Look More but Care Less in Video Recognition

Anonymous Author(s)
Affiliation
Address
email

Abstract

Existing action recognition methods typically sample a few frames to represent each video to avoid the enormous computation, which often limits the recognition performance. To tackle this problem, we propose Ample and Focal Network (AFNet), which is composed of two branches to *utilize more frames but with less computation*. Specifically, the Ample Branch takes all input frames to obtain abundant information with condensed computation and provides the guidance for Focal Branch by the proposed Navigation Module; the Focal Branch squeezes the temporal size to only focus on the salient frames at each convolution block; in the end, the results of two branches are adaptively fused to prevent the loss of information. With this design, we can introduce more frames to the network but cost less computation. Besides, we demonstrate AFNet can *utilize fewer frames while achieving higher accuracy* as the dynamic selection in intermediate features enforces implicit temporal modeling. Further, we show that our method can be extended to reduce spatial redundancy with even less cost. Extensive experiments on five datasets demonstrate the effectiveness and efficiency of our method.

1 Introduction

Online videos have grown wildly in recent years and video analysis is necessary for many applications such as recommendation [4] and surveillance [3]. These applications require not only accurate but also efficient video understanding algorithms. With the introduction of deep learning networks [2] in video action recognition, there has been rapid progress in the performance of the methods in this area. Though successful, these deep learning based methods are often at the cost of huge computation, making them hard to be deployed in the real world.

In video recognition, we need to sample a certain number of frames to represent each video which makes the computational cost scale proportionally to the number of sampled frames. In most cases, a small proportion of all the frames is sampled for each input, which only contains limited information of the original video. A straightforward solution is to sample more frames to the network but the computation expands proportionally to the number of sampled frames.

There are some works proposed recently to dynamically sample salient frames [27, 15]. The selection step of these methods is made before the frames being sent to the classification network, which means the information of those unimportant frames is totally lost and it consumes a considerable time for the sampling procedure. Some other methods proposed to address the spatial redundancy in action recognition by adaptively resizing the resolution based on the importance of each frame [21], or cropping the most salient patch for every frame [26]. However, these methods still completely abandon the information that the network recognizes as unimportant and introduce a policy network to make decisions for each sample which leads to extra computation and complicates training strategies.

In our work, we go from another perspective compared with previous works. We propose a method which makes frame selection within the classification network. Shown in Figure 1, we design an

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

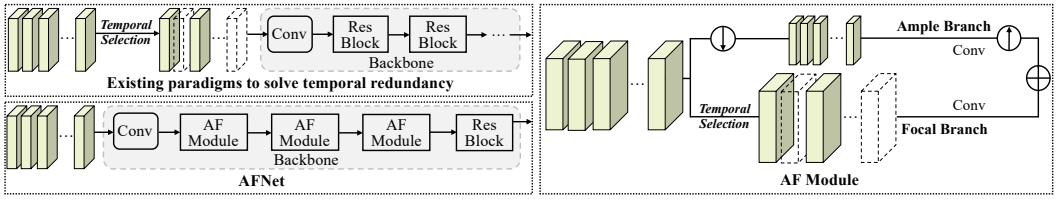

Figure 1: Comparisons between existing methods and our proposed Ample and Focal Network (AFNet). Most existing works reduce the redundancy in data at the beginning of the deep networks which leads to the loss of information. We propose a two-branch design which processes frames with different computational resources within the network and preserves all input information as well.

architecture called Ample and Focal Network (AFNet) which is composed of two branches: the ample branch takes a glimpse of all the input features by lightweight computation as we downsample the features and reduce the channel size; the focal branch receives the guidance from the proposed navigation module to squeeze the temporal size by only computing on the selected frames to save cost; in the end, we adaptively fuse the features of these two branches to prevent the loss of information.

With this design, the two branches are both very lightweight and we enable AFNet to look broadly by sampling more frames and stay focused on the important information for less computation. Considering these two branches in a uniform manner, on the one hand, we can avoid the loss of information compared to other dynamic methods as the ample branch preserves the information of all the input; on the other hand, we can restrain the noise from the unimportant frames by deactivating them in each convolutional block. Further, we have demonstrated that the dynamic selection strategy in intermediate features is beneficial for temporal modeling as it implicitly implements frame-wise attention which can enable our network to use fewer frames while achieving higher accuracy. In addition, instead of introducing a policy network to select frames, we design a lightweight navigation module which can be easily plugged into the network so that our method can be trained in an end-to-end fashion. It is worth noting that AFNet is compatible with spatial adaptive works which can help to further reduce the computations of our method.

The main contributions are summarized as follows:

- We propose an adaptive two-branch framework which enables 2D-CNNs to *process more frames with less computational cost*. With this design, we not only prevent the loss of information but strengthen the representation of essential frames.
- We propose a lightweight navigation module to *dynamically select salient frames* at each convolution block which can be easily trained in an *end-to-end* way.
- The selection strategy at intermediate features not only empowers the model with strong flexibility as different frames will be selected at different layers, but also enforces *implicit temporal modeling* which enables AFNet to *obtain higher accuracy with fewer frames*.
- We have conducted extensive experiments on five video recognition benchmarks. The results show the *superiority of AFNet compared to state-of-the-art methods*.

2 Related Work

2.1 Video Recognition

The development of deep learning in recent years serves as a huge boost to the research of video recognition. A straightforward method for this task is using 2D-CNNs to extract the features of sampled frames and use specific aggregation methods to model the temporal relationships across frames. For instance, [25] proposes to average the temporal information between frames. While [18] shifts channels along the temporal dimension to exchange the information among neighboring frames. Another approach is to use 3D-CNNs to model the temporal and spatial information jointly, such as C3D [24], I3D [2] and ResNet3D [10]. Though being shown effective, methods based on 3D-CNNs are computationally expensive, which brings great difficulty in real-world deployment.

While the two-branch design has been explored by SlowFast [6], but the motivation and detailed structure are different in the following ways: (1) network category: SlowFast is a static 3D model, but AFNet is a dynamic 2D network; (2) motivation: SlowFast is designed to capture semantic information and changing motion with branches at different temporal speed for better performance, while AFNet is aimed to dynamically skip frames to save computation and the design of two-branch structure is to prevent the information loss; (3) specific design: AFNet is designed to downsample

features for efficiency at ample branch while SlowFast processes features in the original resolution; (4) temporal modeling: SlowFast applies 3D convolutions for temporal modeling, AFNet is a 2D model which is enforced with implicit temporal modeling by the designed navigation module.

2.2 Redundancy in Data

The efficiency of 2D-CNNs has been broadly studied in recent years. While some of the works aim at designing efficient network structure [12], there is another line of research focusing on reducing the intrinsic redundancy in image-based data [29, 9]. In video recognition, people usually sample limited number of frames to represent each video to prevent numerous computational costs. Even though, the computation for video recognition is still a heavy burden for researchers and a common strategy to address this problem is reducing the temporal redundancy in videos as not all frames are essential to the final prediction. [30] proposes to use reinforcement learning to skip frames for action detection. There are other works [27, 15] dynamically sampling salient frames to save computational cost. As spatial redundancy widely exists in image-based data, [21] adaptively processes frames with different resolutions. [26] provides the solution as cropping the most salient patch for each frame. However, the unselected regions or frames of these works are completely abandoned. Hence, there will be some information lost in their designed procedures. Moreover, most of these works adopt a policy network to make dynamic decisions, which introduce additional computation somehow and cannot be trained in an end-to-end fashion. In contrast, our method adopts a two-branch design, allocating different computational resources based on the importance of each frame and preventing the loss of information. Besides, we design a lightweight navigation module to guide the network where to look, which can be incorporated to the backbone network and trained in an end-to-end way. Moreover, we demonstrate that the dynamic frame selection at intermediate features will not only empower the model with strong flexibility as different frames will be selected at different layers, but result in learned frame-wise weights which enforce implicit temporal modeling.

3 Methodology

Intuitively, considering more frames enhances the temporal modeling but results in higher computational cost. To efficiently achieve the competitive performance, we propose AFNet to involve more frames but wisely extract information from them to keep the low computational cost. Specifically, we design a two-branch structure to treat frames differently based on their importance and process the features in an adaptive manner which can provide our method with strong flexibility. Besides, we demonstrate that the dynamic selection of frames in the intermediate features results in learned frame-wise weights which can be regarded as implicit temporal modeling.

3.1 Architecture Design

As is shown in Figure 2, we design our Ample and Focal (AF) module as a two-branch structure: the ample branch (top) processes abundant features of all the frames in a lower resolution and a squeezed channel size; the focal branch (bottom) receives the guidance from ample branch generated by the navigation module and makes computation only on the selected frames. Such design can be conveniently plugged into existing CNN architectures to build our network.

Ample Branch. The ample branch is designed to involve all frames with cheap computation, which serves as 1) a guidance to select salient frames to help focal branch to concentrate on important information; 2) a complementary stream with focal branch to prevent the information loss via a specially designed fusion strategy.

Formally, we denote video sample i as v^i, containing T frames as $v^i = \left\{f_1^i, f_2^i, ..., f_T^i\right\}$. For convenience, we omit the superscript i in the following sections if no confusion arises. We denote the input of ample branch as $v_x \in \mathbb{R}^{T\times C\times H\times W}$, where C represents the channel size and $H \times W$ is the resolution. The features generated by the ample branch can be written as:

$$v_{y^a} = F^a\left(v_x\right), \tag{1}$$

where $v_{y^a} \in \mathbb{R}^{T\times(C_o/2)\times(H_o/2)\times(W_o/2)}$ represents the output of ample branch with channel, height, width at focal branch denoted as C_o, H_o, W_o, respectively, and F^a stands for a series of convolution blocks. The stride of the first convolution block is set to 2 to downsample the resolution of this branch and we upsample the feature at the end of this branch by nearest interpolation.

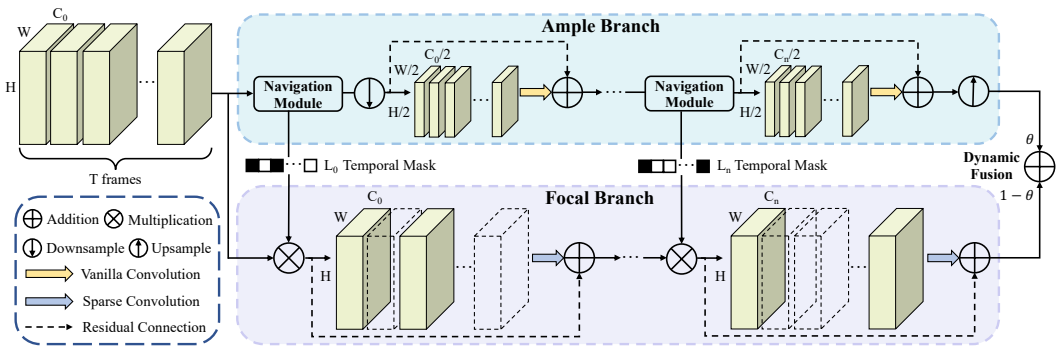

Figure 2: Architecture of AF module. The module is composed of two branches, the ample branch would process all the input features in a lower resolution and reduced channel size; while the focal branch would only compute the features of salient frames (colored features) guided by our proposed navigation module. The results of two branches are adaptively fused at the end of AF module so that we can prevent the loss of information.

Navigation Module. The proposed navigation module is designed to guide the focal branch where to look by adaptively selecting the most salient frames for video v^i.

Specifically, the navigation module generates a binary temporal mask L_n based on the output from the n-th convolution block in ample branch $v_{y_n^a}$. At first, average pooling is applied to $v_{y_n^a}$ to resize the spatial dimension to 1×1, then convolution is performed to transform channel size to 2:

$$\tilde{v}_{y_n^a} = ReLU\left(BN\left(W_1 * Pool\left(v_{y_n^a}\right)\right)\right), \tag{2}$$

where $*$ stands for convolution and W_1 denotes the weights of the 1×1 convolution. After that, we reshape the dimension of feature $\tilde{v}_{y_n^a}$ from $T \times 2 \times 1 \times 1$ to $1 \times (2 \times T) \times 1 \times 1$ so that we can model the temporal relations for each video from channel dimension by:

$$p_n^t = W_2 * \tilde{v}_{y_n^a}, \tag{3}$$

where W_2 represents the weights of the second 1×1 convolution and it will generate a logit $p_n^t \in \mathbb{R}^2$ for each frame t in the output frame map which denotes whether to select it.

However, directly sampling from such discrete distribution is non-differentiable. In this work, we apply Gumbel-Softmax [13] to resolve this non-differentiability. Specifically, we generate a normalized categorical distribution by using Softmax:

$$\pi = \left\{ l_j \mid l_j = \frac{\exp\left(p_n^{t_j}\right)}{\exp\left(p_n^{t_0}\right) + \exp\left(p_n^{t_1}\right)} \right\}, \tag{4}$$

and we draw discrete samples from the distribution π as:

$$L = \arg\max_j \left(\log l_j + G_j\right), \tag{5}$$

where $G_j = -\log(-\log U_j)$ is a standard Gumbel distribution with U_j sampled from a uniform distribution Unif(0,1). As $\arg\max$ cannot be differentiated, we relax the discrete sample L in backpropagation via Softmax:

$$\hat{l}_j = \frac{\exp\left(\left(\log l_j + G_j\right)/\tau\right)}{\sum_{k=1}^{2} \exp\left(\left(\log l_k + G_k\right)/\tau\right)}, \tag{6}$$

the distribution $\hat{l}$ will become a one-hot vector when the temperature factor $\tau \to 0$ and we let τ decrease from 1 to 0.01 during training.

Focal Branch. The focal branch is guided by the navigation module to only consider the selected frames, which diminishes the computational cost and potential noise from redundant frames.

The features at the n-th convolution block in this branch can be denoted as $v_{y_n^f} \in \mathbb{R}^{T \times C_o \times H_o \times W_o}$. Based on the temporal mask L_n from the navigation module, we select frames which have corresponding non-zero values in the binary mask for each video and apply convolutional operations only on these extracted frames $v'_{y_n^f} \in \mathbb{R}^{T_l \times C_o \times H_o \times W_o}$:

$$v'_{y_n^f} = F_n^f\left(v'_{y_{n-1}^f}\right), \tag{7}$$

where F_n^f is the n-th convolution blocks at this branch and we set the group number of convolutions to 2 in order to further reduce the computations. After the convolution operation at n-th block, we generate a zero-tensor which shares the same shape with $v_{y_n^f}$ and fill the value by adding $v'_{y_n^f}$ and $v_{y_{n-1}^f}$ with the residual design following [11].

At the end of these two branches, we generate a weighting factor θ by pooling and linear layers to fuse the features from two branches:

$$v_y = \theta \odot v_{y^a} + (1-\theta) \odot v_{y^f}, \tag{8}$$

where $\odot$ denotes the channel-wise multiplication.

3.2 Implicit Temporal Modeling

While our work is mainly designed to reduce the computational cost in action recognition like [26, 22], we demonstrate that AFNet enforces implicit temporal modeling by the dynamic selection of frames in the intermediate features. Considering a TSN[25] network which adapts vanilla ResNet[11] structure, the feature at the n-th convolutional block in each stage can be written as $v_n \in \mathbb{R}^{T\times C\times H\times W}$. Thus, the feature at $n+1$-th block can be represented as:

$$\begin{aligned} v_{n+1} &= v_n + F_{n+1}(v_n) \\ &= (1+\Delta v_{n+1})\, v_n, \end{aligned} \tag{9}$$

$$\Delta v_{n+1} = \frac{F_{n+1}(v_n)}{v_n}, \tag{10}$$

where F_{n+1} is the $n+1$-th convolutional block and we define Δv_{n+1} as the coefficient learned from this block. By that we can write the output of this stage v_N as:

$$v_N = \left[\prod_{n=2}^{N} (1+\Delta v_n)\right] * v_1. \tag{11}$$

Similarly, we define the features in ample and focal branch as:

$$v_{y_N^a} = \left[\prod_{n=2}^{N} \left(1+\Delta v_{y_n^a}\right)\right] * v_{y_1}, \tag{12}$$

$$v_{y_N^f} = \left[\prod_{n=2}^{N} \left(1+L_n * \Delta v_{y_n^f}\right)\right] * v_{y_1}, \tag{13}$$

where L_n is the binary temporal mask generated by Equation 5 and v_{y_1} denotes the input of this stage. Based on Equation 8, we can get the output of this stage as:

$$\begin{aligned} v_{y_N} &= \theta \odot v_{y_N^a} + (1-\theta) \odot v_{y_N^f} \\ &= \left\{\theta \odot \left[\prod_{n=2}^{N} \left(1+\Delta v_{y_n^a}\right)\right] + (1-\theta) \odot \left[\prod_{n=2}^{N} \left(1+L_n * \Delta v_{y_n^f}\right)\right]\right\} * v_{y_1}. \end{aligned} \tag{14}$$

As L_n is a temporal-wise binary mask, it will decide whether the coefficient $\Delta v_{y_n^f}$ will be calculated in each frame at every convolutional block. Considering the whole stage is made up of multiple convolutional blocks, the series multiplication of focal branch's output with the binary mask L_n will approximate soft weights. This results in learned frame-wise weights in each video which we regard as implicit temporal modeling. Although we do not explicitly build any temporal modeling module, the generation of L_n in Equation 3 has already taken the temporal information into account so that the learned temporal weights equal performing implicit temporal modeling at each stage.

3.3 Spatial Redundancy Reduction

In this part, we show that our approach is compatible with methods that aim to solve the problem of spatial redundancy. We extend the navigation module by applying similar procedures with the

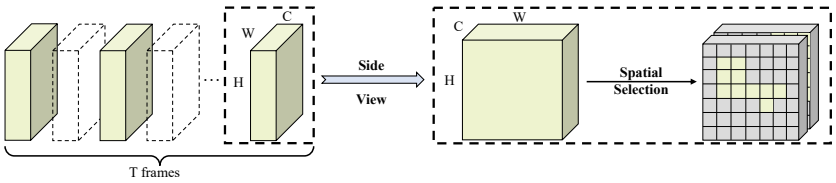

Figure 3: Illustration of extending AFNet to reduce spatial redundancy to further improve the efficiency. Only the colored area will be calculated at the inference stage.

temporal mask generation and the work [9] to generate a spatial logit for the n-th convolution block which is shown in Figure 3:

$$q_n^t = W_4 * \left(Pool\left(ReLU\left(BN\left(W_3 * v_{y_n^a}\right)\right)\right)\right), \tag{15}$$

where W_3 denotes the weights of the 3×3 convolution and W_4 stands for the weights of convolution with kernel size 1×1. After that, we still use Gumbel-Softmax to sample from such discrete distribution to generate spatial mask M_n and guide the focal branch to merely focus on the salient regions of the selected frames to further reduce the cost.

3.4 Loss functions

Inspired by [25], we average the predictions of each frame to represent the final prediction of the corresponding video and our optimization objective is minimizing:

$$\mathcal{L} = \sum_{(v,y)} \left[-y \log\left(P\left(v\right)\right) + \lambda \cdot \sum_{n=1}^{N} \left(r - RT\right)^2 \right]. \tag{16}$$

The first term is the cross-entropy between predictions $P\left(v\right)$ for input video v and the corresponding one-hot label y. We denote r in the second term as the ratio of selected frames in each batch and RT as the target ratio which is set before the training(RS is the target ratio when applying spatial selection module). We let r approximate RT by adding the second loss term and control the trade-off between accuracy and efficiency by introducing a factor λ which balances these two terms.

4 Empirical Validation

In this section, we conduct extensive experiments to validate our method. We first compare our method with plain 2D CNNs to demonstrate that our AF module implements implicit temporal-wise attention which is beneficial for temporal modeling. Then, we validate AFNet's efficiency by introducing more frames but costing less computation compared with other methods. Further, we show AFNet's strong performance compared with other efficient action recognition frameworks. Finally, we provide comprehensive ablation results and qualitative analysis to demonstrate the effectiveness of our proposed navigation module and two-branch design.

Datasets. Our method is evaluated on five widely-used video datasets: (1) ActivityNet-v1.3 [1] is labeled with 200 action categories. It contains 10,024 training videos and 4,926 validation videos in total and the average duration is 117 seconds; (2) Jester is a hand gesture recognition dataset introduced by [20]. The dataset consists of 27 classes, with 119k videos for training and 15k videos for validation; (3) Something-Something V1&V2 [8] are two human action datasets with strong temporal information, including 98k and 194k videos for training and validation respectively; (4) Mini-Kinetics [21, 22] is a subset of Kinetics [14] which selects 200 classes from Kinetics, with 121k videos for training and 10k videos for validation.

Data pre-processing. We uniformly sample 8 frames for each video on Jester, Mini-Kinetics, and 12 frames for ActivityNet and Something-Something to compare with existing works. During training, the training data is randomly cropped to 224×224 following [32], and we perform random flipping except for Something-Something. At inference stage, we use one clip per video for efficiency and all frames are center-cropped to 224×224.

Implementation details. Our method is bulit on ResNet-50 [11] and we replace the first three stages of the network with our proposed AF module. We first train our two-branch network from scratch on ImageNet for fair comparisons with other methods, then add the proposed navigation module and train it along with the backbone network on video datasets. In our implementations, RT denotes the ratio of selected frames while RS represents the ratio of select regions. Due to limited space, more implementation details can be found in supplementary material.

4.1 Comparisons with Existing Methods

Baselines. In this part, we first compare AFNet with TSN [25] on Jester and Something-Something V1 to demonstrate that our AFNet can achieve *Less is More* by using less frames while obtaining higher accuracy. Further, we build our framework on TSM [18] and compare with other works to validate that AFNet can achieve *More is Less* by seeing more frames with less computations. Then we empirically compare AFNet with competing efficient video recognition methods, including AdaFrame [27], LiteEval [28], ListenToLook [7], SCSampler [15], AR-Net [21], VideoIQ [23], AdaFoucs [26] on the dataset of ActivityNet and Mini-Kinetics.

Less is more. At first, we implement AFNet on Something-Something V1 and Jester datasets. We compare it with the baseline method TSN as both methods do not explicitly build temporal modeling module and are built on ResNet-50. In Table 1, our method shows similar performance with TSN when selecting all the frames. Nonetheless, when we select fewer frames in AFNet, it exhibits much higher accuracy compared to TSN. The results may seem counterintuitive as seeing more frames is usually beneficial for video recognition. The explanation is that the two-branch design of AFNet can preserve the information of all input frames and the intermediate selection of salient frames implements implicit temporal modeling compared to TSN which simply averages the logits of all the frames at the end of the network. The results also indicate that the navigation module not only reduces the computational cost but strengthens the representation of salient frames as well. When we decrease the ratio of selected frames from 0.5 to 0.25, there is no accuracy drop on Jester and the accuracy even increases by 0.9% on Something-Something V1, which further demonstrates the effectiveness of our navigation module.

Table 1: Comparisons with baseline method on Something-Something V1 and Jester datasets.

Method	Frame	Sth-Sth V1	Jester
		Top-1 Acc.	Top-1 Acc.
TSN [25](our imp)	8	18.6%	83.5%
AFNet (RT=1.00)	8	**19.2**%	**83.6**%
AFNet (RT=0.50)	8	**26.8**%	**89.2**%
AFNet (RT=0.25)	8	**27.7**%	**89.2**%

Table 2: Performance comparison on Something-Something (Sth-Sth) and Jester datasets. GFLOPs refers to the average computational cost to process a single video.

Method	Dynamic	Frames	Sth-Sth V1		Sth-Sth V2		Jester	
			Top-1 Acc.	GFLOPs	Top-1 Acc.	GFLOPs	Top-1 Acc.	GFLOPs
TRN$_{RGB/Flow}$ [31]	✗	8/8	42.0%	32.0	55.5%	32.0	-	-
ECO [32]	✗	8	39.6%	32.0	-	-	-	
TSM [18]	✗	8	45.6%	32.7	59.1%	32.7	96.0%	32.7
bLVNet-TAM [5]	✗	8×2	47.8%	35.1	60.2%	35.1	-	
TANet [19]	✗	8	47.3%	33.0	60.5%	33.0	-	-
SmallBig [16]	✗	8	47.0%	52.0	59.7%	52.0	-	-
TEA [17]	✗	8	48.9%	35.0	60.9%	35.0	-	-
SlowFast [6]	✗	8×8	-	-	61.7%	66.6×3	-	-
AdaFuse-TSM [22]	✓	8	46.8%	31.5	59.8%	31.3	-	-
AdaFocus-TSM [26]	✓	8+12	48.1%	33.7	60.7%	33.7	-	-
TSM [18](our imp)	✗	12	47.9%	49.1	60.9%	49.1	96.4%	49.1
AFNet-TSM (RT=0.4)	✓	12	**49.0**%	**27.9** (↓1.76x)	**61.3**%	**27.8** (↓1.77x)	**97.0**%	**27.6** (↓1.78x)
AFNet-TSM (RT=0.8)	✓	12	**49.9**%	31.7	**62.5**%	31.7	**97.1**%	31.5

More is less. In this segment, we incorporate our method with temporal shift module (TSM [18]) to validate that AFNet can further reduce the redundancy of such competing methods and achieve *More is Less* by seeing more frames with less computation. We implement our method on the datasets of Something-Something V1&V2 and Jester which contain strong temporal information and relevant results are shown in Table 2. Compared to TSM which samples 8 frames, our method shows significant advantages in performance as we introduce more frames. Yet, our computation cost is much less than TSM as we allocate frames with different computation resources by this two-branch design. Compared to other competitive dynamic methods like AdaFuse and AdaFocus, our method shows an obviously better performance both in accuracy and computations. When costing similar computation, AFNet outperforms AdaFuse and AdaFocus by 3.1% and 1.8% respectively on Something-Something v1. Moreover, we compare our method with TSM when both sampling 12 frames and AFNet still outperforms TSM which achieves *Less is More* on the TSM structure as the navigation module can restrain the noise of unimportant frames and enforces frame-wise attention which can be regarded as implicit temporal modeling.

Table 3: Comparisons with state-of-the-art efficient video recognition methods on Mini-Kinetics. AFNet achieves the best trade-off compared to existing works. GFLOPs refers to the average computation to process a single video.

Method	Mini-Kinetics	
	mAP	GFLOPs
LiteEval [28]	61.0%	99.0
SCSampler [15]	70.8%	42.0
AR-Net [21]	71.7%	32.0
AdaFuse [22]	72.3%	23.0
AdaFocus [26]	72.2%	26.6
VideoIQ [23]	72.3%	20.4
AFNet (RT=0.4)	**72.8%**	**19.4**
AFNet (RT=0.8)	**73.5%**	22.0

Table 4: Comparisons with state-of-the-art efficient video recognition methods on ActivityNet. AFNet achieves the best trade-off compared to existing works. GFLOPs refers to the average computational cost to process a single video.

Method	ActivityNet	
	mAP	GFLOPs
AdaFrame [27]	71.5%	79.0
LiteEval [28]	72.7%	95.1
ListenToLook [7]	72.3%	81.4
SCSampler [15]	72.9%	42.0
AR-Net [21]	73.8%	33.5
VideoIQ [23]	74.8%	28.1
AdaFocus [26]	75.0%	26.6
AFNet (RS=0.4,RT=0.8)	**75.6%**	**24.6**

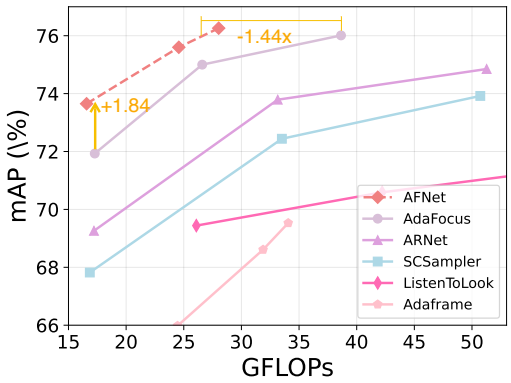

Figure 4: Computational cost vs mean Average Precision on ActivityNet. Our method is implemented with different ratio of selected frames $RT \in \{0.6, 0.8, 1.0\}$ and a fixed ratio of selected regions $RS \in \{0.4\}$.

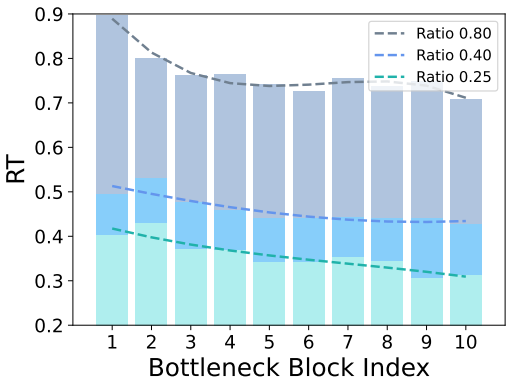

Figure 5: Distribution of $RT \in \{0.25, 0.4, 0.8\}$ for each convolution block in ResNet on Something-Something V1. Dash lines denote using 3rd-order polynomials to estimate the trend of distribution.

Comparisons with state-of-the-art methods. Then, we implement our method on the dataset of Mini-Kinetics, ActivityNet-v1.3, and compare AFNet with other efficient video recognition approaches. At first, we evaluate our method on Mini-Kinetics and our method shows the best performance both in accuracy and computations compared with other efficient methods. To demonstrate that AFNet can further reduce spatial redundancy, we extend AFNet by adding the spatial selection module on ActivityNet. We move the navigation module to the first layer of the network to avoid huge variance in features when incorporating spatial selection module and note that we only apply this procedure in this part. As shown in Table 4, our method achieves the best performance while costing the least computation compared to other works. Moreover, we change the ratio of selected frames and plot the mean Average Precision and computational cost of various methods in Figure 4. One can observe that AFNet exhibits a better trade-off between accuracy and efficiency compared to other works.

4.2 Visualizations and Ablation Study

In this part, we implement our method on ActivityNet with 12 sampled frames to analyze the results and conduct extensive ablation study to prove the effectiveness of our design.

Visualization results. We show the distribution of RT among different convolution blocks under different selection ratios in Figure 5 and use 3rd-order polynomials to estimate the trend of distribution (shown in dash lines). One can see a decreased trend in RT for all the curves with the increased index in convolution blocks and this can be explained that earlier layers mostly capture low-level information which has relatively large divergence among different frames. While high-level semantics between different frames are more similar, therefore AFNet tends to skip more at later convolution blocks. In Figure 6, we visualize the frames selected by the 3rd-block of our AFNet with RT=0.5 on

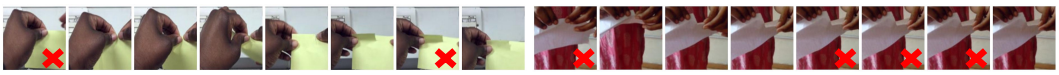

(a) Tearing something just a little bit

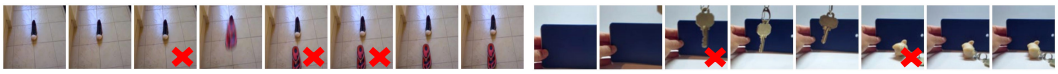

(b) Dropping something in front of something

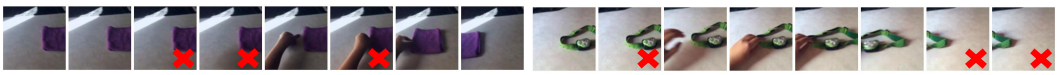

(c) Pulling something from right to left

Figure 6: Visualization results on Something-Something V1 dataset. Frames annotated with red crosses denote the ones that AFNet does not select.

Table 5: Ablation study on navigation module and two-branch architecture. Different sampling policies are compared with our dynamic temporal selection module in various selection ratios.

Ablation			mAP		
Structure	Temporal Sampling Policy	Spatial Sampling Policy	0.25	0.50	0.75
Single Branch	Navigation Module	-	46.8%	60.4%	71.3%
Two Branch	Random Sampling	-	62.8%	71.3%	74.2%
	Uniform Sampling		59.0%	71.7%	74.7%
	Normal Sampling		61.5%	71.4%	74.6%
Two Branch	Navigation Module	-	**72.2**%	**74.3**%	**75.8**%
Two Branch	-	Random Sampling	69.2%	71.4%	74.1%
		Center Cropping	68.9%	72.0%	73.9%
Two Branch	-	Navigation Module	**72.0**%	**74.3**%	**75.9**%

the validation set of Something-Something V1 where we uniformly sample 8 frames. Our navigation module effectively guides the focal branch to concentrate on frames which are more task-relevant and deactivate the frames that contain similar information.

Effect of two branch design. We further incorporate our navigation module into ResNet50 and compare it with AFNet to prove the strength of our designed two-branch architecture. From Table 5, AFNet shows substantial advantages in accuracy under different ratios of select frames. Aside from it, models which adopted our structure but with a fixed sampling policy also show significantly better performance compared with the network based on single branch which can further demonstrate the effectiveness of our two-branch structure and the necessity to preserve the information of all frames.

Effect of navigation module. In this part, we compare our proposed navigation module with three alternative sampling strategies in different selection ratios: (1) random sampling; (2) uniform sampling: sample frames in equal step; (3) normal sampling: sample frames from a standard gaussian distribution. As shown in Table 5, our proposed strategy consistently outperforms other fixed sampling policies under different selection ratios which validates the effectiveness of the navigation module. Moreover, the advantage of our method is more obvious when the ratio of selected frames is small which demonstrates that our selected frames are more task-relevant and contain essential information for the recognition. Further, we evaluate the extension of navigation module which can reduce spatial redundancy, and compare it with: (1) random sampling; (2) center cropping. Our method shows better performance compared with fixed sampling strategies under various selection ratios.

5 Conclusion

This paper proposes an adaptive Ample and Focal Network (AFNet) to reduce temporal redundancy in videos with the consideration of architecture design and the intrinsic redundancy in data. Our method enables 2D-CNNs to have access to more frames to look broadly but with less computation by staying focused on the salient information. AFNet exhibits promising performance as our two-branch design preserves the information of all the input frames instead of discarding part of the knowledge at the beginning of the network. Moreover, the dynamic temporal selection within the network not only restrains the noise of unimportant frames but enforces implicit temporal modeling as well. This enables AFNet to obtain higher accuracy when using fewer frames compared with static methods. We further show that our method can be extended to reduce spatial redundancy by only computing important regions of the selected frames. Extensive experiments have shown that our method outperforms existing works both in accuracy and computational efficiency.

References

[1] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015.

[2] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.

[3] J. Chen, K. Li, Q. Deng, K. Li, and S. Y. Philip. Distributed deep learning model for intelligent video surveillance systems with edge computing. *IEEE Transactions on Industrial Informatics*, 2019.

[4] J. Davidson, B. Liebald, J. Liu, P. Nandy, T. Van Vleet, U. Gargi, S. Gupta, Y. He, M. Lambert, B. Livingston, et al. The youtube video recommendation system. In *ACM conference on Recommender systems*, 2010.

[5] Q. Fan, C.-F. Chen, H. Kuehne, M. Pistoia, and D. Cox. More is less: Learning efficient video representations by big-little network and depthwise temporal aggregation. *arXiv preprint arXiv:1912.00869*, 2019.

[6] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition. In *ICCV*, 2019.

[7] R. Gao, T.-H. Oh, K. Grauman, and L. Torresani. Listen to look: Action recognition by previewing audio. In *CVPR*, 2020.

[8] R. Goyal, S. Ebrahimi Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag, et al. The" something something" video database for learning and evaluating visual common sense. In *ICCV*, 2017.

[9] Y. Han, G. Huang, S. Song, L. Yang, Y. Zhang, and H. Jiang. Spatially adaptive feature refinement for efficient inference. *IEEE Transactions on Image Processing*, 30:9345–9358, 2021.

[10] K. Hara, H. Kataoka, and Y. Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *CVPR*, 2018.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[12] 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.

[13] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

[14] 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.

[15] B. Korbar, D. Tran, and L. Torresani. Scsampler: Sampling salient clips from video for efficient action recognition. In *ICCV*, 2019.

[16] X. Li, Y. Wang, Z. Zhou, and Y. Qiao. Smallbignet: Integrating core and contextual views for video classification. In *CVPR*, 2020.

[17] Y. Li, B. Ji, X. Shi, J. Zhang, B. Kang, and L. Wang. Tea: Temporal excitation and aggregation for action recognition. In *CVPR*, 2020.

[18] J. Lin, C. Gan, and S. Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019.

[19] Z. Liu, L. Wang, W. Wu, C. Qian, and T. Lu. Tam: Temporal adaptive module for video recognition. In *CVPR*, 2021.

[20] J. Materzynska, G. Berger, I. Bax, and R. Memisevic. The jester dataset: A large-scale video dataset of human gestures. In *ICCVW*, 2019.

[21] Y. Meng, C.-C. Lin, R. Panda, P. Sattigeri, L. Karlinsky, A. Oliva, K. Saenko, and R. Feris. Ar-net: Adaptive frame resolution for efficient action recognition. In *ECCV*, 2020.

[22] Y. Meng, R. Panda, C.-C. Lin, P. Sattigeri, L. Karlinsky, K. Saenko, A. Oliva, and R. Feris. Adafuse: Adaptive temporal fusion network for efficient action recognition. *arXiv preprint arXiv:2102.05775*, 2021.

[23] X. Sun, R. Panda, C.-F. R. Chen, A. Oliva, R. Feris, and K. Saenko. Dynamic network quantization for efficient video inference. In *ICCV*, 2021.

[24] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.

[25] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016.

[26] Y. Wang, Z. Chen, H. Jiang, S. Song, Y. Han, and G. Huang. Adaptive focus for efficient video recognition. *arXiv preprint arXiv:2105.03245*, 2021.

[27] Z. Wu, H. Li, C. Xiong, Y.-G. Jiang, and L. S. Davis. A dynamic frame selection framework for fast video recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[28] Z. Wu, C. Xiong, Y.-G. Jiang, and L. S. Davis. Liteeval: A coarse-to-fine framework for resource efficient video recognition. *arXiv preprint arXiv:1912.01601*, 2019.

[29] L. Yang, Y. Han, X. Chen, S. Song, J. Dai, and G. Huang. Resolution adaptive networks for efficient inference. In *CVPR*, 2020.

[30] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, 2016.

[31] B. Zhou, A. Andonian, A. Oliva, and A. Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018.

[32] M. Zolfaghari, K. Singh, and T. Brox. Eco: Efficient convolutional network for online video understanding. In *ECCV*, 2018.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See supplementary material.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See supplementary material.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 4.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 4.
   (b) Did you mention the license of the assets? [Yes] See Section 4.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]