

# AdaFocus V2: End-to-End Training of Spatial Dynamic Networks for Video Recognition

Yulin Wang<sup>1\*</sup> Yang Yue<sup>1\*</sup> Yuanze Lin<sup>2</sup> Haojun Jiang<sup>1</sup> Zihang Lai<sup>3</sup>  
 Victor Kulikov<sup>4</sup> Nikita Orlov<sup>4</sup> Humphrey Shi<sup>4,5</sup> Gao Huang<sup>1</sup>  
<sup>1</sup>BNRist, Tsinghua University <sup>2</sup>U of Washington <sup>3</sup>CMU  
<sup>4</sup>Picsart AI Research (PAIR) <sup>5</sup>U of Oregon

## Abstract

Recent works have shown that the computational efficiency of video recognition can be significantly improved by reducing the spatial redundancy. As a representative work, the adaptive focus method (**AdaFocus**) has achieved a favorable trade-off between accuracy and inference speed by dynamically identifying and attending to the informative regions in each video frame. However, **AdaFocus** requires a complicated three-stage training pipeline (involving reinforcement learning), leading to slow convergence and is unfriendly to practitioners. This work reformulates the training of **AdaFocus** as a simple one-stage algorithm by introducing a differentiable interpolation-based patch selection operation, enabling efficient end-to-end optimization. We further present an improved training scheme to address the issues introduced by the one-stage formulation, including the lack of supervision, input diversity and training stability. Moreover, a conditional-exit technique is proposed to perform temporal adaptive computation on top of **AdaFocus** without additional training. Extensive experiments on six benchmark datasets (i.e., ActivityNet, FCVID, Mini-Kinetics, Something-Something V1&V2, and Jester) demonstrate that our model significantly outperforms the original **AdaFocus** and other competitive baselines, while being considerably more simple and efficient to train. Code is available at <https://github.com/LeapLabTHU/AdaFocusV2>.

## 1. Introduction

Deep networks have achieved remarkable success in large-scale video recognition tasks [3, 14, 16, 25, 58]. Their high accuracy has fueled the desire to deploy them for automatically recognizing the actions, events, or other contents within the explosively growing online videos in recent years (e.g., on YouTube). However, the models with state-of-the-

Table 1. A comparison of training the original **AdaFocus** model (**AdaFocusV1**, from ❶ to ❸) and **AdaFocusV2** (end-to-end) on **Something-Something (Sth-Sth) V1** dataset. Both procedures start from the same initial backbone networks. Herein,  $f_G$ ,  $f_L$ ,  $f_C$  and  $\pi$  are the model components (see Section 3.1 for details).

	AdaFocusV1	AdaFocusV2
<b>Pre-training</b>	❶ Pre-train $f_G$ on Sth-Sth V1. ❷ Pre-train $f_L$ on Sth-Sth V1.	<b>End-to-End Training</b> ( $f_G$ , $f_L$ , $f_C$ , $\pi$ )
<b>Stage-1</b>	❸ Train $f_L$ and $f_C$ using random patches.	
<b>Stage-2</b>	❹ Train $\pi$ using reinforcement learning.	
<b>Stage-3</b>	❺ Fine-tune $f_L$ and $f_C$ .	




Figure 1. **Comparisons of AdaFocusV1 and AdaFocusV2 on Sth-Sth V1 in terms of accuracy v.s. training cost.** The training time is measured based on 4 NVIDIA 3090 GPUs. The two sides of grey arrows correspond to the same network architecture (i.e., the same inference cost). Our **AdaFocusV2** accelerates the training by **2.2-2.4×**, while boosting the accuracy by **1.0-1.5%**.

art performance [1, 13, 20, 42, 51, 59] tend to be computationally intensive during inference. In real-world applications such as recommendation [8, 9, 18], surveillance [4, 7] and content-based searching [30], computation translates into power consumption and latency, both of which should be minimized for environmental, safety or economic reasons.

Several algorithms have been proposed to reduce the *temporal redundancy* of videos [19, 21, 35, 36, 47, 57, 67, 68, 73] by allocating the majority of computation to the

\*Equal contribution.

✉Corresponding authors.

most task-relevant video frames rather than all. Orthogonal to these approaches, the recently proposed adaptive focus network (AdaFocus) [64] reveals that reducing the *spatial redundancy* in video analysis yields promising results for efficient video recognition. The AdaFocus framework is also compatible with the aforementioned temporal-adaptive methods to realize highly efficient spatial-temporal computation. Specifically, AdaFocus reduces the computational cost by applying the expensive high-capacity network only on some relatively small patches. These patches are strategically selected to capture the most informative regions of each video frame. In particular, the patch localization task is formulated as a non-differentiable *discrete* decision task, which is further solved with reinforcement learning. As a consequence, AdaFocus needs to be trained with a complicated three-stage training pipeline (see Table 1), resulting in long training time and being unfriendly to users.

This paper seeks to simplify the training process of AdaFocus. We first introduce a *differentiable* interpolation-based formulation for patch selecting, allowing gradient back-propagation throughout the whole model. We note that a straightforward implementation of end-to-end training leads to optimization issues, including the lack of supervision, input diversity and training stability, which severely degrade the performance. Therefore, we further propose three tailored training techniques: auxiliary supervision, diversity augmentation and stopping gradient, to address the aforementioned issues. These simple but effective techniques enable the simple one-stage formulation of our AdaFocus algorithm to be trained effectively, and eventually outperform the three-stage counterparts in terms of both test accuracy and training cost. The experimental comparisons are presented in Figure 1. Our proposed method is referred to as AdaFocusV2.

An additional advantage of AdaFocus is that it can be easily improved by further considering *temporal redundancy*. The original paper implements this idea by dynamically skipping less valuable frames with reinforcement learning. In contrast, this work proposes a simplified early-exit algorithm that removes the requirement of introducing additional training, but achieves competitive performance.

The effectiveness of AdaFocusV2 is extensively evaluated on six video recognition benchmarks (*i.e.*, ActivityNet, FCVID, Mini-Kinetics, Sth-Sth V1&V2, and Jester). Experimental results show that the training of AdaFocusV2 is  $2\times$  faster (measured in wall-time) than the original counterpart, while achieving consistently higher accuracy.

## 2. Related Works

**Video recognition.** In recent years, convolutional networks (ConvNets) have achieved remarkable performance for video recognition. One of the prominent approaches is to capture the spatial/temporal information jointly us-

ing 3D ConvNets. Representative works include C3D [58], I3D [3], ResNet3D [25], X3D [13], etc. Some other works focus on first extracting frame-wise features, and then aggregating temporal information with specialized architectures, such as temporal averaging [63], deploying recurrent networks [10, 38, 74], and temporal channel shift [12, 40, 48, 56]. Another line of works leverage two-stream architectures to model short-term and long-term temporal relationships respectively [14–16, 22]. In addition, as processing videos with ConvNets, especially 3D ConvNets, tends to be computationally intensive, recent research starts to pay attention to designing efficient video recognition models [41, 43, 44, 59, 60, 77].

**Temporal redundancy.** A popular approach for facilitating efficient video recognition is to reduce the temporal redundancy in videos [19, 21, 35, 36, 47, 57, 67, 68, 73]. Since not all frames are equally important for a given task, the model should ideally allocate less computation on less informative frames [24]. Several effective algorithms have been proposed in this direction. For example, VideoIQ [57] processes video frames using different precision according to their relative importance. FrameExit [21] learns to terminate the inference process after seeing a few sufficiently informative frames. The AdaFocusV2+ algorithm, which we propose to model temporal redundancy on top of AdaFocusV2, is related to FrameExit on the spirit of early-exit. However, our method is easier to implement since it does not need to learn an additional conditional-exit policy.

**Spatial-wise dynamic networks** perform computation adaptively on top of different spatial locations of the inputs [24]. The AdaFocusV2 network studied in this paper can be classified into this category as well. Many of the spatially adaptive networks are designed from the lens of inference efficiency [5, 24, 52, 62, 72]. For example, recent works have revealed that 2D images can be efficiently processed via attending to the task-relevant or more informative image regions [17, 61, 66, 70, 71]. In the context of video recognition, how to exploit this spatial redundancy for reducing the computational cost is still an under-explored topic. It has been shown by the attention-based methods [46] that the contributions of different frame regions to the recognition task are not equivalent. Preliminary works like AdaFocus [64] have demonstrated the potentials of this direction.

The spatial transformer networks [31] are trained based on a similar interpolation-based mechanism to us. However, they focus on actively transforming the feature maps for learning spatially invariant representations, while we aim to localize and attend to the task-relevant regions of the video frames for improving the inference efficiency. Moreover, we show that a straightforward implementation of this mechanism fails to yield competitive results in our problem. Special designs need to be introduced by our algorithm to solve the optimization difficulties.

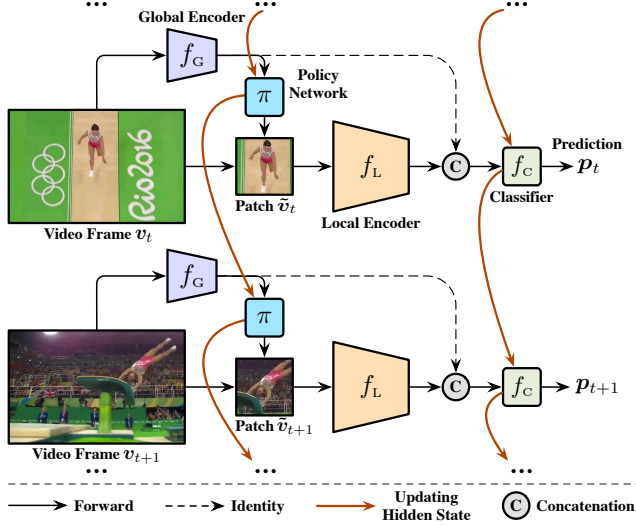


Figure 2. **An overview of adaptive focus networks (AdaFocus).** The global and local encoders  $f_G$  and  $f_L$  are two deep networks (e.g., ConvNets). The former catches a glimpse of each frame, enabling the recurrent policy network  $\pi$  to localize an informative image patch  $\tilde{v}_t$ . The latter is designed to extract discriminative representations by processing  $\tilde{v}_t$  with its large and high-capacity architecture. A classifier  $f_C$  aggregates the information from all the processed frames to obtain the prediction.

### 3. Method

The AdaFocus network [64] facilitates efficient video recognition via reducing the spatial-wise redundant computation. However, it suffers from a complicated three-stage training procedure. This section introduces the details of our end-to-end trainable AdaFocusV2 approach, which consistently outperforms the original AdaFocus with a simpler and more efficient training procedure.

#### 3.1. Preliminaries of AdaFocusV1 [64]

We start by giving an overview of AdaFocus (see Figure 2), laying the basis for the discussions on training efficiency. Assume that a stream of video frames  $\{v_1, v_2, \dots\}$  comes in sequentially. AdaFocus first takes a quick glance at each frame with a light-weighted global encoder  $f_G$ , aiming to extract cheap and coarse global features

$$e_t^G = f_G(v_t), \quad t = 1, 2, \dots, \quad (1)$$

where  $e_t^G$  denotes the feature maps of the  $t^{\text{th}}$  frame. Then a recurrent policy network  $\pi$  is learned to aggregate the features of all previous frames  $\{e_1^G, \dots, e_t^G\}$ , and accordingly determine the location of a small image patch  $\tilde{v}_t$  to capture the most informative region in  $v_t$  for a given task.

The selected patch  $\tilde{v}_t$  will be fed into a high-capacity, accurate but computationally more expensive local encoder  $f_L$  to extract the local feature maps  $e_t^L$ :

$$e_t^L = f_L(\tilde{v}_t), \quad t = 1, 2, \dots \quad (2)$$

Importantly, the computational cost introduced by Eq.(2) is considerably smaller than activating  $f_L$  for processing the whole frame due to the reduced size of  $\tilde{v}_t$ . AdaFocus is designed to unevenly allocate the computation across the spatial dimension to improve inference efficiency.

Finally, a classifier  $f_C$  integrates the information from all the processed frames, and produces the softmax prediction  $p_t$  at  $t^{\text{th}}$  step, i.e.,

$$p_t = f_C(\text{cat}(e_1^G, e_1^L), \dots, \text{cat}(e_t^G, e_t^L)), \quad (3)$$

where  $\text{cat}(\cdot)$  is the concatenation operation<sup>1</sup>. Note that  $e_t^G$  is leveraged for both localizing the informative patches and classification, under the goal of facilitating efficient feature reuse. This design is natural since it has been observed that deep networks (e.g., ConvNets and Vision Transformers) excel at learning representations for both recognition and localization simultaneously [11, 55, 76]. In addition, the architecture of  $f_C$  may have different choices, such as recurrent networks [6, 27], averaging the frame-wise predictions [40, 47, 48], and accumulated feature pooling [21].

**Training algorithm.** In the original AdaFocus [64] (referred to as AdaFocusV1), the task of selecting task-relevant patches is modeled as a non-differentiable decision problem on several pre-defined patch candidates. Hence, the training of AdaFocusV1 includes both continuous (i.e., video recognition) and discrete (i.e., patch localization) optimization, resulting in a 3-stage algorithm to solve it alternatively. They first train the classification components (i.e.,  $f_G$ ,  $f_L$  and  $f_C$ ) with random patches, and then fix them to learn a patch selection strategy (i.e.,  $\pi$ ) using reinforcement learning. The final stage further fine-tunes the model with the learned policy.

**Limitations of AdaFocusV1.** The underlying logic of the 3-stage training is straightforward. However, this procedure is unfriendly for practitioners. First, effectively deploying the reinforcement learning algorithm is nontrivial. It requires considerable efforts for properly designing the key components (e.g., the action space and the reward function), and implementing specialized optimization techniques (e.g., deep Q-Network [49] or proximal policy optimization [54]). Second, the 3-stage alternative algorithm is an indirect formulation for optimizing the recognition objective, which tends to be time-consuming, and may result in sub-optimal solutions. Third, the performance of AdaFocusV1 largely depends on a number of implementation configurations (e.g., performing pre-training, freezing some components in different stages, and stage-wise hyperparameter searching) that need to be carefully tuned on a per-dataset or per-backbone basis.

In the following, we present an end-to-end trainable formulation for AdaFocus to address the issue of inefficient

<sup>1</sup>Typically, global average pooling will be performed on  $e_t^G, e_t^L$  before concatenation to obtain feature vectors. Here we omit this for simplicity.




Figure 3. **Illustration of interpolation-based patch selection.** This operation is differentiable, *i.e.*, the gradients can be directly back-propagated into the policy network  $\pi$  through the selected image patch  $\tilde{v}_t$ . Consequently, integrating the learning of  $\pi$  into a unified end-to-end training paradigm turns out to be feasible.

training. The proposed network, AdaFocusV2, can be conveniently implemented to achieve consistently better performance than AdaFocusV1 with reduced training cost.

### 3.2. Interpolation-based Patch Selection

To enable end-to-end training, we propose a differentiable solution to obtain  $\tilde{v}_t$ . Suppose that the size of the original frame  $v_t$  and the patch  $\tilde{v}_t$  is  $H \times W$  and  $P \times P$  ( $P < H, W$ ), respectively<sup>2</sup>. We assume that  $\pi$  outputs the continuous centre coordinates  $(\tilde{x}_c^t, \tilde{y}_c^t)$  of  $\tilde{v}_t$ , namely

$$\begin{aligned} (\tilde{x}_c^t, \tilde{y}_c^t) &= \pi(e_1^G, \dots, e_t^G), \\ \tilde{x}_c^t &\in \left[\frac{P}{2}, W - \frac{P}{2}\right], \quad \tilde{y}_c^t \in \left[\frac{P}{2}, H - \frac{P}{2}\right], \end{aligned} \quad (4)$$

where  $e_1^G, \dots, e_t^G$  are the global features of 1<sup>st</sup> –  $t^{\text{th}}$  frames extracted by the global encoder  $f_G$ . Notably, we refer to the coordinates of the top-left corner of the frame as  $(0, 0)$ , and Eq. (4) ensures that  $\tilde{v}_t$  will never go outside of  $v_t$ . Our aim is to calculate the values of all pixels in  $\tilde{v}_t$ , while allowing the gradients to be back-propagated through  $(\tilde{x}_c^t, \tilde{y}_c^t)$ .

**Feed-forward.** We first introduce the feed-forward process of our method. Formally, the coordinates of a pixel in the patch  $\tilde{v}_t$  can be expressed as the addition of  $(\tilde{x}_c^t, \tilde{y}_c^t)$  and a fixed offset:

$$\begin{aligned} (\tilde{x}_{ij}^t, \tilde{y}_{ij}^t) &= (\tilde{x}_c^t, \tilde{y}_c^t) + \mathbf{o}_{ij}, \\ \mathbf{o}_{ij} &\in \left\{ -\frac{P}{2}, -\frac{P}{2} + 1, \dots, \frac{P}{2} \right\}^2. \end{aligned} \quad (5)$$

Herein,  $(\tilde{x}_{ij}^t, \tilde{y}_{ij}^t)$  denotes the horizontal and vertical coordinates in the original frame  $v_t$  corresponding to the pixel in the  $i^{\text{th}}$  row and  $j^{\text{th}}$  column of  $\tilde{v}_t$ , while  $\mathbf{o}_{ij}$  represents the

<sup>2</sup>In our implementation, the height/width/coordinates are correspondingly normalized using the linear projection  $[0, H] \rightarrow [0, 1]$  and  $[0, W] \rightarrow [0, 1]$ . Here we use the original values for the ease of understanding.

vector from the patch centre  $(\tilde{x}_c^t, \tilde{y}_c^t)$  to this pixel. Given a fixed patch size,  $\mathbf{o}_{ij}$  is a constant conditioned only on  $i, j$ , regardless of  $t$  or the inputs of  $\pi$ .

Since the values of  $(\tilde{x}_c^t, \tilde{y}_c^t)$  are continuous, there does not exist a pixel of  $v_t$  exactly located at  $(\tilde{x}_{ij}^t, \tilde{y}_{ij}^t)$  to directly get the pixel value. Alternatively, as illustrated in Figure 3, we can always find that the location  $(\tilde{x}_{ij}^t, \tilde{y}_{ij}^t)$  is surrounded by four adjacent pixels of  $v_t$ , forming a grid. The coordinates are  $(\lfloor \tilde{x}_{ij}^t \rfloor, \lfloor \tilde{y}_{ij}^t \rfloor)$ ,  $(\lfloor \tilde{x}_{ij}^t \rfloor + 1, \lfloor \tilde{y}_{ij}^t \rfloor)$  and  $(\lfloor \tilde{x}_{ij}^t \rfloor + 1, \lfloor \tilde{y}_{ij}^t \rfloor + 1)$ , respectively, where  $\lfloor \cdot \rfloor$  denotes the rounding-down operation. By assuming that the corresponding pixel values of these four pixels are  $(m_{ij}^t)_{00}$ ,  $(m_{ij}^t)_{01}$ ,  $(m_{ij}^t)_{10}$ , and  $(m_{ij}^t)_{11}$ , the pixel value at  $(\tilde{x}_{ij}^t, \tilde{y}_{ij}^t)$  (referred to as  $\tilde{m}_{ij}^t$ ) can be obtained via interpolation algorithms. In this paper, we simply adopt the differentiable bilinear interpolation:

$$\begin{aligned} \tilde{m}_{ij}^t &= (m_{ij}^t)_{00}(\lfloor \tilde{x}_{ij}^t \rfloor - \tilde{x}_{ij}^t + 1)(\lfloor \tilde{y}_{ij}^t \rfloor - \tilde{y}_{ij}^t + 1) \\ &\quad + (m_{ij}^t)_{01}(\tilde{x}_{ij}^t - \lfloor \tilde{x}_{ij}^t \rfloor)(\lfloor \tilde{y}_{ij}^t \rfloor - \tilde{y}_{ij}^t + 1) \\ &\quad + (m_{ij}^t)_{10}(\lfloor \tilde{x}_{ij}^t \rfloor - \tilde{x}_{ij}^t + 1)(\tilde{y}_{ij}^t - \lfloor \tilde{y}_{ij}^t \rfloor) \\ &\quad + (m_{ij}^t)_{11}(\tilde{x}_{ij}^t - \lfloor \tilde{x}_{ij}^t \rfloor)(\tilde{y}_{ij}^t - \lfloor \tilde{y}_{ij}^t \rfloor). \end{aligned} \quad (6)$$

Consequently, we can obtain the image patch  $\tilde{v}_t$  by traversing all possible  $i, j$  in Eq. (6).

**Back-propagation.** Give the training loss  $\mathcal{L}$ , it is easy to compute the gradient  $\partial \mathcal{L} / \partial \tilde{m}_{ij}^t$  with standard back-propagation. Then, following on the chain rule, we have

$$\frac{\partial \mathcal{L}}{\partial \tilde{x}_c^t} = \sum_{i,j} \frac{\partial \mathcal{L}}{\partial \tilde{m}_{ij}^t} \frac{\partial \tilde{m}_{ij}^t}{\partial \tilde{x}_c^t}, \quad \frac{\partial \mathcal{L}}{\partial \tilde{y}_c^t} = \sum_{i,j} \frac{\partial \mathcal{L}}{\partial \tilde{m}_{ij}^t} \frac{\partial \tilde{m}_{ij}^t}{\partial \tilde{y}_c^t}. \quad (7)$$

Combining Eq. (5) and Eq. (7), we can further derive

$$\frac{\partial \tilde{m}_{ij}^t}{\partial \tilde{x}_c^t} = \frac{\partial \tilde{m}_{ij}^t}{\partial \tilde{x}_{ij}^t}, \quad \frac{\partial \tilde{m}_{ij}^t}{\partial \tilde{y}_c^t} = \frac{\partial \tilde{m}_{ij}^t}{\partial \tilde{y}_{ij}^t}. \quad (8)$$

Eq. (8) can be solved by leveraging Eq. (6), such that we can obtain the gradients  $\partial \mathcal{L} / \partial \tilde{x}_c^t$  and  $\partial \mathcal{L} / \partial \tilde{y}_c^t$ . Given that  $\tilde{x}_c^t$  and  $\tilde{y}_c^t$  are the outputs of the policy network  $\pi$ , the back-propagation process is able to proceed in an ordinary way.

### 3.3. Training Techniques

**Naive implementation.** Thus far, we have enabled the gradients to be back-propagated throughout the whole AdaFocus network for updating all trainable parameters simultaneously. Consequently, end-to-end training has been feasible. For example, one can minimize the frame-wise cross-entropy loss  $L_{CE}(\cdot)$  in AdaFocusV1 [64]:

$$\underset{f_G, f_L, f_C, \pi}{\text{minimize}} \quad \mathcal{L} = \mathbb{E}_{\{v_1, v_2, \dots\}} \left[ \frac{1}{T} \sum_{t=1}^T L_{CE}(\mathbf{p}_t, y) \right], \quad (9)$$

where  $T$  and  $y$  denote the length and the label of the video  $\{v_1, v_2, \dots\}$ , and  $\mathbf{p}_t$  is the softmax prediction at  $t^{\text{th}}$  frame, as stated in Section 3.1.



However, importantly, such a straightforward implementation leads to the severely degraded performance (see Table 6 for experimental evidence). We attribute this issue to the absence of some appealing optimization properties introduced by 3-stage training, namely the lack of *supervision*, *input diversity* and *training stability*. To solve these problems, we develop three simple but effective training techniques, with which end-to-end training can significantly outperform the 3-stage counterpart. These techniques do not introduce additional tunable hyper-parameters, while achieving consistent improvements across varying datasets, backbone architectures, and model configurations.

**Lack of supervision: auxiliary supervision.** AdaFocusV1 first pre-trains the two encoders (*i.e.*,  $f_G, f_L$ ) separately using a direct frame-wise recognition loss (by simply appending a fully-connected layer), aiming to obtain a proper initialization. In contrast, when solving problem (9),  $f_G$  and  $f_L$  are trained from scratch, while they are only indirectly supervised by the gradients from the classifier  $f_C$ . We find that explicitly introducing auxiliary supervision on  $f_G$  and  $f_L$  effectively facilitates the efficient end-to-end training of AdaFocusV2. In specific, we attach two linear classifiers,  $FC_G(\cdot)$  and  $FC_L(\cdot)$ , to the outputs of  $f_G$  and  $f_L$ , and replace the loss function  $\mathcal{L}$  in (9) by  $\mathcal{L}'$ :

$$\mathcal{L}' = \mathbb{E}_{\{v_1, v_2, \dots\}} \left\{ \frac{1}{T} \sum_{t=1}^T [L_{CE}(\mathbf{p}_t, y) + L_{CE}(\text{SoftMax}(FC_G(\bar{\mathbf{e}}_t^G)), y) + L_{CE}(\text{SoftMax}(FC_L(\bar{\mathbf{e}}_t^L)), y)] \right\}, \quad (10)$$

where  $\bar{\mathbf{e}}_t^G$  and  $\bar{\mathbf{e}}_t^L$  are the feature vectors after performing global average pooling on the feature maps  $\mathbf{e}_t^G$  and  $\mathbf{e}_t^L$  output by  $f_G$  and  $f_L$ . Intuitively, minimizing  $\mathcal{L}'$  enforces the two encoders to learn linearized deep representations, which has been widely verified as an efficient approach for training deep networks [11, 26, 29]. This paradigm benefits the learning of  $f_C$  as well, since its inputs are explicitly regularized to be linearly separable.

**Lack of input diversity: diversity augmentation.** In the stage-1 for training AdaFocusV1, image patches are randomly cropped, yielding highly diversified inputs for learning well-generalized local encoder  $f_L$ . However, the patch selection process presented in Section 3.2 is deterministic. In Eq. (10), given a video frame, the local encoder  $f_L$  only has access to the patch specified by the policy network  $\pi$ . This procedure leads to the limited diversity of training data for the inputs of  $f_L$ . Empirically, we observe that it results in the inferior performance of  $f_L$ . We address this issue by proposing a straightforward diversity augmentation approach. For each video, we first compute  $\mathcal{L}'$  by activating  $\pi$  as aforementioned. Then we infer  $f_L$  and the classifier  $f_C$  for a second time using randomly cropped patches, obtaining an additional loss  $\mathcal{L}'_{\text{random}}$ , which follows Eq. (10)




Figure 4. Illustration of AdaFocusV2+.

as well. Our final optimization objective is to minimize the combination of  $\mathcal{L}'$  and  $\mathcal{L}'_{\text{random}}$ :

$$\underset{f_G, f_L, f_C, \pi}{\text{minimize}} \quad \frac{1}{2} (\mathcal{L}' + \mathcal{L}'_{\text{random}}). \quad (11)$$

**Lack of training stability: stop-gradient.** In AdaFocusV1, the policy network  $\pi$  is learned on top of the fixed and completely trained global encoder  $f_G$ . When it comes to end-to-end training,  $\pi$  and  $f_G$  are simultaneously updated. In this case, we observe that the gradients back-propagated from  $\pi$  interfere with the learning of  $f_G$ , leading to an unstable training process with slow convergence speed. We find that this problem can be solved by simply stopping the gradients before the inputs of  $\pi$ . In other words, we propose to train  $f_G$  using the pure classification objective without any effect from  $\pi$ , as done in AdaFocusV1. This design is rational since previous works have revealed that the representations extracted by deep recognition networks can naturally be leveraged for localizing task-relevant regions [11, 55, 76].

### 3.4. Reducing Temporal Redundancy

The basic formulation of AdaFocus processes each frame using the same amount of computation, and hence it can be improved by further reducing temporal redundancy. AdaFocusV1 achieves this via skipping less informative frames with reinforcement learning. In contrast, we propose a simple confidence-based early-exit algorithm that achieves competitive performance. Our approach can be directly deployed on AdaFocusV2 trained following the aforementioned paradigm, without any additional training process. We refer to this extension of AdaFocusV2 as AdaFocusV2+, as shown in Figure 4.

It has been widely observed that there exist a considerable number of “easier” samples in datasets [21, 28, 37, 65, 66, 72], which can be accurately recognized with smaller computational cost than others. In the context of videos, we assume that processing a subset of frames (from the beginning) rather than all may be sufficient for these “easier” samples. To implement this idea, at test time, we propose to compare the largest entry of the softmax prediction  $\mathbf{p}_t$  (defined as confidence in previous works [28, 65, 66, 72]) at  $t^{\text{th}}$  frame with a pre-defined threshold  $\eta_t$ . Once  $\max_j p_{tj} \geq \eta_t$ , the prediction will be postulated to be reliable enough, and the inference will be terminated by outputting  $\mathbf{p}_t$ . We always adopt a zero-threshold for at final frame.

The values of  $\{\eta_1, \eta_2, \dots\}$  are solved on the validation set. Suppose that the model needs to classify a set of

Table 2. **Comparisons of AdaFocusV1 and AdaFocusV2 on training efficiency.** ActivityNet mAP and wall-clock training time are reported. The latter is obtained using 4 NVIDIA 3090 GPUs. The better results are **bold-faced** (E2E: End-to-End).

Methods	3-stage Training	mAP on ActivityNet			Training Wall-time		
		96 <sup>2</sup>	128 <sup>2</sup>	160 <sup>2</sup>	96 <sup>2</sup>	128 <sup>2</sup>	160 <sup>2</sup>
AdaFocusV1-MN2/RN	✓	71.9%	75.0%	76.0%	6.4h	7.2h	8.6h
AdaFocusV2-MN2/RN	✗	<b>77.4%</b>	<b>79.0%</b>	<b>79.9%</b>	<b>3.4h</b>	<b>3.7h</b>	<b>4.3h</b>

Table 3. **AdaFocusV1+ v.s. AdaFocusV2+.** The two methods are deployed on the same AdaFocusV1 base model to reduce temporal redundancy. ActivityNet mAP on varying inference cost (GFLOPs/video) is reported. The better results are **bold-faced** (RL: reinforcement learning).

Methods	Additional Training	On AdaFocusV1 (128 <sup>2</sup> )		On AdaFocusV1 (160 <sup>2</sup> )	
		19.6G/video	24.4G/video	27.2G/video	34.6G/video
AdaFocusV1-MN2/RN+	✓(RL)	74.2%	74.8%	75.5%	75.9%
AdaFocusV2-MN2/RN+	✗	<b>74.6%</b>	<b>75.0%</b>	75.5%	75.9%

Table 4. **Comparisons of AdaFocusV2 and state-of-the-art baselines on ActivityNet, FCVID and Mini-Kinetics.** GFLOPs refers to the average computational cost for processing a single video. The best two results are **bold-faced** and underlined, respectively.

Methods	Published on	Backbones	ActivityNet		FCVID		Mini-Kinetics	
			mAP	GFLOPs	mAP	GFLOPs	Top-1 Acc.	GFLOPs
AdaFuse [48]	<i>ICLR'21</i>	ResNet	73.1%	61.4	<u>81.6%</u>	45.0	72.3%	23.0
Dynamic-STE [35]	<i>ICCV'21</i>	ResNet	75.9%	30.5	—	—	72.7%	<u>18.3</u>
FrameExit [21]	<i>CVPR'21</i>	ResNet	<u>76.1%</u>	<u>26.1</u>	—	—	72.8%	19.7
AdaFocusV2-RN (128 <sup>2</sup> )	—	ResNet	<b>78.9%</b>	34.1	<b>84.5%</b>	<u>34.1</u>	<b>74.0%</b>	34.1
AdaFocusV2-RN+ (128 <sup>2</sup> )	—	ResNet	<u>76.1%</u>	<b>15.3</b>	<u>81.6%</u>	<b>12.0</b>	<u>72.8%</u>	<b>13.4</b>
LiteEval [69]	<i>NeurIPS'19</i>	MobileNet-V2 + ResNet	72.7%	95.1	80.0%	94.3	61.0%	99.0
SCSampler [36]	<i>ICCV'19</i>	MobileNet-V2 + ResNet	72.9%	42.0	81.0%	42.0	70.8%	42.0
ListenToLook [19]	<i>CVPR'20</i>	MobileNet-V2 + ResNet	72.3%	81.4	—	—	—	—
AR-Net [47]	<i>ECCV'20</i>	MobileNet-V2 + ResNet	73.8%	33.5	81.3%	35.1	71.7%	32.0
AdaFrame [68]	<i>T-PAMI'21</i>	MobileNet-V2 + ResNet	71.5%	79.0	80.2%	75.1	—	—
VideoIQ [57]	<i>ICCV'21</i>	MobileNet-V2 + ResNet	<u>74.8%</u>	28.1	<u>82.7%</u>	<u>27.0</u>	<u>72.3%</u>	<u>20.4</u>
AdaFocusV2-MN2/RN (128 <sup>2</sup> )	—	MobileNet-V2 + ResNet	<b>79.0%</b>	<u>27.0</u>	<b>85.0%</b>	<u>27.0</u>	<b>75.4%</b>	27.0
AdaFocusV2-MN2/RN+ (128 <sup>2</sup> )	—	MobileNet-V2 + ResNet	<u>74.8%</u>	<b>9.9</b>	<u>82.7%</u>	<b>10.1</b>	<u>72.3%</u>	<b>6.3</b>

samples  $\mathcal{D}_{\text{val}}$  within a given computational budget  $B > 0$  [28, 65]. One can obtain the thresholds through

$$\begin{aligned} & \underset{\eta_1, \eta_2, \dots}{\text{maximize}} \quad \text{Acc}(\eta_1, \eta_2, \dots | \mathcal{D}_{\text{val}}) \\ & \text{subject to} \quad \text{FLOPs}(\eta_1, \eta_2, \dots | \mathcal{D}_{\text{val}}) \leq B. \end{aligned} \quad (12)$$

Here  $\text{Acc}(\eta_1, \eta_2, \dots | \mathcal{D}_{\text{val}})$  and  $\text{FLOPs}(\eta_1, \eta_2, \dots | \mathcal{D}_{\text{val}})$  refer to the accuracy and computational cost on  $\mathcal{D}_{\text{val}}$  using the thresholds  $\{\eta_1, \eta_2, \dots\}$ . Notably, by changing  $B$ , one can obtain varying values of  $\{\eta_1, \eta_2, \dots\}$ . The computational cost of AdaFocusV2+ can be flexibly adjusted without additional training by simply adjusting these thresholds. In our implementation, we solve problem (12) following the method proposed in [28] on training set, which we find performs on par with using cross-validation.

## 4. Experiment

In this section, we empirically compare AdaFocusV2 with both AdaFocusV1 and state-of-the-art efficient video recognition frameworks. We also demonstrate that AdaFocusV2 is able to effectively reduce the inference cost on top of recent representative light-weighted deep networks. In addition, ablation studies and visualization results are provided for more insights. Code and pre-trained models are available at <https://github.com/LeapLabTHU/AdaFocusV2>.

**Datasets.** Six large-scale video recognition datasets are used, *i.e.*, ActivityNet [2], FCVID [33], Mini-Kinetics [34, 47], Something-Something (Sth-Sth) V1&V2 [23] and Jester [45]. The official training-validation split is adopted. Given the limited space, we introduce them in Appendix A.

Following the common practice [19, 40, 47, 48, 64, 68, 69], we evaluate the performance of different methods via mean average precision (mAP) and Top-1 accuracy (Top-1 Acc.) on ActivityNet/FCVID and other datasets, respectively.

**Setups.** Unless otherwise specified, we uniformly sample 16 frames from each video on ActivityNet, FCVID and Mini-Kinetics, while sampling 8/12 frames on Sth-Sth. Offline recognition is considered, where we obtain a single prediction for each video. We adopt the final prediction  $p_T$  after processing all frames for AdaFocusV2, and use the early-exited prediction for AdaFocusV2+. We follow the data pre-processing pipeline in [40, 47, 64]. For inference, we resize all frames to  $256^2$  and perform  $224^2$  centre-crop.

### 4.1. Comparisons with State-of-the-art Baselines

**Baselines.** In this subsection, we compare AdaFocusV2 with a number of competitive baselines on ActivityNet, FCVID and Mini-Kinetics. AdaFocusV1 and several state-of-the-art frameworks for efficient video recognition are considered, including MultiAgent [67], LiteEval [69], SC-Sampler [36], ListenToLook [19], AR-Net [47], AdaFrame [68], AdaFuse [48], VideoIQ [57], Dynamic-STE [35] and FrameExit [21]. An introduction is given in Appendix A.

**Implementation details.** The comparisons are presented on top of the same backbone networks. Following AdaFocusV1, in most cases, we adopt MobileNet-V2 (MN2) [53] and ResNet-50 (RN) [26] as the global encoder  $f_G$  and local encoder  $f_L$  of AdaFocusV2. When comparing with the methods that only leverage ResNet (*i.e.*, FrameExit [21], Dynamic-STE [35] and AdaFuse [48]), we replace the MobileNet-V2 with a ResNet-50 with down-sampled inputs (96<sup>2</sup>). The two variants are referred to as AdaFocusV2-

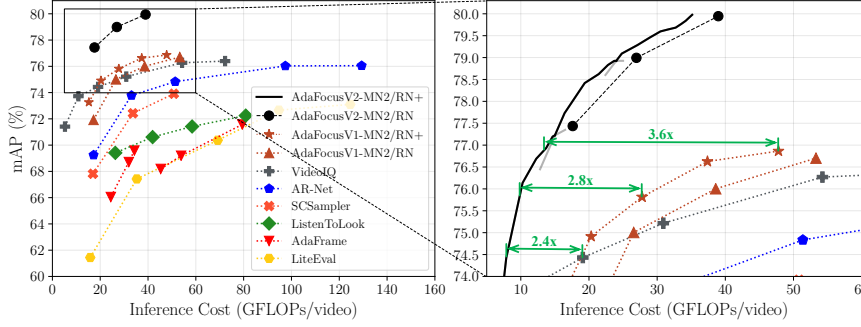


Figure 5. **Comparisons of AdaFocusV2, AdaFocusV1 and state-of-the-art baselines on ActivityNet in terms of inference efficiency.** MobileNetV2 (MN2) and ResNet (RN) are deployed as backbones in all methods. AdaFocusV2 adopts  $96^2$ ,  $128^2$ , and  $160^2$  patches. Notably, AdaFocusV2+ can switch within each black curve without additional training.




Figure 6. **AdaFocusV2 ( $96^2$ ,  $128^2$ ,  $160^2$ ) v.s. state-of-the-art efficient video recognition frameworks on top of ResNet (RN).** Results on ActivityNet are reported.




Figure 7. **Comparisons of training efficiency of AdaFocusV1 and AdaFocusV2 on Sth-Sth V2.** The wall-clock training time is obtained based on 4 NVIDIA 3090 GPUs.





Figure 8. **AdaFocusV1 v.s. AdaFocusV2 on Sth-Sth V1 (left) and Sth-Sth V2 (right) in terms of inference efficiency.** The two algorithms are implemented on top of TSM [40] with the patch size of  $\{144^2, 160^2, 176^2\}$  and  $\{128^2, 144^2, 160^2, 176^2\}$ , respectively. AdaFocusV2 outperforms AdaFocusV1 significantly with the same network architecture.



MN2/RN and AdaFocusV2-RN, respectively. A one-layer gated recurrent unit (GRU) [6] with a hidden size of 2048 is used as the main body of the policy network  $\pi$ , while the accumulated max-pooling module [21] is deployed as the classifier  $f_C$ . The input frames are re-ordered following the sampling policy in [21] (see Table 6 for ablation). Due to spatial limitations, more details on network architecture and training hyper-parameters are deferred to Appendix B.

**AdaFocusV2 v.s. AdaFocusV1.** The comparisons of AdaFocusV2 and AdaFocusV1 in terms of mean average precision (mAP) v.s. training/inference cost are presented in Table 2 and Figure 5, respectively. Our method is implemented with the patch size of  $96^2$ ,  $128^2$ , and  $160^2$ . One can observe that AdaFocusV2 reduces the time consumption for training by  $\sim 2\times$ , while dramatically improving the mAP using the same patch size (by 4-5%).

**Performance of AdaFocusV2+** is depicted in Figures 5, 6. The curves correspond to reducing temporal redundancy on top of AdaFocusV2 ( $96^2$ ,  $128^2$ ,  $160^2$ ), i.e., the three black dots. As stated in Section 3.4, we vary the average computational budget, solve the confidence thresholds, and evaluate the corresponding validation accuracy. One can observe that AdaFocusV2+ effectively improves the inference efficiency of AdaFocusV2. In Table 3, we compare

AdaFocusV2+ with AdaFocusV1+ by deploying them on top of the same AdaFocusV1 base model. AdaFocusV2+ achieves slightly better performance, and removes the requirement of additional reinforcement learning.

**Comparisons with state-of-the-art baselines** on ActivityNet, FCVID and Mini-Kinetics are shown in Table 4. It is clear that AdaFocusV2 ( $128^2$ ) outperforms all the competitive efficient video recognition methods by large margins. For example, it achieves 4.2% higher mAP (79.0% v.s. 74.8%) than VideoIQ [57] on ActivityNet with similar GFLOPs. We further present the variants of the baselines with varying computational costs in Figure 5 for a comprehensive comparison. It can be observed that our method leads to a considerably better efficiency-accuracy trade-off. With the same mAP, the number of the required GFLOPs per video for AdaFocusV2+ is 1.9-3.6 $\times$  less than the strongest baselines.

## 4.2. Deploying on Top of Light-weighted Models

**Setup.** In this subsection, we implement AdaFocusV2 on top of a representative efficient network architecture, CNNs with temporal shift module (TSM) [40]. We adopt the same network architectures and experimental protocols as AdaFocusV1. The MobileNet-V2 and ResNet-50 with

Table 5. **Performance of AdaFocusV2-TSM and representative efficient video recognition models.** MN2, R18/R34/R50 and BN-Inc. denote MobileNet-V2, ResNet-18/34/50 and BN-Inception, respectively. TSM+ refers to the augmented TSM baseline with the same network architecture as our method except for the policy network  $\pi$ . We uniformly sample 8/12 frames for the MobileNetV2/ResNet-50 in our models. The throughput is tested on an NVIDIA GeForce RTX 3090 GPU with a batch size of 128. The best results are **bold-faced**.

Method	Backbones	#Frames	Sth-Sth V1		Sth-Sth V2		Jester		Throughput	
			Top-1 Acc.	GFLOPs	Top-1 Acc.	GFLOPs	Top-1 Acc.	GFLOPs	(NVIDIA 3090, bs=128)	
TSN [63]	R50	8	19.7%	33.2	27.8%	33.2	82.6%	33.2	-	-
AR-Net [47]	MN2 + R18/34/50	8	18.9%	41.4	-	-	87.8%	21.2	-	-
TRN <sub>RGB/Flow</sub> [75]	BN-Inc.	8/8	42.0%	32.0	55.5%	32.0	-	-	-	-
ECO [77]	BN-Inc. + 3DR18	8	39.6%	32.0	-	-	-	-	-	-
TANet [43]	R50	8	47.3%	33.0	60.5%	33.0	-	-	-	-
STM [32]	R50	8	47.5%	33.3	-	-	-	-	-	-
TEA [39]	R50	8	48.9%	35.0	60.9%	35.0	-	-	-	-
TSM [40]	R50	8	46.1%	32.7	59.1%	32.7	96.0%	32.7	162.7 Videos/s	-
AdaFuse-TSM [48]	R50	8	46.8%	31.5	59.8%	31.3	-	-	-	-
TSM+ [40]	MN2 + R50	8+8	47.0%	35.1	59.6%	35.1	96.2%	35.1	123.0 Videos/s	-
AdaFocusV2-TSM (128 <sup>2</sup> )	MN2 + R50	8+12	47.0%	<b>18.5 (↓1.90x)</b>	59.6%	<b>18.5 (↓1.90x)</b>	96.6%	<b>18.5 (↓1.90x)</b>	<b>197.0 Videos/s (↑1.60x)</b>	-
AdaFocusV2-TSM (144 <sup>2</sup> )	MN2 + R50	8+12	48.2%	23.5	60.5%	23.5	-	-	159.6 Videos/s	-
AdaFocusV2-TSM (160 <sup>2</sup> )	MN2 + R50	8+12	48.6%	27.5	60.8%	27.5	-	-	143.6 Videos/s	-
AdaFocusV2-TSM (176 <sup>2</sup> )	MN2 + R50	8+12	<b>49.6% (↑2.6%)</b>	33.7	<b>61.3% (↑1.7%)</b>	33.7	<b>96.9% (↑0.7%)</b>	33.7	123.5 Videos/s	-

Table 6. **Ablation study of the training techniques.** Three representative conditions with different backbones, datasets and varying patch sizes are considered. We report the results of AdaFocusV2-MN2/RN and AdaFocusV2-TSM on ActivityNet and Sth-Sth V1.

Auxiliary supervision	Diversity augmentation	Stop-gradient	Frame sampling policy from [21]	ActivityNet (mAP)		Sth-Sth V1 (Top-1 Acc.)			
				128 <sup>2</sup>	$\Delta$	128 <sup>2</sup>	$\Delta$	144 <sup>2</sup>	$\Delta$
✓				69.4%	-	38.2%	-	40.7%	-
✓				72.5%	+3.1%	42.9%	+4.7%	44.6%	+3.9%
✓	✓			74.8%	+2.3%	46.1%	+3.2%	47.6%	+3.0%
✓	✓	✓		78.7%	+3.9%	<b>47.0%</b>	+0.9%	<b>48.2%</b>	+0.6%
✓	✓	✓	✓	<b>79.0%</b>	+0.3%	-	-	-	-




Figure 9. **Ablation study of AdaFocusV2+.**

TSM are used as  $f_G$  and  $f_L$ . A fully-connected layer is deployed as  $f_C$  to average the frame-wise predictions as outputs. The policy network  $\pi$  generates a single patch location for the whole video after aggregating the information of all frames, which is found important for high generalization performance [64]. Training details can be found in Appendix B. Besides, for fair comparisons, the vanilla TSM is augmented by exploiting the same two backbone networks as ours (named as TSM+). TSM+ differentiates itself from AdaFocusV2 only in that it feeds the whole frames into ResNet-50, while we feed the selected image patches.

**Comparisons of AdaFocusV2 and AdaFocusV1 on Sth-Sth** in terms of training/inference efficiency are shown in Figures 1, 7 and Figure 8. The end-to-end trainable AdaFocusV2 accelerates the training by 2.2-2.4 $\times$ , and improves the accuracy by 1-1.5%/0.6-0.8% on Sth-Sth V1/V2.

**Main results on Sth-Sth and Jester** are presented in Table 5. With the reduced input size, AdaFocusV2 enables TSM to process more frames in the task-relevant region of each video, and effectively improves the inference efficiency. For example, AdaFocusV2-TSM reduces the computational cost of TSM+ by 1.9 $\times$  without sacrificing accuracy. Notably, the practical speedup is significant as well.

### 4.3. Analytical Results

**Effectiveness of the proposed training techniques** are validated in Table 6. One can observe that all of the three techniques significantly improve the performance of

AdaFocusV2 across different experimental settings.

**Ablation study of AdaFocusV2+** is presented in Figure 9. Two variants are considered: (1) random early-exit with the same exit proportion as AdaFocusV2+; (2) early-exit with fixed frame length. Our confidence-based adaptive early-exit mechanism outperforms both of them.

**Ablation study and the visualization results of the learned patch selection policy** are deferred to Appendix C due to the limited space.

## 5. Conclusion

In this paper, we enabled the end-to-end training of adaptive focus video recognition networks (AdaFocus). We first proposed a differentiable interpolation-based operation for selecting patches, allowing the gradient back-propagation throughout the whole model. Then we present three tailored training techniques to address the optimization issues introduced by end-to-end training. Experimental results on six benchmarks demonstrated that our AdaFocusV2 network is considerably more efficient to train than the original AdaFocus model, while achieving state-of-the-art performance.

## Acknowledgements

This work is supported in part by the National Science and Technology Major Project of the Ministry of Science and Technology of China (2018AAA0100701), the National Natural Science Foundation of China (61906106, 62022048), Picsart AI Research and Beijing Academy of Artificial Intelligence.



## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. *arXiv preprint arXiv:2103.15691*, 2021. 1
- [2] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. 6, 12
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. 1, 2
- [4] Jianguo Chen, Kenli Li, Qingying Deng, Keqin Li, and S Yu Philip. Distributed deep learning model for intelligent video surveillance systems with edge computing. *IEEE Transactions on Industrial Informatics*, 2019. 1
- [5] Jin Chen, Xijun Wang, Zichao Guo, Xiangyu Zhang, and Jian Sun. Dynamic region-aware convolution. In *CVPR*, pages 8064–8073, 2021. 2
- [6] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *EMNLP*, pages 1724–1734, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. 3, 7, 12
- [7] Robert T Collins, Alan J Lipton, Takeo Kanade, Hironobu Fujiyoshi, David Duggins, Yanghai Tsin, David Tolliver, Nobuyoshi Enomoto, Osamu Hasegawa, Peter Burt, et al. A system for video surveillance and monitoring. *VSAM final report*, 2000(1-68):1, 2000. 1
- [8] James Davidson, Benjamin Liebald, Junning Liu, Palash Nandy, Taylor Van Vleet, Ullas Gargi, Sujoy Gupta, Yu He, Mike Lambert, Blake Livingston, et al. The youtube video recommendation system. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 293–296, 2010. 1
- [9] Yashar Deldjoo, Mehdi Elahi, Paolo Cremonesi, Franca Gatzotto, Pietro Piazzolla, and Massimo Quadrona. Content-based video recommendation system based on stylistic visual features. *Journal on Data Semantics*, 5(2):99–113, 2016. 1
- [10] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pages 2625–2634, 2015. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3, 5
- [12] Linxi Fan, Shyamal Buch, Guanzhi Wang, Ryan Cao, Yuke Zhu, Juan Carlos Niebles, and Li Fei-Fei. Rubiksnet: Learnable 3d-shift for efficient video action recognition. In *ECCV*, pages 505–521. Springer, 2020. 2
- [13] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *CVPR*, pages 203–213, 2020. 1, 2
- [14] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019. 1, 2
- [15] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Spatiotemporal multiplier networks for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4768–4777, 2017. 2
- [16] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, pages 1933–1941, 2016. 1, 2
- [17] Michael Figurnov, Maxwell D Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *CVPR*, pages 1039–1048, 2017. 2
- [18] Junyu Gao, Tianzhu Zhang, and Changsheng Xu. A unified personalized video recommendation via dynamic recurrent neural networks. In *ACM MM*, pages 127–135, 2017. 1
- [19] Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, and Lorenzo Torresani. Listen to look: Action recognition by previewing audio. In *CVPR*, pages 10457–10467, 2020. 1, 2, 6, 12
- [20] Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *CVPR*, pages 12046–12055, 2019. 1
- [21] Amir Ghodrati, Babak Ehteshami Bejnordi, and Amirhossein Habibian. Frameexit: Conditional early exiting for efficient video recognition. In *CVPR*, pages 15608–15618, 2021. 1, 2, 3, 5, 6, 7, 8, 12
- [22] Xinyu Gong, Heng Wang, Mike Zheng Shou, Matt Feiszli, Zhangyang Wang, and Zhicheng Yan. Searching for two-stream models in multivariate space for video recognition. In *ICCV*, pages 8033–8042, 2021. 2
- [23] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, pages 5842–5850, 2017. 6, 12
- [24] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021. 2
- [25] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *CVPR*, pages 6546–6555, 2018. 1, 2
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5, 6
- [27] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 3
- [28] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. In *ICLR*, 2018. 5, 6
- [29] Gao Huang, Zhuang Liu, Geoff Pleiss, Laurens Van Der Maaten, and Kilian Weinberger. Convolutional networks with dense connectivity. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 5

- [30] Nazli Ikizler and David Forsyth. Searching video for complex activities with finite state models. In *CVPR*, pages 1–8. IEEE, 2007. 1
- [31] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *NeurIPS*, 28:2017–2025, 2015. 2
- [32] Boyuan Jiang, MengMeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. Stm: Spatiotemporal and motion encoding for action recognition. In *ICCV*, pages 2000–2009, 2019. 8
- [33] Y.-G. Jiang, Z. Wu, J. Wang, X. Xue, and S.-F. Chang. Exploiting feature and class relationships in video categorization with regularized deep neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(2):352–364, 2018. 6, 12
- [34] 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. 6, 12
- [35] Hanul Kim, Mihir Jain, Jun-Tae Lee, Sungrack Yun, and Fatih Porikli. Efficient action recognition via dynamic knowledge propagation. In *ICCV*, pages 13719–13728, 2021. 1, 2, 6, 12
- [36] Bruno Korbar, Du Tran, and Lorenzo Torresani. Scsampler: Sampling salient clips from video for efficient action recognition. In *ICCV*, pages 6232–6242, 2019. 1, 2, 6, 12
- [37] Changlin Li, Guangrun Wang, Bing Wang, Xiaodan Liang, Zhihui Li, and Xiaojun Chang. Ds-net++: Dynamic weight slicing for efficient inference in cnns and transformers. *arXiv preprint arXiv:2109.10060*, 2021. 5
- [38] Dong Li, Zhaofan Qiu, Qi Dai, Ting Yao, and Tao Mei. Recurrent tubelet proposal and recognition networks for action detection. In *ECCV*, pages 303–318, 2018. 2
- [39] Yan Li, Bin Ji, Xintian Shi, Jianguo Zhang, Bin Kang, and Limin Wang. Tea: Temporal excitation and aggregation for action recognition. In *CVPR*, pages 909–918, 2020. 8
- [40] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. 2, 3, 6, 7, 8, 12, 13
- [41] Zhaoyang Liu, Donghao Luo, Yabiao Wang, Limin Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Tong Lu. Teinet: Towards an efficient architecture for video recognition. In *AAAI*, volume 34, pages 11669–11676, 2020. 2
- [42] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 1
- [43] Zhaoyang Liu, Limin Wang, Wayne Wu, Chen Qian, and Tong Lu. Tam: Temporal adaptive module for video recognition. In *ICCV*, pages 13708–13718, 2021. 2, 8
- [44] Chenxu Luo and Alan L Yuille. Grouped spatial-temporal aggregation for efficient action recognition. In *ICCV*, pages 5512–5521, 2019. 2
- [45] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. The jester dataset: A large-scale video dataset of human gestures. In *ICCVW*, 2019. 6, 12
- [46] Lili Meng, Bo Zhao, Bo Chang, Gao Huang, Wei Sun, Frederick Tung, and Leonid Sigal. Interpretable spatio-temporal attention for video action recognition. In *ICCV Workshops*, pages 0–0, 2019. 2
- [47] Yue Meng, Chung-Ching Lin, Rameswar Panda, Prasanna Sattigeri, Leonid Karlinsky, Aude Oliva, Kate Saenko, and Rogerio Feris. Ar-net: Adaptive frame resolution for efficient action recognition. In *ECCV*, pages 86–104. Springer, 2020. 1, 2, 3, 6, 8, 12
- [48] Yue Meng, Rameswar Panda, Chung-Ching Lin, Prasanna Sattigeri, Leonid Karlinsky, Kate Saenko, Aude Oliva, and Rogerio Feris. Adafuse: Adaptive temporal fusion network for efficient action recognition. In *ICLR*, 2021. 2, 3, 6, 8, 12
- [49] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013. 3
- [50] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32:8026–8037, 2019. 13
- [51] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *CVPR*, pages 12056–12065, 2019. 1
- [52] Mengye Ren, Andrei Pokrovsky, Bin Yang, and Raquel Urtasun. Sbnet: Sparse blocks network for fast inference. In *CVPR*, pages 8711–8720, 2018. 2
- [53] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520, 2018. 6
- [54] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 3
- [55] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017. 3, 5
- [56] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift networks for video action recognition. In *CVPR*, pages 1102–1111, 2020. 2
- [57] Ximeng Sun, Rameswar Panda, Chun-Fu Richard Chen, Aude Oliva, Rogerio Feris, and Kate Saenko. Dynamic network quantization for efficient video inference. In *ICCV*, pages 7375–7385, 2021. 1, 2, 6, 7, 12
- [58] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015. 1, 2
- [59] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *ICCV*, pages 5552–5561, 2019. 1, 2
- [60] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018. 2

- [61] Thomas Verelst and Tinne Tuytelaars. Dynamic convolutions: Exploiting spatial sparsity for faster inference. In *CVPR*, pages 2320–2329, 2020. 2
- [62] Guangrun Wang, Keze Wang, and Liang Lin. Adaptively connected neural networks. In *CVPR*, pages 1781–1790, 2019. 2
- [63] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36. Springer, 2016. 2, 8
- [64] Yulin Wang, Zhaoxi Chen, Haojun Jiang, Shiji Song, Yizeng Han, and Gao Huang. Adaptive focus for efficient video recognition. In *ICCV*, October 2021. 2, 3, 4, 6, 8, 12, 13
- [65] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. In *NeurIPS*, 2021. 5, 6
- [66] Yulin Wang, Kangchen Lv, Rui Huang, Shiji Song, Le Yang, and Gao Huang. Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. In *NeurIPS*, 2020. 2, 5, 13
- [67] Wenhao Wu, Dongliang He, Xiao Tan, Shifeng Chen, and Shilei Wen. Multi-agent reinforcement learning based frame sampling for effective untrimmed video recognition. In *ICCV*, pages 6222–6231, 2019a. 1, 2, 6, 12
- [68] Zuxuan Wu, Hengduo Li, Caiming Xiong, Yu-Gang Jiang, and Larry Steven Davis. A dynamic frame selection framework for fast video recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020b. 1, 2, 6, 12
- [69] Zuxuan Wu, Caiming Xiong, Yu-Gang Jiang, and Larry S Davis. Liteeval: A coarse-to-fine framework for resource efficient video recognition. In *NeurIPS*, 2019b. 6, 12
- [70] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *CVPR*, 2022. 2
- [71] Zhenda Xie, Zheng Zhang, Xizhou Zhu, Gao Huang, and Stephen Lin. Spatially adaptive inference with stochastic feature sampling and interpolation. In *ECCV*, pages 531–548. Springer, 2020. 2
- [72] Le Yang, Yizeng Han, Xi Chen, Shiji Song, Jifeng Dai, and Gao Huang. Resolution adaptive networks for efficient inference. In *CVPR*, pages 2369–2378, 2020. 2, 5
- [73] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, pages 2678–2687, 2016. 1, 2
- [74] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, pages 4694–4702, 2015. 2
- [75] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, pages 803–818, 2018. 8
- [76] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 3, 5
- [77] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. Eco: Efficient convolutional network for online video understanding. In *ECCV*, pages 695–712, 2018. 2, 8