

Recurrent Residual Module for Fast Inference in Videos

Bowen Pan^{1†}, Wuwei Lin^{1†}, Xiaolin Fang^{2¶}, Chaoqin Huang^{1†}, Bolei Zhou^{3§}, Cewu Lu^{1‡*}

¹Shanghai Jiao Tong University, ²Zhejiang University, ³Massachusetts Institute of Technology

[†]{googletornado, linwuwei13, huangchaoqin}@sjtu.edu.cn, [¶]fxlfang@gmail.com

[§]bzhou@csail.mit.edu; [‡]lucewu@sjtu.edu.cn

Abstract

Deep convolutional neural networks (CNNs) have made impressive progress in many video recognition tasks such as video pose estimation and video object detection. However, CNN inference on video is computationally expensive due to processing dense frames individually. In this work, we propose a framework called Recurrent Residual Module (RRM) to accelerate the CNN inference for video recognition tasks. This framework has a novel design of using the similarity of the intermediate feature maps of two consecutive frames, to largely reduce the redundant computation. One unique property of the proposed method compared to previous work is that feature maps of each frame are precisely computed. The experiments show that, while maintaining the similar recognition performance, our RRM yields averagely $2\times$ acceleration on the commonly used CNNs such as AlexNet, ResNet, deep compression model (thus $8-12\times$ faster than the original dense models using the efficient inference engine), and impressively $9\times$ acceleration on some binary networks such as XNOR-Nets (thus $500\times$ faster than the original model). We further verify the effectiveness of the RRM on speeding up CNNs for video pose estimation and video object detection.

1. Introduction

Video understanding is one of the long-standing topics in computer vision. Recently, deep convolutional neural networks (CNNs) advanced different tasks of video understanding, such as video classification [33, 59, 58, 60], video pose estimation [16, 5], and video object detection [18, 17, 47, 39, 45, 46]. However, using CNNs to process the dense frames of videos is computationally expensive while it becomes unaffordable as the video goes longer. Meanwhile, millions of videos are shared on the Internet, where processing and extracting useful information remains a challenge. With the video datasets becoming larger and

larger [49, 1, 33, 34, 15, 41], training and evaluating neural networks for video recognition are more challenging. For example, for Youtube-8M dataset [1] with over 8 million video clips, it will take 50 years for a CPU to extract the deep features using a standard CNN model.

One of the bottlenecks for video understanding using CNNs is the frame-by-frame CNN inference. A one-minute video contains thousands of frames thus the model inference becomes much slower in comparison with processing a single image. However, different from a set of independent images, consecutive frames in a video clip are usually similar. Thus, the high-level semantic feature maps in the deep convolutional neural networks of the consecutive frames will also be similar. Intuitively, we can leverage the frame similarity to reduce some redundant computation in the frame-by-frame video CNN inference. An attractive recursive schema is as follows:

$$\mathcal{R}(I_t) = \mathcal{R}(I_{t-1}) + \mathcal{G}(I_t - I_{t-1}), \quad (1)$$

where \mathcal{R} is the deep CNN feature, \mathcal{G} is a fast and shallow network that only processes the frame difference between frame I_t and I_{t-1} in a video clip. Ideally, \mathcal{G} should be both efficient and accurate to extract the residual feature. However, it remains challenging to implement such a schema due to the nonlinearity of CNNs.

Some previous works have tried to address this nonlinearity. Zhu *et al.* [61] proposed deep feature flow framework which utilizes the flow field to propagate the deep feature maps. However, these estimated feature maps will cause a drop on performance compared to the original feature maps. Kang *et al.* [32] developed a NoScope system to perform the fast binary query of the absence of a specific category. It is fast but not generic enough for other video recognition tasks.

We propose the framework of Recurrent Residual Module (RRM) to thoroughly address the nonlinear issue of CNNs in Eq. 1. The nonlinearity of CNNs results from the pooling layers and activation functions, while the computationally expensive layers such as convolution layer and fully-connected layer are linear. Thus for two consecutive frame inferences, if we are able to share the overlapped calculation of these linear layers, a large amount of the computation can be elim-

*The corresponding author is Cewu Lu. Cewu Lu is also member of AI research institution at SJTU. This work was done when Xiaolin Fang was an intern at MVIG lab of Shanghai Jiao Tong University.

inated. To this end, we snapshot the input and output feature maps of convolution layers and fully-connected layers for the inference on the next frame. Consequently, we only need to forward pass the frame difference region with the feature maps of the previous frame in each layer, which leads to the sparsity matrix multiplication that can be largely accelerated by the EIE techniques [22]. In general, our RRM can dramatically reduce the computation cost from the convolution layers and fully-connected layers, while still maintains the nonlinearity of the whole network.

The main contribution of this work is the framework of Recurrent Residual Module, which is able to speed up almost any CNN-based models for video recognition without extra training cost. To the best of our knowledge, this is the first acceleration method that can compute the feature maps precisely when deep CNNs process videos. We evaluate the proposed method and verify its effectiveness on accelerating CNNs for video recognition tasks such as video pose estimation and the video object detection.

2. Related Work

We have a brief survey on the related work of improving the neural network efficiency as below.

Network weight pruning. It is known that removing the redundant model parameters reduces the computational complexity of networks [36, 25, 26, 55, 9]. At the very beginning, Hanson & Pratt [25] applied the weight decay method to prune the network, then Optimal Brain Damage (OBD) [36] and Optimal Brain Surgeon (OBS) [26] pruned the parameters using the Hessian of the loss function. Recently, Han *et al.* [24, 23] showed that they could even reduce the model parameters by an order of magnitude in deep CNN models while maintaining the performance. They devised an efficient inference engine [22] to speed up the models. Instead of pruning model weights, our RRM framework focuses on factorizing the input at each layer, then further speeds up the model based on the pruning methods.

Network quantization. Quantizing network weight is to replace the high-precision float numbers of the weights with several limited integers, such as $+1/-1$ [54, 10, 11, 43, 37] or $+1/0/-1$ [4]. Rastegari *et al.* [43] proposed XNOR-Networks that use both binary weights and binary inputs to achieve $58\times$ faster convolution operations on a CNN trained on ImageNet. Yet, applying these quantization methods requires retraining the model and also results in a loss of accuracy.

Low rank acceleration. Decomposing weight tensor based on low-rank methods are used to accelerate deep convolutional networks. Both [13, 31] reduced the redundancy of the weight tensors through the low-rank approximation. Yang *et al.* [57] showed that they can use a single Fast-food layer to replace the FC layer. Liu *et al.* [38] reduced the computation complexity using a sparse decomposition. All of these methods speed up the test-time evaluation of

convolutional networks with some sacrifice in precision.

Filter optimization. Reducing the filter redundancy in convolution layers is an effective method to simplify the CNN models [40, 28, 29]. Luo *et al.* [40] pruned filters and set the output feature maps as the optimization objective to minimize the loss of information. Howard *et al.* [29] developed MobileNet which applied depth-wise separable convolution to decompose a standard convolution operation and showed an effectiveness. He *et al.* [28] proposed an iterative algorithm to jointly learn additional filters for filter selection and scalar masks for each output channel. They achieved $13\times$ speedup on AlexNet.

Sparsity. It is most related to our method. Obviously, sparsity can significantly accelerate the convolutional networks both in training and testing [38, 6, 21, 56]. There are many previous works showing that they can save the energy [8, 44] and accelerate the convolution [2, 48, 14] by skipping the zeros or elements close to zero in the sparse input. Albericio *et al.* [2] proposed an efficient convolution accelerator utilizing the sparsity of inputs, while Shi & Chu [48] sped up the convolution on CPUs by eliminating the zero values in the output of ReLUs. Graham & Maaten [20, 19] introduced a sparse convolution that eliminated the computation of values in some inactive output positions by recognizing the input cells in the ground state. Recently, Han *et al.* [22] devised an efficient inference engine (EIE) that can exploit the dynamic sparsity of the input feature maps to accelerate the inference. Our RRM integrates EIE as a step to further optimize the model weight.

Our Recurrent Residual Module works in a recurrent manner. The most similar architecture to ours is the Predictive-Corrective Networks [12], which derives a series of recurrent neural networks to make prediction about feature and then correct them with some bottom-up observations. The key difference, also the most innovative point of our model, is that we utilize the recurrent framework to accelerate CNN models using sparsity and Efficient Inference Engine, which is much more efficient than the Predictive-Corrective Networks [12]. Besides our method is a generic framework that could be plugged in a variety of CNN models without retraining to speed up the forward pass.

3. Recurrent Residual Module Framework

The key idea of the Recurrent Residual Module is to utilize the similarity between the consecutive frames in a video clip to accelerate the model inference. To be more specific, we first improve the sparsity of the input to each **linear layer** (layers with linearity, including convolution layer and FC layer), then use the sparse matrix-vector multiplication accelerators (SPMV) to further speed up the forward pass.

We will first introduce some preliminary concepts and discuss the linearity of convolution layers and FC layers. Then the recurrent residual module will be introduced in

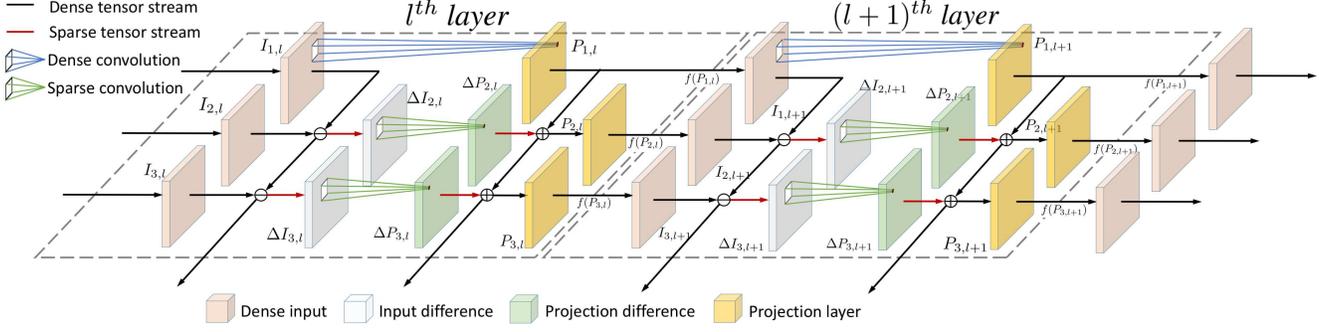


Figure 1: Illustration of the Recurrent Residual Module on two layers. *Dense convolution* operation represents the standard convolution operation. *Sparse convolution* operation represents the SPMV-equipped convolution operation which will deliver speed up to the sparse input. *Sparse convolution* has no bias term and shares the same weight filters with *Dense convolution*. Mapping f represents some nonlinear operator combinations in CNNs such as ReLUs and max poolings.

detail, followed by the analysis of computation complexity, sparsity enhancement, and accumulated error. Last but not least, we integrate the efficient inference engine [22] (EIE) to further improve the framework’s efficiency.

3.1. Preliminary

We denote a standard neural network using the notion set $\langle \mathcal{I}, \mathcal{F}, *, \mathcal{W}, f \rangle$, where \mathcal{I} represents the set of input tensor (it could be the input image or the output from the previous layer), \mathcal{F} is the set of weight filters in convolution layers, $*$ denotes the convolution operations, \mathcal{W} represents the set of weight tensors in FC layers, and f represents some nonlinear operators. In convolution phase, f can be a ReLU [42] or a pooling operator. And in the fully-connected phase, it can be a short-cut function.

We use $I_{tl} \in \mathcal{I}$ to denote the input tensor to the l^{th} linear layer when we process the t^{th} frame in the video, $W_l \in \mathcal{W}$ to represent the weight tensor of the l^{th} layer if it is FC layer, $F_l \in \mathcal{F}$ to represent the weight filter of the l^{th} layer if it is convolution layer. When processing the t^{th} frame, the l^{th} layer performs the following operation:

$$I_{t(l+1)} = \begin{cases} f(F_l * I_{tl} + b_l) & \text{if } l^{th} \text{ layer is convolution} \\ f(W_l I_{tl} + b_l) & \text{if } l^{th} \text{ layer is FC} \end{cases}, \quad (2)$$

where b_l is the bias term of the l^{th} layer. And we define the **projection layer** P_{tl} as:

$$P_{tl} = \begin{cases} F_l * I_{tl} + b_l & \text{if } l^{th} \text{ layer is convolution} \\ W_l I_{tl} + b_l & \text{if } l^{th} \text{ layer is FC} \end{cases}. \quad (3)$$

Due to the linearity of convolution operation and multiplication operation, given the difference of P_{tl} and $P_{(t-1)l}$, we have:

$$P_{tl} - P_{(t-1)l} = \begin{cases} F_l * \Delta I_{tl} & \text{if } l^{th} \text{ layer is convolution} \\ W_l \Delta I_{tl} & \text{if } l^{th} \text{ layer is FC} \end{cases}, \quad (4)$$

where $\Delta I_{tl} = I_{tl} - I_{(t-1)l}$. Thus Eq. 2 can be written as:

$$I_{t(l+1)} = \begin{cases} f(P_{(t-1)l} + F_l * \Delta I_{tl}) & \text{if } l^{th} \text{ layer is convolution} \\ f(P_{(t-1)l} + W_l \Delta I_{tl}) & \text{if } l^{th} \text{ layer is FC} \end{cases}. \quad (5)$$

Eq. 5 is the key point in our RRM framework. $P_{(t-1)l}$ has been obtained and preserved during the inference phase of the last frame. Evidently, the computation mainly falls on $F_l * \Delta I_{tl}$ or $W_l \Delta I_{tl}$. Due to the similarity between the consecutive frames, ΔI_{tl} is usually highly sparse (This will be verified in our experiment). As a result, to obtain the final result, we just need to work on a rather sparse tensor ΔI_{tl} instead of the original one I_{tl} , which is dense and computationally expensive. With the help of sparse matrix-vector multiplication accelerators (SPMV), the calculations of zero elements can be skipped, thus inference speed is improved.

3.2. Recurrent Residual Module for Fast Inference

The illustration of the recurrent residual module (RRM) is shown in Fig. 1. In order to preserve the information of the last frame and obtain the efficient \mathcal{G} which is introduced in Sec. 1, the information of input tensor to each linear layer $\mathcal{I}_{(t-1)}$ and the corresponding projection layer set $\mathcal{P}_{(t-1)}$ of each linear layer is saved. The preserved information can be applied during the inference phase for the following frame.

As shown in the Fig. 1, in the inference stream of frame 2, when the input tensor $I_{2,l}$ is fed to the convolution layer (the l^{th} layer), we first subtract $I_{1,l}$ from $I_{2,l}$ to obtain $\Delta I_{2,l}$, where $I_{1,l}$ is the input tensor to the l^{th} layer of frame 1 and was snapshotted when processing frame 1. As illustrated in the previous discussion, $\Delta I_{2,l}$ is a sparse tensor. Apply the sparse matrix-vector multiplication accelerator to the l^{th} layer, we can skip the zero elements and get the convolution result within a short time. Next, the output of the convolution layer is snapshotted. Add the output to projection layer $P_{1,l}$, we can obtain the intact tensor that is exactly the same as the output of a normal convolution layer which is fed $I_{2,l}$. After that, we perform the nonlinear mapping f to $P_{1,l}$. In this manner, the final result is obtained. To some extent, it is similar to the distributive law of multiplication.

The specific procedure of the inference with Recurrent

Residual Module is listed in Algorithm 1.

Algorithm 1 Inference with Recurrent Residual Module

Input: A video clip $\mathcal{X} := \{x_t | t = 1, 2, \dots, T\}$, where x_t is the frame at time t , a pre-trained neural network \mathcal{M} .

Output: Frame-level feature set $\mathcal{F} := \{f_t | t = 1, 2, \dots, T\}$, where f_t is the deep feature of the frame x_t .

- 1: $I_{0,:} \leftarrow 0$
 - 2: **for** $t = 1$ to T **do**
 - 3: $I_{t,0} \leftarrow x_t$
 - 4: **for** l, L in $\text{enumerate}(\mathcal{M})$ **do**
 - 5: $\hat{I}_{t,l} \leftarrow I_{t,l} - I_{(t-1),l}$
 - 6: **if** L is convolution layer **then**
 - 7: $P_{tl} = P_{(t-1)l} + F_l * \hat{I}_{t,l}$
 - 8: **else**
 - 9: $P_{tl} = P_{(t-1)l} + W_l \hat{I}_{t,l}$
 - 10: $I_{t,l+1} \leftarrow f(P_{tl})$
 - 11: $f_t \leftarrow I_{t,T}$
-

One drawback of the RRM is that we can only forward pass frames with the help of the feature snapshots of the previous frames, which limits doing inference in parallel for the whole video. To address this we can split the video into several chunks then process each chunk with RRM-equipped CNN in parallel.

3.3. Analyzing computational complexity

Layer Type	Complexity
Convolution layer	$O(W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F)$
Convolution layer + SPMV	$O(\rho_{c_i} W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F)$
FC layer	$O(C_{f_j}^{in} C_{f_j}^{out})$
FC layer + SPMV	$O(\rho_{f_j} C_{f_j}^{in} C_{f_j}^{out})$

Table 1: Ablation analysis of computational complexity. Layers equipped with SPMV will skip the calculations of zero elements.

The computational complexity of the neural network with the recurrent residual module in test-phase is analyzed. In a sequence of convolution layers $M_{c_1}, M_{c_2}, \dots, M_{c_n}$, suppose that for layer M_{c_i} , the density (the proportion of non-zero elements) of the input tensor $I_{c_i} \in R^{C_{c_i}^{in} \times W_{c_i} \times H_{c_i}}$ is ρ_{c_i} , the weight matrices is $F_{c_i} \in R^{C_{c_i}^{in} \times C_{c_i}^{out} \times W_F \times H_F}$. Similarly, for an FC layer M_{f_j} , we have the density ρ_{f_j} , the input vector $I_{f_j} \in R^{C_{f_j}^{in}}$ and the weight tensor $W_{f_j} \in R^{C_{f_j}^{in} \times C_{f_j}^{out}}$.

In our Recurrent Residual Module, compared to the multiplication operation, both execution time and computational cost of add operation are trivial. Hence, to analyze the computation complexity, the following discussion will only focus on the multiplication complexity in the original linear layer

and in our RRM framework. Table 1 shows the multiplication complexity of a single layer. For the entire neural network, the computational complexity after utilizing the sparsity can be calculated as follows (assume that the stride is 1):

$$O\left(\sum_i \rho_{c_i} W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F + \sum_j \rho_{f_j} C_{f_j}^{in} C_{f_j}^{out}\right). \quad (6)$$

Eq. 6 illustrates that the sparsity (the proportion of zero elements) of the input tensor to each layer is the key to reduce the computation cost. In terms of the sparsity, some networks equipped with ReLU activation functions already have many zero elements in their feature maps. In our recurrent residual architecture, the sparsity can be further improved as discussed below.

3.4. Improving sparsity

Our framework can obtain the inference output identical to the original model without any approximation. And we could further improve the sparsity of the intermediate feature map to approximate the inference output as a trade-off to further accelerate inference. However, it would possibly lead to the issue of error accumulation over time. To address this issue, we estimate the accumulated error given by accumulated truncated values. First, the accumulated truncated values are obtained by

$$e_t = \sum_t \sum_j \ell_2(u_{t,j}), \quad (7)$$

where $u_{t,j}$ is the truncated map to the j^{th} linear layer in the inference stream of t^{th} frame. We denote accumulated accuracy error by

$$e_c = \mathcal{H}(e_t, \mu). \quad (8)$$

\mathcal{H} is a fourth order Polynomial function regression with the parameter μ , which is fitted from large amount of data pairs of accumulated truncate value and accumulated error. If it is larger than a certain threshold, a new precise inference will be carried out to clear accumulated error and a new round of fast inference will start.

3.5. Efficient inference engine

To implement the RRM framework efficiently, we utilize dynamic sparse matrix-vector multiplication (DSPMV) technique. While there are a number of existing off-the-shelf DSPMV techniques [22, 48], the most efficient one among them is the efficient inference engine (EIE) proposed by Han *et al.* [22].

EIE is the first accelerator which exploits the dynamic sparsity in the matrix-vector multiplications. When performing multiplication between matrix W and sparse vector a , the vector a is scanned and a Leading Non-zero Detection Node (LNZD Node) is applied to recursively look for the next non-zero element a_j . Once found, EIE broadcasts a_j along with its index j to the processing elements (PEs) which

hold the weight tensor in the CSC format. Then weights column W_j with the corresponding index j in all PEs will be multiplied by a_j and the results will be summed into the corresponding row accumulator. These accumulators finally output the resulting vector b .

Since the multiplication between matrix and matrix can be decomposed into several matrix-vector multiplication processes, by decomposing the input tensor to several dynamically sparse vectors, we embed the EIE to our RRM framework conveniently.

4. Experiments

In this section, we first verify that our recurrent residual module can consistently improve the sparsity of the input tensor to each layer in Sec. 4.1 across different network architectures. We measure the overall sparsity of the whole network to estimate the improvement. The overall sparsity is calculated as the ratio of zero-value elements in the inputs of all linear layers, which is:

$$S = \frac{\sum_i s_{c_i} W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F + \sum_j s_{f_j} C_{f_j}^{in} C_{f_j}^{out}}{\sum_i W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F + \sum_j C_{f_j}^{in} C_{f_j}^{out}}, \quad (9)$$

where s_{c_i} and s_{f_i} are the sparsity of the input tensor to the convolution layer M_{c_i} and the FC layer M_{f_j} respectively. Then, we show the speed and accuracy trade-off in our RRM framework. After that, we combine our RRM framework with some classical model acceleration techniques such as the XNOR-Net [43] and the Deep Compression models [23] to further accelerate the model inference. Finally, we demonstrate that we can accelerate several off-the-shelf CNN-based models, here we take the detectors in the field of pose estimation and object detection for examples. In this section, we provide a theoretical speedup ratio by computing the theoretical computational time of the EIE [22], which is calculated by dividing the total workload GOPs by the peak throughput. The actual computation time is around 10% more than the theoretical time due to the load imbalance. Yet, this bias will not affect our speedup ratio. For an uncompressed model, EIE has an impressive processing power of 3 TOP/s. We utilize its feature that it can exploit the dynamic sparsity of the activations. When both are equipped with EIE, the speedup ratio η of the model accelerated by RRM compared to the original model can be calculated as:

$$\eta = \frac{\sum_i \rho_{c_i} W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F + \sum_j \rho_{f_j} C_{f_j}^{in} C_{f_j}^{out}}{\sum_i \hat{\rho}_{c_i} W_{c_i} H_{c_i} C_{c_i}^{in} C_{c_i}^{out} w_F h_F + \sum_j \hat{\rho}_{f_j} C_{f_j}^{in} C_{f_j}^{out}}, \quad (10)$$

where $\hat{\rho}_{c_i}$ and $\hat{\rho}_{f_i}$ are the density of the input tensor in our RRM.

4.1. Results on the sparsity

To show that our RRM framework is able to generally improve the overall sparsity, we evaluate our method on three different real-time video benchmark datasets: Charades [50],

Model	Charades	UCF-101	MERL
AlexNet [35]	35.7%	35.4%	34.8%
AlexNet + RRM	57.5%	60.1%	71.8%
Improvement	21.8%	24.7%	37.0%
Speedup ratio	146%	154%	211%
VGG-16 [51]	50.4%	51.3%	53.2%
VGG-16 + RRM	66.4%	70.1%	75.2%
Improvement	16.0%	18.8%	22.0%
Speedup ratio	124%	128%	136%
ResNet-18 [27]	40.5%	40.4%	40.0%
ResNet-18 + RRM	58.0%	58.4%	73.6%
Improvement	17.5%	18.0%	33.6%
Speedup ratio	126%	130%	191%

Table 2: Overall sparsity improvement and the speedup ratio of each model evaluated on three benchmark datasets. Model plus RRM means that we apply our recurrent residual module to the original baseline model. RRM clearly brings significant improvement over the baselines.

UCF-101 [53], MERL [52], and choose three classical deep networks: AlexNet [35], VGG-16 [51], ResNet-18 [27] to be our base networks. In order to formulate the real-time analysis on videos, we sample the video frames at 24 FPS, which is the original frame rate in Charades, and then perform inference that extracts the deep features of these video frames. We measure the overall sparsity improvement of each network when performing inference with our RRM on these three datasets, during which the threshold ϵ in RRM (as is illustrated in Sec. 3.4) is set to be 10^{-2} . And the results are recorded in Table 2. It can be seen that our RRM framework can generally improve the overall sparsity of the input feature maps in DNNs and deliver a speedup as calculated by Eq. 10. This sparsity improvement comparison between datasets indicates that the similarity property of video frames is efficiently exploited by our RRM framework.

Here we also want to clarify the threshold setting. In fact, it makes no difference to treat such small-value elements as zero elements. The L_2 distance between the feature extracted under this setting and the original feature is generally around 10^{-6} . This is a trivial deviation for that, in contrast, translating the cropped image by one pixel can result in an L_2 error around 10^{-2} . As shown in Fig. 2, features extracted under this threshold setting have no difference with the original features.

4.2. Trade-off between accuracy and speedup

In Sec. 3.4, we introduced a sparsity enhancement scheme, which truncates some small values into zero. It can further accelerate the model, but bring some deviation between the calculated feature maps and the original feature

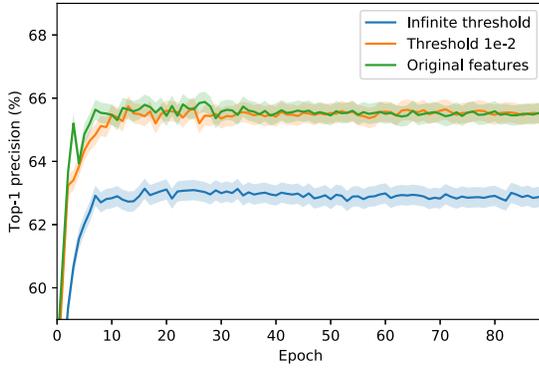


Figure 2: We train these features for the action recognition task on UCF-101 (as is discussed in Sec. 4.2). The blue curve represents the performance of the feature extracted with an infinite threshold. We add this ablation curve to show that precisely-computed features are necessary to obtain a good performance.

maps. Thus, there naturally exists a trade-off between speed and accuracy by adjusting the threshold ϵ .

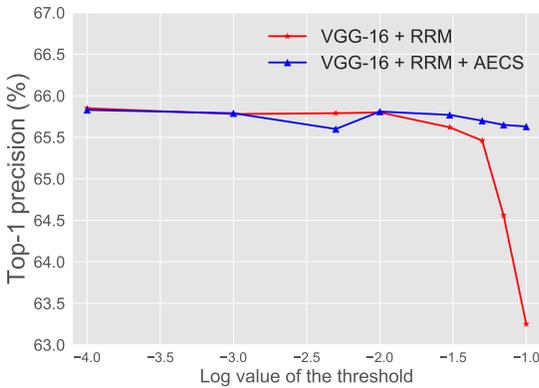


Figure 3: We measure the accuracy of features extracted under different threshold settings and validate the effectiveness of the accumulated error control scheme.

Threshold ϵ	10^{-2}	3×10^{-2}	5×10^{-2}	10^{-1}
Speedup ratio	$1.3\times$	$1.4\times$	$1.6\times$	$1.8\times$

Table 3: Different speedup ratios of VGG-16 at different level of thresholds.

We explore this trade-off by performing the action recognition task on UCF-101 dataset [53]. For each video, we first extract the VGG-16 feature vectors of its frames. Then, we perform the average pool on these feature vectors to obtain a video-level feature vector in 4096 dimensions to represent this video. With these video-level features, we train a two-layer MLP to recognize the actions in these videos and evaluate the top-1 precision. As is shown in Fig. 3, by gradually amplifying the threshold ϵ when extracting the feature, the speed up ratio increases while the accuracy drops due to the exploded accumulated error.

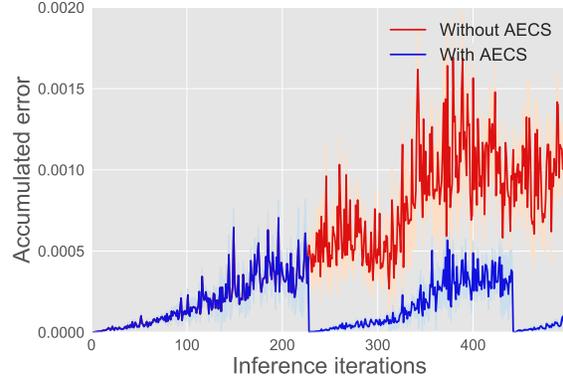


Figure 4: The dynamic accumulated error with the threshold ϵ set as 3×10^{-2} .

We then validate the effectiveness of accumulated error control scheme (AECS), which is introduced in Sec. 3.4. With the protection of AECS, the precision is maintained as the ϵ grows up. Dynamic accumulated error during inference is shown in Fig. 4. We can see that, with a moderate ϵ , the inference speed will not be affected since the expensive original inference is rare.

4.3. Speed up deeply compressed models

Model	Charades	MERL
Deep Compression	41.0%	37.3%
Deep Compression + RRM	59.4%	70.2%
Improvement	18.4%	32.9%
Speedup ratio	145%	210%
XNOR-Net	0.1%	0.1%
XNOR-Net + RRM	83.2%	89.2%
Improvement	83.1%	89.1%
Speedup ratio	598%	927%

Table 4: Comparison of overall sparsity of XNOR-Nets and deep compression models with RRM.

We examine the performance of RRM on some already-accelerated models and show that these models can be further accelerated by our RRM framework on video inference.

Deep compression model. Han *et al.* [23] proposed the deep compression model, which effectively reduces the model size and the energy consumption. There is a three-stage pipeline that prunes redundant connections between layers, quantizes parameters and compresses model with Huffman encoding. Deep compression model can be largely accelerated in efficient inference engine [22]. Efficient inference engine is a general methodology that compresses and accelerates DNNs. We show we can further accelerate the model when processing video frames.

XNOR-Net. Deep CNN models can be sped up by binarizing the input and the weight of the network. Rastegari *et al.* [43] devised the XNOR-Nets which approximated

the original model with binarized input and parameters and achieved a $58\times$ faster convolution operation. Value of elements in both the input and the weight of the XNOR-Net is transformed to $+1$ or -1 by taking their signs. Consequently, convolution operation can be implemented with only additions. The sparsity of feature maps in XNOR-Net is very poor due to the binarization. With RRM applied, the overall sparsity is significantly improved. Besides, after skipping zero-value input elements, the elements remained to be calculated are all $+2$ or -2 , where the advantages of binary convolution operation can still be maintained by scaling a factor 0.5.

Experiment results can be referred in Table 4. It demonstrates that our RRM is able to achieve an impressive speedup ratio on these compressed models.

4.4. Video pose estimation and object detection

In this section, we apply our RRM framework to several mainstream visual systems to improve the efficiency of their backbone CNN models. We choose two video recognition tasks, video pose estimation and video object detection, to verify the effectiveness of our RRM framework. We set the threshold ϵ as 10^{-2} in the experiments. It is a precise setting which has been validated by preceding experiments in Sec. 4.1 so that the output features are almost the same as the original model and the recognition performance will not be affected. Some qualitative results are shown in Fig. 6.

Model	MPII Video Pose	BBC Pose
rt-Pose[5]	78.5%	79.6%
rt-Pose + RRM	91.0%	93.3%
Improvement	12.5%	13.7%
Speedup ratio	213.7%	291.2%

Table 5: Comparison of the overall sparsity of pose estimator rt-Pose with RRM.

Video pose estimation. Real-time video pose estimation is a rising topic in computer vision. To meet the requirement of inference speed, our RRM can be applied for acceleration. Currently, the fastest multi-person pose estimator is the rt-Pose model proposed by Cao *et al.* [5], which can reach a speed of 8.8 FPS with one NVIDIA GeForce GTX-1080 GPU. In this part, we apply our RRM framework to further accelerate the rt-Pose model. We evaluate the models on two video pose datasets, BBC Pose[7] and MPII-Video-Pose [30]. The BBC Pose dataset consists of 20 TV broadcast videos (each 0.5h-1.5h in length) while the MPII Video Pose dataset is composed of 28 sequences which contains some challenging frames in the MPII dataset [3]. The experiment results are shown in Table 5, we can see that by applying our RRM, pose estimation in videos are significantly accelerated.

Video object detection. Majority of the work on object detection is focused on image rather than videos. Redmon

Model	Charades	UCF-101	MERL
YOLOv2[46]	2.45%	2.55%	2.46%
YOLOv2 + RRM	50.01%	52.26%	48.21%
Improvement	47.56%	49.71%	45.75%
Speedup ratio	204.9%	200.8%	197.3%

Table 6: Comparison of overall sparsity of object detection model YOLOv2 with RRM.

et al. [45, 46] created YOLO network, which achieved very efficient end-to-end training and testing for object detection. We apply our RRM framework to accelerates the YOLO network to realize a faster real-time detection in videos. We evaluate the models on video object detection on Charades, UCF-101, and MERL. YOLOv2 uses the Leaky-ReLU as the activation function, thus it prevents the sparsity of the original model. By applying our RRM, there brings a huge improvement. As shown in Table 6, the sparsity of original model ranges between 2% and 3%. With our RRM, the sparsity increases to 48%-52%. In total, our RRM brings a speedup ratio around 200%.

OBJECTS	mAP	KEYPOINTS	mAP
YOLOv2	61.2	rt-Pose	46.2
YOLOv2+RRM	61.1	rt-Pose+RRM	46.2

Table 7: Detection and pose estimation performance results.

Recognition accuracy. To prove that our method is able to maintain performance while greatly accelerate the model inference, we conduct the detection experiments on the Youtube-BB dataset using YOLOv2 and the pose estimation experiments on MPII video pose dataset using rt-Pose. We keep all the training conditions as the same. And the accuracy results are shown in Table 7.

4.5. Discussion

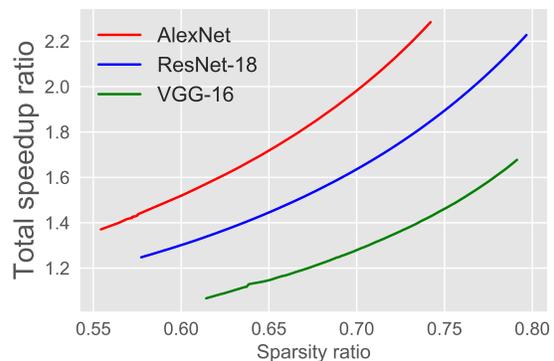


Figure 5: Trade-off comparison between speedup and sparsity ratio. We do experiments on the UCF-101 dataset.

Theoretical vs. Actual speedup. Hardware designing to evaluate actual speedup is beyond the scope of the current



Figure 6: Qualitative results of the object detection and pose estimation in videos. frames with green border are the original results evaluated in a frame-by-frame manner and frames with blue border are the results of our RRM framework. The performance is not affected, and interestingly, we find that our RRM framework sometimes gets a more reliable result than the original model. It could be that our RRM framework can utilize the temporal context information across consecutive frames.

work, while according to Table III in [22] actual speedup can be well estimated by the sparsity of weight and activation on EIE engine. It can be seen from Table III in [22] that the relationship between density of the layer ($\text{Weight}\% \times \text{Act}\%$) and the speedup of layer inference ($\text{FLOP}\%$) is near-linear. Thus, it can be inferred that, with well-designed hardwares, there won't be a significant performance gap between these theoretical numbers and those in real application.

Batch Normalization. Several studies have shown that the linear layer calculation only occupied part of total inference time, some other non-linear layers are also time-consuming, especially the BN layer. Thus, here we compare the trade-off between total speedup (with all overhead considered) and sparsity ratio among AlexNet (no BN), VGG-16 (no BN) and ResNet-18 (with BN) in Fig. 5.

5. Conclusion

We proposed the Recurrent Residual Module for fast inference in videos. We have shown that the overall sparsity of different CNN models can be generally improved by our RRM framework. Meanwhile, applying our RRM framework to some already-accelerated models, such as XNOR-Net and Deep Compression Model, they can achieve further speedup. Experiments showed that the proposed RRM framework speeds up the visual recognition systems YOLOv2 and rt-Pose for real-time video understanding, delivering impressive speedup without a loss in recognition accuracy.

Acknowledgements. This work is supported in part by the National Natural Science Foundation of China under Grants 61772332.

References

- [1] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. 1
- [2] J. Albericio, P. Judd, T. Hetherington, T. Aamodt, N. E. Jerger, and A. Moshovos. Cnvlutin: ineffectual-neuron-free deep neural network computing. In *Computer Architecture (ISCA), 2016 ACM/IEEE 43rd Annual International Symposium on*, pages 1–13. IEEE, 2016. 2
- [3] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 7
- [4] S. Arora, A. Bhaskara, R. Ge, and T. Ma. Provable bounds for learning some deep representations. In *International Conference on Machine Learning*, pages 584–592, 2014. 2
- [5] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime Multi-person 2D Pose Estimation Using Part Affinity Fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 7
- [6] S. Changpinyo, M. Sandler, and A. Zhmoginov. The power of sparsity in convolutional neural networks. *arXiv preprint arXiv:1702.06257*, 2017. 2
- [7] J. Charles, T. Pfister, M. Everingham, and A. Zisserman. Automatic and efficient human pose estimation for sign language videos. *International Journal of Computer Vision*, 2013. 7
- [8] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE Journal of Solid-State Circuits*, 52(1):127–138, 2017. 2
- [9] M. D. Collins and P. Kohli. Memory bounded deep convolutional networks. *arXiv preprint arXiv:1412.1442*, 2014. 2
- [10] M. Courbariaux, Y. Bengio, and J.-P. David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in Neural Information Processing Systems*, pages 3123–3131, 2015. 2
- [11] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016. 2
- [12] A. Dave, O. Russakovsky, and D. Ramanan. Predictive-corrective networks for action detection. *arXiv preprint arXiv:1704.03615*, 2017. 2
- [13] E. L. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in Neural Information Processing Systems*, pages 1269–1277, 2014. 2
- [14] X. Dong, J. Huang, Y. Yang, and S. Yan. More is less: A more complicated network with less inference complexity. *arXiv preprint arXiv:1703.08651*, 2017. 2
- [15] B. G. Fabian Caba Heilbron, Victor Escorcia and J. C. Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970, 2015. 1
- [16] H.-S. Fang, S. Xie, Y.-W. Tai, and C. Lu. RMPE: Regional multi-person pose estimation. In *ICCV*, 2017. 1
- [17] R. Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 1
- [18] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 1
- [19] B. Graham. Sparse 3D convolutional neural networks. *BMVC*, 2015. 2
- [20] B. Graham and L. van der Maaten. Submanifold sparse convolutional networks. *CoRR*, abs/1706.01307, 2017. 2
- [21] Y. Guo, A. Yao, and Y. Chen. Dynamic network surgery for efficient dnns. In *Advances In Neural Information Processing Systems*, pages 1379–1387, 2016. 2
- [22] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally. Eie: Efficient inference engine on compressed deep neural network. *SIGARCH Comput. Archit. News*, 44(3):243–254, June 2016. 2, 3, 4, 5, 6, 8
- [23] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015. 2, 5, 6
- [24] S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, pages 1135–1143, 2015. 2
- [25] S. J. Hanson and L. Y. Pratt. Comparing biases for minimal network construction with back-propagation. In *Advances in neural information processing systems*, pages 177–185, 1989. 2
- [26] B. Hassibi and D. G. Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in neural information processing systems*, pages 164–171, 1993. 2
- [27] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [28] Y. He, X. Zhang, and J. Sun. Channel pruning for accelerating very deep neural networks. *arXiv preprint arXiv:1707.06168*, 2017. 2
- [29] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 2
- [30] E. Insafutdinov, M. Andriluka, L. Pishchulin, S. Tang, E. Levinkov, B. Andres, and B. Schiele. ArtTrack: Articulated Multi-person Tracking in the Wild. In *CVPR*, 2017. 7
- [31] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014. 2
- [32] D. Kang, J. Emmons, F. Abuzaid, P. Bailis, and M. Zaharia. Optimizing deep cnn-based queries over video streams at scale. *arXiv preprint arXiv:1703.02529*, 2017. 1

- [33] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 1
- [34] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 1
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012. 5
- [36] Y. LeCun, J. S. Denker, and S. A. Solla. Optimal brain damage. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 598–605. Morgan-Kaufmann, 1990. 2
- [37] Z. Li, B. Ni, W. Zhang, X. Yang, and W. Gao. Performance guaranteed network acceleration via high-order residual quantization. *arXiv preprint arXiv:1708.08687*, 2017. 2
- [38] B. Liu, M. Wang, H. Foroosh, M. Tappen, and M. Pensky. Sparse convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 806–814, 2015. 2
- [39] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 1
- [40] J.-H. Luo, J. Wu, and W. Lin. Thinet: A filter level pruning method for deep neural network compression. *arXiv preprint arXiv:1707.06342*, 2017. 2
- [41] M. Monfort, B. Zhou, S. A. Bargal, A. Andonian, T. Yan, K. Ramakrishnan, L. Brown, Q. Fan, D. Gutfrueud, C. Vondrick, et al. Moments in time dataset: one million videos for event understanding. *arXiv preprint arXiv:1801.03150*, 2018. 1
- [42] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In J. Frnkranz and T. Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 807–814. Omnipress, 2010. 3
- [43] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. Xnornet: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pages 525–542. Springer, 2016. 2, 5, 6
- [44] B. Reagen, P. Whatmough, R. Adolf, S. Rama, H. Lee, S. K. Lee, J. M. Hernández-Lobato, G.-Y. Wei, and D. Brooks. Minerva: Enabling low-power, highly-accurate deep neural network accelerators. In *Proceedings of the 43rd International Symposium on Computer Architecture*, pages 267–278. IEEE Press, 2016. 2
- [45] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016. 1, 7
- [46] J. Redmon and A. Farhadi. Yolo9000: better, faster, stronger. *arXiv preprint arXiv:1612.08242*, 2016. 1, 7
- [47] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1
- [48] S. Shi and X. Chu. Speeding up Convolutional Neural Networks By Exploiting the Sparsity of Rectifier Units. *arXiv.org*, Apr. 2017. 2, 4
- [49] G. A. Sigurdsson, O. Russakovsky, and A. Gupta. What actions are needed for understanding human actions in videos? *arXiv preprint arXiv:1708.02696*, 2017. 1
- [50] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, 2016. 5
- [51] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [52] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1961–1970, 2016. 5
- [53] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5, 6
- [54] D. Soudry, I. Hubara, and R. Meir. Expectation backpropagation: Parameter-free training of multilayer neural networks with continuous or discrete weights. In *Advances in Neural Information Processing Systems*, pages 963–971, 2014. 2
- [55] N. Ström. Phoneme probability estimation with dynamic sparsely connected artificial neural networks. *The Free Speech Journal*, 5:1–41, 1997. 2
- [56] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 2074–2082, 2016. 2
- [57] Z. Yang, M. Moczulski, M. Denil, N. de Freitas, A. Smola, L. Song, and Z. Wang. Deep fried convnets. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1476–1483, 2015. 2
- [58] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 1
- [59] S. Zha, F. Luisier, W. Andrews, N. Srivastava, and R. Salakhutdinov. Exploiting image-trained cnn architectures for unconstrained video classification. *arXiv preprint arXiv:1503.04144*, 2015. 1
- [60] B. Zhou, A. Andonian, and A. Torralba. Temporal relational reasoning in videos. *arXiv preprint arXiv:1711.08496*, 2017. 1
- [61] X. Zhu, Y. Xiong, J. Dai, L. Yuan, and Y. Wei. Deep feature flow for video recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1