i^0 = X_i, \quad X_i^{+1} = X_{i+1} \quad (1)$$

and the *multiply-accumulate* operation is:

$$Y = w_1 X^{-1} + w_2 X^0 + w_3 X^{+1} \quad (2)$$

The first step *shift* can be conducted without any multiplication. While the second step is more computationally expensive, our Temporal Shift module *merges* the *multiply-accumulate* into the following 2D convolution, so it introduces no extra cost compared to 2D CNN based models.

The proposed Temporal Shift module is described in Figure 1. In Figure 1a, we describe a tensor with C channels and T frames. The features at different time stamps are denoted as different colors in each row. Along the temporal dimension, we shift part of the channels by -1 , another part by $+1$, leaving the rest un-shifted (Figure 1b). For online video recognition setting, we also provide an online version of TSM (Figure 1c). In the online setting, we cannot access



(a) Overhead vs. proportion. (b) Residual vs. in-place.

Figure 2. (a) Latency overhead of TSM due to data movement. (b) Residual TSM achieve better performance than in-place shift. We choose 1/4 proportion residual shift as our default setting. It achieves higher accuracy with a negligible overhead.

future frames, therefore, we only shift from past frames to future frames in a uni-directional fashion.

3.2. Naive Shift Does Not Work

Despite the simple philosophy behind the proposed module, we find that directly applying the spatial shift strategy [51] to the temporal dimension cannot provide high performance nor efficiency. To be specific, if we shift all or most of the channels, it brings two disasters: (1) **Worse efficiency due to large data movement.** The shift operation enjoys no computation, but it involves data movement. Data movement increases the memory footprint and inference latency on hardware. Worse still, such effect is exacerbated in the video understanding networks due to large activation size (5D tensor). When using the naive shift strategy shifting every map, we observe a 13.7% increase in CPU latency and 12.4% increase in GPU latency, making the overall inference slow. (2) **Performance degradation due to worse spatial modeling ability.** By shifting part of the channels to neighboring frames, the information contained in the channels is no longer accessible for the current frame, which may harm the spatial modeling ability of the 2D CNN backbone. We observe a 2.6% accuracy drop when using the naive shift implementation compared to the 2D CNN baseline (TSN).

3.3. Module Design

To tackle the two problem from naive shift implementation, we discuss two technical contributions.

Reducing Data Movement. To study the effect of data movement, we first measured the inference latency of TSM models and 2D baseline on different hardware devices. We shifted different proportion of the channels and measured the latency. We measured models with ResNet-50 backbone and 8-frame input using no shift (2D baseline), partial shift (1/8, 1/4, 1/2) and all shift (shift all the channels). The timing was measure on server GPU (NVIDIA Tesla P100), mobile GPU (NVIDIA Jetson TX2) and CPU (Intel Xeon E5-2690). We report the average latency from 1000 runs after

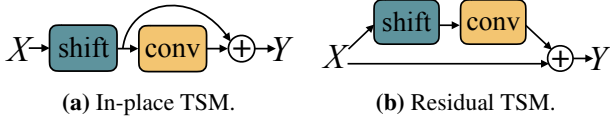


Figure 3. Residual shift is better than in-place shift. In-place shift happens before a convolution layer (or a residual block). Residual shift fuses temporal information inside a residual branch.

200 warm-up runs. We show the overhead of the shift operation as the percentage of the original 2D CNN inference time in 2a. We observe the same overhead trend for different devices. If we shift all the channels, the latency overhead takes up to **13.7%** of the inference time on CPU, which is definitely **non-negligible** during inference. On the other hand, if we only shift a small proportion of the channels, *e.g.*, 1/8, we can limit the latency overhead to **only 3%**. Therefore, we use *partial shift* strategy in our TSM implementation to significantly bring down the memory movement cost.

Keeping Spatial Feature Learning Capacity. We need to balance the model capacity for spatial feature learning and temporal feature learning. A straight-forward way to apply TSM is to insert it before each convolutional layer or residual block, as illustrated in Figure 3a. We call such implementation *in-place shift*. It harms the spatial feature learning capability of the backbone model, especially when we shift a large amount of channels, since the information stored in the shifted channels is lost for the current frame.

To address such issue, we propose a variant of the shift module. Instead of inserting it in-place, we put the TSM *inside* the residual branch in a residual block. We denote such version of shift as *residual shift* as shown in 3b. Residual shift can address the degraded spatial feature learning problem, as all the information in the original activation is still accessible after temporal shift through identity mapping.

To verify our assumption, we compared the performance of in-place shift and residual shift on Kinetics [25] dataset. We studied the experiments under different shift proportion setting. The results are shown in 2b. We can see that residual shift achieves better performance than in-place shift for all shift proportion. Even we shift all the channels to neighboring frames, due to the shortcut connection, residual shift still achieves better performance than the 2D baseline. Another finding is that the performance is related to the proportion of shifted channels: if the proportion is too small, the ability of temporal reasoning may not be enough to handle complicated temporal relationships; if too large, the spatial feature learning ability may be hurt. For residual shift, we found that the performance reaches the peak when 1/4 (1/8 for each direction) of the channels are shifted. Therefore, we use this setting for the rest of the paper.

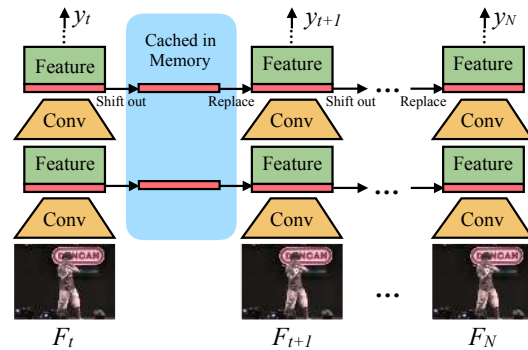


Figure 4. Uni-directional TSM for online video recognition.

4. TSM Video Network

4.1. Offline Models with Bi-directional TSM

We insert bi-directional TSM to build offline video recognition models. Given a video V , we first sample T frames F_i, F_1, \dots, F_T from the video. After frame sampling, 2D CNN baselines process each of the frames individually, and the output logits are averaged to give the final prediction. Our proposed TSM model has exactly the same parameters and computation cost as 2D model. During the inference of convolution layers, the frames are still running independently just like the 2D CNNs. The difference is that TSM is inserted for each residual block, which enables temporal information fusion at no computation. For each inserted temporal shift module, the temporal receptive field will be enlarged by 2, as if running a convolution with the kernel size of 3 along the temporal dimension. Therefore, our TSM model has a very large temporal receptive field to conduct highly complicated temporal modeling. In this paper, we used ResNet-50 [17] as the backbone unless otherwise specified.

A unique advantage of TSM is that it can easily convert any off-the-shelf 2D CNN model into a pseudo-3D model that can handle both spatial and temporal information, without adding additional computation. Thus the deployment of our framework is hardware friendly: we only need to support the operations in 2D CNNs, which are already well-optimized at both framework level (CuDNN [6], MKL-DNN, TVM [5]) and hardware level (CPU/GPU/TPU/FPGA).

4.2. Online Models with Uni-directional TSM

Video understanding from online video streams is important in real-life scenarios. Many real-time applications requires online video recognition with low latency, such as AR/VR and self-driving. In this section, we show that we can adapt TSM to achieve online video recognition while with multi-level temporal fusion.

As shown in Figure 1, offline TSM shifts part of the channels bi-directionally, which requires features from future frames to replace the features in the current frame. If we only shift the feature from previous frames to current frames, we can achieve online recognition with uni-directional TSM.

The inference graph of uni-directional TSM for online video recognition is shown in Figure 4. During inference, for each frame, we save the first 1/8 feature maps of each residual block and cache it in the memory. For the next frame, we replace the first 1/8 of the current feature maps with the cached feature maps. We use the combination of 7/8 current feature maps and 1/8 old feature maps to generate the next layer, and repeat. Using the uni-directional TSM for online video recognition shares several unique advantages:

1. Low latency inference. For each frame, we only need to replace and cache 1/8 of the features, without incurring any extra computations. Therefore, the latency of giving per-frame prediction is almost the same as the 2D CNN baseline. Existing methods like [61] use multiple frames to give one prediction, which may leads to large latency.

2. Low memory consumption. Since we only cache a small portion of the features in the memory, the memory consumption is low. For ResNet-50, we only need 0.9MB memory cache to store the intermediate feature.

3. Multi-level temporal fusion. Most of the online method only enables late temporal fusion after feature extraction like [58] or mid level temporal fusion [61], while our TSM enables all levels of temporal fusion. Through experiments (Table 2) we find that multi-level temporal fusion is very important for complex temporal modeling.

5. Experiments

We first show that TSM can significantly improve the performance of 2D CNN on video recognition while being computationally free and hardware efficient. It further demonstrated state-of-the-art performance on temporal-related datasets, arriving at a much better accuracy-computation pareto curve. TSM models achieve an order of magnitude speed up in measured GPU throughput compared to conventional I3D model from [50]. Finally, we leverage uni-directional TSM to conduct low-latency and real-time online prediction on both video recognition and object detection.

5.1. Setups

Training & Testing. We conducted experiments on video action recognition tasks. The training parameters for the Kinetics dataset are: 100 training epochs, initial learning rate 0.01 (decays by 0.1 at epoch 40&80), weight decay $1e-4$, batch size 64, and dropout 0.5. For other datasets, we scale the training epochs by half. For most of the datasets, the model is fine-tuned from ImageNet pre-trained weights; while HMDB-51 [26] and UCF-101 [40] are too small and prone to over-fitting [48], we followed the common practice [48, 49] to fine-tune from Kinetics [25] pre-trained weights and freeze the Batch Normalization [22] layers. For testing, when pursue high accuracy, we followed the common setting in [49, 50] to sample multiple clips per video (10 for Kinetics, 2 for others) and use the full resolution image

Table 1. Our method consistently outperforms 2D counterparts on multiple datasets at zero extra computation (protocol: ResNet-50 8f input, 10 clips for Kinetics, 2 for others, full-resolution).

	Dataset	Model	Acc1	Acc5	Δ Acc1
Less Temporal	Kinetics	TSN	70.6	89.2	
		Ours	74.1	91.2	+3.5
	UCF101	TSN	91.7	99.2	
		Ours	95.9	99.7	+4.2
	HMDB51	TSN	64.7	89.9	
		Ours	73.5	94.3	+8.8
More Temporal	Something V1	TSN	20.5	47.5	
		Ours	47.3	76.2	+28.0
	Something V2	TSN	30.4	61.0	
		Ours	61.7	87.4	+31.3
	Jester	TSN	83.9	99.6	
		Ours	97.0	99.9	+11.7

with shorter side 256 for evaluation, so that we can give a direct comparison; when we consider the efficiency (*e.g.*, as in Table 2), we used just 1 clip per video and the center 224×224 crop for evaluation. We keep the same protocol for the methods compared in the same table.

Model. To have an apple-to-apple comparison with the state-of-the-art method [50], we used the same backbone (ResNet-50) on the dataset (Something-Something-V1 [14]). This dataset focuses on temporal modeling. The difference is that [50] used 3D ResNet-50, while we used 2D ResNet-50 as the backbone to demonstrate efficiency.

Datasets. Kinetics dataset [25] is a large-scale action recognition dataset with 400 classes. As pointed in [58, 53], datasets like Something-Something (V1&V2) [14], Charades [38], and Jester [1] are more focused on modeling the temporal relationships, while UCF101 [40], HMDB51 [26], and Kinetics [25] are less sensitive to temporal relationships. Since TSM focuses on temporal modeling, we mainly focus on datasets with stronger temporal relationships like Something-Something. Nevertheless, we also observed strong results on the other datasets and reported it.

5.2. Improving 2D CNN Baselines

We can seamlessly inject TSM into a normal 2D CNN and improve its performance on video recognition. In this section, we demonstrate a 2D CNN baseline can significantly benefit from TSM with double-digits accuracy improvement. We chose TSN [48] as the 2D CNN baseline. We used the same training and testing protocol for TSN and our TSM. The only difference is with or without TSM.

Comparing Different Datasets. We compare the results on several action recognition datasets in Table 1. The chart is split into two parts. The upper part contains datasets Kinetics [25], UCF101 [40], HMDB51 [26], where temporal relationships are less important, while our TSM still

Table 2. Comparing TSM against other methods on Something-Something dataset (center crop, 1 clip/video unless otherwise specified).

Model	Backbone	#Frame	FLOPs/Video	#Param.	Val Top-1	Val Top-5	Test Top-1
TSN [58]	BNInception	8	16G	10.7M	19.5	-	-
TSN (our impl.)	ResNet-50	8	33G	24.3M	19.7	46.6	-
TRN-Multiscale [58]	BNInception	8	16G	18.3M	34.4	-	33.6
TRN-Multiscale (our impl.)	ResNet-50	8	33G	31.8M	38.9	68.1	-
Two-stream TRN _{RGB+Flow} [58]	BNInception	8+8	-	36.6M	42.0	-	40.7
ECO [61]	BNIncep+3D Res18	8	32G	47.5M	39.6	-	-
ECO [61]	BNIncep+3D Res18	16	64G	47.5M	41.4	-	-
ECO _{EnLite} [61]	BNIncep+3D Res18	92	267G	150M	46.4	-	42.3
ECO _{EnLite} _{RGB+Flow} [61]	BNIncep+3D Res18	92+92	-	300M	49.5	-	43.9
I3D from [50]	3D ResNet-50	32×2clip	153G ¹ ×2	28.0M	41.6	72.2	-
Non-local I3D from [50]	3D ResNet-50	32×2clip	168G ¹ ×2	35.3M	44.4	76.0	-
Non-local I3D + GCN [50]	3D ResNet-50+GCN	32×2clip	303G ² ×2	62.2M ²	46.1	76.8	45.0
TSM	ResNet-50	8	33G	24.3M	45.6	74.2	-
TSM	ResNet-50	16	65G	24.3M	47.2	77.1	46.0
TSM _{En}	ResNet-50	24	98G	48.6M	49.7	78.5	-
TSM _{RGB+Flow}	ResNet-50	16+16	-	48.6M	52.6	81.9	50.7

Table 3. TSM can consistently improve the performance over different backbones on Kinetics dataset.

	Mb-V2	R-50	RX-101	NL R-50
TSN	66.5	70.7	72.4	74.6
TSM	69.5	74.1	76.3	75.7
ΔAcc.	+3.0	+3.4	+3.9	+1.1

consistently outperforms the 2D TSN baseline at no extra computation. For the lower part, we present the results on Something-Something V1 and V2 [14] and Jester [1], which depend heavily on temporal relationships. 2D CNN baseline cannot achieve a good accuracy, but once equipped with TSM, the performance improved by double digits.

Scaling over Backbones. TSM scales well to backbones of different sizes. We show the Kinetics top-1 accuracy with MobileNet-V2 [36], ResNet-50 [17], ResNext-101 [52] and ResNet-50 + Non-local module [49] backbones in Table 3. TSM consistently improves the accuracy over different backbones, even for NL R-50, which already has temporal modeling ability.

5.3. Comparison with State-of-the-Arts

TSM not only significantly improves the 2D baseline but also outperforms state-of-the-art methods, which heavily rely on 3D convolutions. We compared the performance of our TSM model with state-of-the-art methods on both Something-Something V1&V2 because these two datasets focus on temporal modeling.

¹ We reported the performance of NL I3D described in [50], which is a variant of the original NL I3D [49]. It uses fewer temporal dimension pooling to achieve good performance, but also incur larger computation.

²Includes parameters and FLOPs of the Region Proposal Network.

Something-Something-V1. Something-Something-V1 is a challenging dataset, as activity cannot be inferred merely from individual frames (*e.g.*, pushing something from *right to left*). We compared TSM with current state-of-the-art methods in Table 2. We only applied center crop during testing to ensure the efficiency unless otherwise specified. TSM achieves *the first place* on the leaderboard upon publication.

We first show the results of the 2D based methods TSN [48] and TRN [58]. TSN with different backbones fails to achieve decent performance (<20% Top-1) due to the lack of temporal modeling. For TRN, although **late temporal fusion** is added after feature extraction, the performance is still significantly lower than state-of-the-art methods, showing the importance of temporal fusion across all levels.

The second section shows the state-of-the-art efficient video understanding framework ECO [61]. ECO uses an early 2D + late 3D architecture which enables **medium-level temporal fusion**. Compared to ECO, our method achieves better performance at a smaller FLOPs. For example, when using 8 frames as input, our TSM achieves 45.6% top-1 accuracy with 33G FLOPs, which is 4.2% higher accuracy than ECO with 1.9× less computation. The ensemble versions of ECO (ECO_{EnLite} and ECO_{EnLite}_{RGB+Flow}, using an ensemble of {16, 20, 24, 32} frames as input) did achieve competitive results, but the computation and parameters are too large for deployment. While our model is much more efficient: we only used {8, 16} frames model for ensemble (TSM_{En}), and the model achieves better performance using 2.7× less computation and 3.1× fewer parameters.

The third section contains previous state-of-the-art methods: Non-local I3D + GCN [50], that enables **all-level temporal fusion**. The GCN needs a Region Proposal Network [34] trained on MSCOCO object detection dataset [30]

Table 4. Results on Something-Something-V2. Our TSM achieves state-of-the-art performance.

Method	Val		Test	
	Top-1	Top-5	Top-1	Top-5
TSN (our impl.)	30.0	60.5	-	-
MultiScale TRN [58]	48.8	77.6	50.9	79.3
2-Stream TRN [58]	55.5	83.1	56.2	83.2
TSM _{8F}	59.1	85.6	-	-
TSM _{16F}	63.4	88.5	64.3	89.6
TSM _{RGB+Flow}	66.0	90.5	66.6	91.3

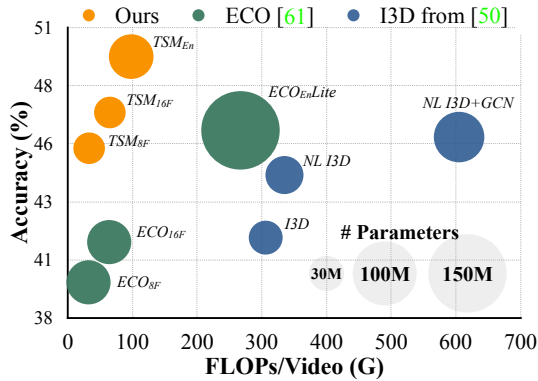


Figure 5. TSM enjoys better accuracy-cost trade-off than I3D family and ECO family on Something-Something-V1 [14] dataset. (GCN includes the cost of ResNet-50 RPN to generate region proposals.)

to generate the bounding boxes, which is unfair to compare since external data (MSCOCO) and extra training cost is introduced. Thus we compared TSM to its CNN part: Non-local I3D. Our TSM (8f) achieves 1.2% better accuracy with 10× fewer FLOPs on the validation set compared to the Non-local I3D network. Note that techniques like Non-local module [49] are orthogonal to our work, which could also be added to our framework to boost the performance further.

Generalize to Other Modalities. We also show that our proposed method can generalize to other modalities like optical flow. To extract the optical flow information between frames, we followed [48] to use the TVL1 optical flow algorithm [55] implemented in OpenCV with CUDA. We conducted two-stream experiments on both Something-Something V1 and V2 datasets, and it consistently improves over the RGB performance: introducing optical flow branch brings 5.4% and 2.6% top-1 improvement on V1 and V2.

Something-Something-V2. We also show the result on Something-Something-V2 dataset, which is a newer release to its previous version. The results compared to other state-of-the-art methods are shown in Table 4. On Something-Something-V2 dataset, we achieved state-of-the-art performance while only using RGB input.

Cost vs. Accuracy. Our TSM model achieves very competitive performance while enjoying high efficiency and low computation cost for fast inference. We show the FLOPs for each model in Table 2. Although GCN itself is light, the method used a ResNet-50 based Region Proposal Net-

Table 5. TSM enjoys low GPU inference latency and high throughput. V/s means videos per second, higher the better (Measured on NVIDIA Tesla P100 GPU).

Model	Efficiency Statistics				Accuracy	
	FLOPs	Param.	Latency	Thruput.	Sth.	Kinetics
I3D from [50]	306G	35.3M	165.3ms	6.1V/s	41.6%	-
ECO _{16F} [61]	64G	47.5M	30.6ms	45.6V/s	41.4%	-
I3D from [49]	33G	29.3M	25.8ms	42.4V/s	-	73.3%
I3D _{replace}	48G	33.0M	28.0ms	37.9V/s	44.9%	-
TSM _{8F}	33G	24.3M	17.4ms	77.4V/s	45.6%	74.1%
TSM _{16F}	65G	24.3M	29.0ms	39.5V/s	47.2%	74.7%

work [34] to extract bounding boxes, whose cost is also considered in the chart. Note that the computation cost of optical flow extraction is usually larger than the video recognition model itself. Therefore, we do not report the FLOPs of two-stream based methods.

We show the accuracy, FLOPs, and number of parameters trade-off in Figure 5. The accuracy is tested on the validation set of Something-Something-V1 dataset, and the number of parameters is indicated by the area of the circles. We can see that our TSM based methods have a better Pareto curve than both previous state-of-the-art efficient models (ECO based models) and high-performance models (non-local I3D based models). TSM models are both efficient and accurate. It can achieve state-of-the-art accuracy at high efficiency: it achieves better performance while consuming 3× less computation than the ECO family. Considering that ECO is already an efficiency-oriented design, our method enjoys highly competitive hardware efficiency.

5.4. Latency and Throughput Speedup

The measured inference latency and throughput are important for the large-scale video understanding. TSM has low latency and high throughput. We performed measurement on a single NVIDIA Tesla P100 GPU. We used batch size of 1 for latency measurement; batch size of 16 for throughput measurement. We made two comparisons:

(1) Compared with the I3D model from [50], our method is faster by an order of magnitude at 1.8% higher accuracy (Table 5). We also compared our method to the state-of-the-art efficient model ECO [61]: Our TSM model has 1.75× lower latency (17.4ms vs. 30.6ms), 1.7× higher throughput, and achieves 2% better accuracy. ECO has a two-branch (2D+3D) architecture, while TSM only needs the in-expensive 2D backbone.

(2) We then compared TSM to efficient 3D model designs. One way is to only inflate the first 1 × 1 convolution in each of the block as in [49], denoted as "I3D from [49]" in the table. Although the FLOPs are similar due to pooling, it suffers from 1.5× higher latency and only 55% the throughput compared with TSM, with worse accuracy. We speculate the reason is that TSM model only uses 2D convolution which is highly optimized for hardware. To exclude the factors

Table 6. Comparing the accuracy of offline TSM and online TSM on different datasets. Online TSM brings negligible latency overhead.

Model	Latency	Kinetics	UCF101	HMDB51	Something
TSN	4.7ms	70.6%	91.7%	64.7%	20.5%
+Offline	-	74.1%	95.9%	73.5%	47.3%
+Online	4.8ms	74.3%	95.5%	73.6%	46.3%

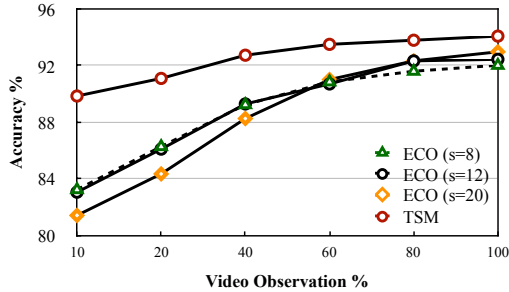


Figure 6. Early recognition on UCF101. TSM gives high prediction accuracy after only observing a small portion of the video.

of backbone design, we replace every TSM primitive with $3 \times 1 \times 1$ convolution and denote this model as $I3D_{\text{replace}}$. It is still much slower than TSM and performs worse.

5.5. Online Recognition with TSM

Online vs. Offline Online TSM models shift the feature maps uni-directionally so that it can give predictions in real time. We compare the performance of offline and online TSM models to show that online TSM can still achieve comparable performance. Follow [61], we use the prediction averaged from all the frames to compare with offline models, *i.e.*, we compare the performance after observing the whole videos. The performance is provided in Table 6. We can see that for less temporal related datasets like Kinetics, UCF101 and HMDB51, the online models achieve comparable and sometimes even better performance compared to the offline models. While for more temporal related datasets Something-Something, online model performs worse than offline model by 1.0%. Nevertheless, the performance of online model is still significantly better than the 2D baseline.

We also compare the per-frame prediction latency of pure 2D backbone (TSN) and our online TSM model. We compile both models with TVM [5] on GPU. Our online TSM model only adds to less than 0.1ms latency overhead per frame while bringing up to 25% accuracy improvement. It demonstrates online TSM is hardware-efficient for latency-critical real-time applications.

Early Recognition Early recognition aims to classify the video while only observing a small portion of the frames. It gives fast response to the input video stream. Here we compare the early video recognition performance on UCF101 dataset (Figure 6). Compared to ECO, TSM gives much higher accuracy, especially when only observing a small portion of the frames. For example, when only observing the first 10% of video frames, TSM model can achieve 90% accuracy, which is 6.6% higher than the best ECO model.

Table 7. Video detection results on ImageNet-VID.

Model	Online	Need Flow	Latency	mAP			
				Overall	Slow	Medium	Fast
R-FCN [23]	✓		1×	74.7	83.6	72.5	51.4
FGFA [60]		✓	2.5×	75.9	84.0	74.4	55.6
Online TSM	✓		1×	76.3	83.4	74.8	56.0

Table 8. TSM efficiently runs on edge devices with low latency.

Devices	Jetson Nano		Jetson TX2		Rasp.	Note8	Pixel1
	CPU	GPU	CPU	GPU			
Latency (ms)	47.8	13.4	36.4	8.5	69.6	34.5	47.4
Power (watt)	4.8	4.5	5.6	5.8	3.8	-	-

Online Object Detection Real-time online video object detection is an important application in self-driving vehicles, robotics, *etc.* By injecting our online TSM into the backbone, we can easily take the temporal cues into consideration at negligible overhead, so that the model can handle poor object appearance like motion blur, occlusion, defocus, *etc.* We conducted experiments on R-FCN [23] detector with ResNet-101 backbone on ImageNet-VID [35] dataset. We inserted the uni-directional TSM to the backbone, while keeping other settings the same. The results are shown in Table 7. Compared to 2D baseline R-FCN [23], our online TSM model significantly improves the performance, especially on the fast moving objects, where TSM increases mAP by 4.6%. We also compare to a strong baseline FGFA [60] that uses optical flow to aggregate the temporal information from 21 frames (past 10 frames and future 10 frames) for offline video detection. Compared to FGFA, TSM can achieve similar or higher performance while enabling online recognition at much smaller latency. We visualize some video clips in the supplementary material to show that online TSM can leverage the temporal consistency to correct mis-predictions.

Edge Deployment TSM is mobile device friendly. We build an online TSM model with MobileNet-V2 backbone, which achieves 69.5% accuracy on Kinetics. The latency and energy on NVIDIA Jetson Nano & TX2, Raspberry Pi 4B, Samsung Galaxy Note8, Google Pixel-1 is shown in Table 8. The models are compiled using TVM [5]. Power is measured with a power meter, subtracting the static power. TSM achieves low latency and low power on edge devices.

6. Conclusion

We propose Temporal Shift Module for hardware-efficient video recognition. It can be inserted into 2D CNN backbone to enable joint spatial-temporal modeling at no additional cost. The module shifts part of the channels along temporal dimension to exchange information with neighboring frames. Our framework is both efficient and accurate, enabling low-latency video recognition on edge devices.

Acknowledgments We thank MIT Quest for Intelligence, MIT-IBM Watson AI Lab, MIT-SenseTime Alliance, Samsung, SONY, AWS, Google for supporting this research. We thank Oak Ridge National Lab for Summit supercomputer.

References

- [1] The 20bn-jester dataset v1. <https://20bn.com/datasets/jester>. 5, 6
- [2] Hakan Bilen, Basura Fernando, Efstratios Gavves, Andrea Vedaldi, and Stephen Gould. Dynamic image networks for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3034–3042, 2016. 2
- [3] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. 3
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 4724–4733. IEEE, 2017. 1, 2
- [5] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. {TVM}: An automated end-to-end optimizing compiler for deep learning. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 578–594, 2018. 4, 8
- [6] Sharan Chetlur, Cliff Woolley, Philippe Vandermersch, Jonathan Cohen, John Tran, Bryan Catanzaro, and Evan Shelhamer. cudnn: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759*, 2014. 4
- [7] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. 1, 2
- [8] Christoph Feichtenhofer, Axel Pinz, and Richard Wildes. Spatiotemporal residual networks for video action recognition. In *Advances in neural information processing systems*, pages 3468–3476, 2016. 2
- [9] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1933–1941, 2016. 1, 2
- [10] Chuang Gan, Chen Sun, Lixin Duan, and Boqing Gong. Webly-supervised video recognition by mutually voting for relevant web images and web video frames. In *European Conference on Computer Vision*, pages 849–866. Springer, 2016. 2
- [11] Chuang Gan, Naiyan Wang, Yi Yang, Dit-Yan Yeung, and Alex G Hauptmann. Devnet: A deep event network for multimedia event detection and evidence recounting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2568–2577, 2015. 2
- [12] Chuang Gan, Ting Yao, Kuiyuan Yang, Yi Yang, and Tao Mei. You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 923–932, 2016. 2
- [13] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *CVPR*, volume 2, page 3, 2017. 1, 2
- [14] 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 *The IEEE International Conference on Computer Vision (ICCV)*, volume 1, page 3, 2017. 5, 6, 7
- [15] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations*, 2016. 3
- [16] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pages 1135–1143, 2015. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4, 6
- [18] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–800, 2018. 3
- [19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [20] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 3
- [21] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016. 3
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 5
- [23] Kaiming He Jian Sun Jifeng Dai, Yi Li. R-FCN: Object detection via region-based fully convolutional networks. 2016. 8
- [24] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. 1, 2
- [25] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 4, 5
- [26] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video

- database for human motion recognition. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2556–2563. IEEE, 2011. 5
- [27] Myunggi Lee, Seungeui Lee, Sungjoon Son, Gyutae Park, and Nojun Kwak. Motion feature network: Fixed motion filter for action recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 387–403, 2018. 2
- [28] Zhenyang Li, Kirill Gavriluk, Efstratios Gavves, Mihir Jain, and Cees GM Snoek. Videolstm convolves, attends and flows for action recognition. *Computer Vision and Image Understanding*, 166:41–50, 2018. 2
- [29] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *Advances in Neural Information Processing Systems*, pages 2181–2191, 2017. 3
- [30] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [31] Chenxi Liu, Barret Zoph, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. *arXiv preprint arXiv:1712.00559*, 2017. 3
- [32] Xiang Long, Chuang Gan, Gerard de Melo, Jiajun Wu, Xiao Liu, and Shilei Wen. Attention clusters: Purely attention based local feature integration for video classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7834–7843, 2018. 2
- [33] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5534–5542. IEEE, 2017. 2
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 6, 7
- [35] 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. 8
- [36] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 3, 6
- [37] Shikhar Sharma, Ryan Kiros, and Ruslan Salakhutdinov. Action recognition using visual attention. *arXiv preprint arXiv:1511.04119*, 2015. 2
- [38] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, pages 510–526. Springer, 2016. 5
- [39] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014. 1, 2
- [40] 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. 5
- [41] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhutdinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015. 2
- [42] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4597–4605, 2015. 2
- [43] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 2
- [44] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. *arXiv preprint arXiv:1807.11626*, 2018. 3
- [45] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 1, 2
- [46] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. 1, 2
- [47] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization. *arXiv preprint arXiv:1811.08886*, 2018. 3
- [48] 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 *European Conference on Computer Vision*, pages 20–36. Springer, 2016. 1, 2, 5, 6, 7
- [49] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. *arXiv preprint arXiv:1711.07971*, 10, 2017. 1, 2, 5, 6, 7
- [50] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. *arXiv preprint arXiv:1806.01810*, 2018. 2, 5, 6, 7
- [51] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, and Kurt Keutzer. Shift: A zero flop, zero parameter alternative to spatial convolutions. *arXiv preprint arXiv:1711.08141*, 2017. 2, 3
- [52] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 6
- [53] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning:

- Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 305–321, 2018. 1, 2, 5
- [54] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015. 2
- [55] 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. Springer, 2007. 2, 7
- [56] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *CoRR*, abs/1707.01083, 2017. 3
- [57] Huasong Zhong, Xianggen Liu, Yihui He, Yuchun Ma, and Kris Kitani. Shift-based primitives for efficient convolutional neural networks. *arXiv preprint arXiv:1809.08458*, 2018. 3
- [58] Bolei Zhou, Alex Andonian, and Antonio Torralba. Temporal relational reasoning in videos. *arXiv preprint arXiv:1711.08496*, 2017. 1, 2, 5, 6, 7
- [59] Chenzhuo Zhu, Song Han, Huizi Mao, and William J Dally. Trained ternary quantization. *International Conference on Learning Representations*, 2016. 3
- [60] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 408–417, 2017. 8
- [61] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. Eco: Efficient convolutional network for online video understanding. *arXiv preprint arXiv:1804.09066*, 2018. 1, 2, 5, 6, 7, 8
- [62] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 3
- [63] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. *arXiv preprint arXiv:1707.07012*, 2(6), 2017. 3