

Frame Flexible Network

Yitian Zhang¹ Yue Bai¹ Chang Liu¹ Huan Wang¹ Sheng Li² Yun Fu¹
¹Northeastern University ²University of Virginia

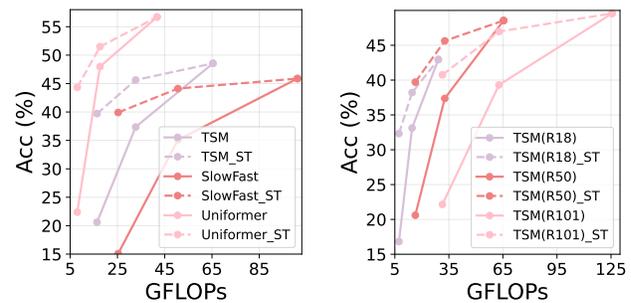
{zhang.yitian, bai.yue, liu.chang6, wang.huan}@northeastern.edu
 shengli@virginia.edu, yunfu@ece.neu.edu

Abstract

Existing video recognition algorithms always conduct different training pipelines for inputs with different frame numbers, which requires repetitive training operations and multiplying storage costs. If we evaluate the model using other frames which are not used in training, we observe the performance will drop significantly (see Fig. 1), which is summarized as Temporal Frequency Deviation phenomenon. To fix this issue, we propose a general framework, named Frame Flexible Network (FFN), which not only enables the model to be evaluated at different frames to adjust its computation, but also reduces the memory costs of storing multiple models significantly. Concretely, FFN integrates several sets of training sequences, involves Multi-Frequency Alignment (MFAL) to learn temporal frequency invariant representations, and leverages Multi-Frequency Adaptation (MFAD) to further strengthen the representation abilities. Comprehensive empirical validations using various architectures and popular benchmarks solidly demonstrate the effectiveness and generalization of FFN (e.g., 7.08/5.15/2.17% performance gain at Frame 4/8/16 on Something-Something V1 dataset over Uniformer). Code is available at <https://github.com/BeSpontaneous/FFN>.

1. Introduction

The growing number of online videos boosts the research on video recognition, laying a solid foundation for deep learning which requires massive data. Compared with image classification, video recognition methods need a series of frames to represent the video which scales the computation. Thus, the efficiency of video recognition methods has always been an essential factor in evaluating these approaches. One existing direction to explore efficiency is designing lightweight networks [9, 40] which are hardware friendly. Even if they increase the efficiency with an acceptable performance trade-off, these methods cannot make further customized adjustments to meet the dynamic-changing resource constraint in real scenarios. In community, there



(a) Temporal Frequency Deviation phenomenon exists in various video recognition architectures.

(b) Temporal Frequency Deviation phenomenon exists in different depths of deep networks.

Figure 1. Temporal Frequency Deviation phenomenon widely exists in video recognition. All methods are trained with high frame number and evaluated at other frames to compare with Separated Training (ST) which individually trains the model at different frames on Something-Something V1 dataset.

are two lines of research being proposed to resolve this issue. The first one is to design networks that can execute at various depths [10] or widths [37] to adjust the computation from the model perspective. The other line of research considers modifying the resolutions of input data [15, 34] to accommodate the cost from the data aspect. However, these methods are carefully designed for 2D CNNs, which may hinder their applications on video recognition where 3D CNNs and Transformer methods are crucial components.

Different from image-related tasks, we need to sample multiple frames to represent the video, and the computational costs will grow proportionally to the number of sampled frames. Concretely, standard protocol trains the same network with different frames separately to obtain multiple models with different performances and computations. This brings challenges to applying these networks on edge devices as the parameters will be multiplied if we store all models, and downloading and offloading models to switch them will cost non-negligible time. Moreover, the same video may be sampled at various temporal rates on different platforms, employing a single network that is trained at a certain frame number for inference cannot resist the variance of frame numbers in real scenarios.

Training the model with a high frame number (i.e., high temporal frequency) and directly evaluating it at fewer frames (i.e., low temporal frequency) to adjust the cost is a naive and straightforward solution. To test its effectiveness, we compare it with Separated Training (ST) which trains the model at different temporal frequency individually and tests it with the corresponding frame. We conduct experiments on 2D-network TSM [18], 3D-network Slow-Fast [6] and Transformer-network Uniformer [16], and find obvious performance gaps between the inference results and ST from Fig. 1, which means these methods will exhibit significantly inferior performance if they are not evaluated at the frame number used in training. Further, we conduct the same experiments on different depths of deep networks and a similar phenomenon appears. We denote this generally existing phenomenon as Temporal Frequency Deviation.

The potential reason for Temporal Frequency Deviation has been explored in Sec. 3 and briefly summarized as the shift in normalization statistics. To address this issue, we propose a general framework, named Frame Flexible Network (FFN), which only requires one-time training, but can be evaluated at multiple frame numbers with great flexibility. We import several input sequences with different sampled frames to FFN during training and propose Multi-Frequency Alignment (MFAL) to learn the temporal frequency invariant representations for robustness towards frame change. Moreover, we present Multi-Frequency Adaptation (MFAD) to further strengthen the representation abilities of the sub-networks which helps FFN to exhibit strong performance at different frames during inference.

Although normalization shifting problem [36, 37] and resolution-adaptive networks [15, 34] have been studied, we stress that designing frame flexible video recognition frameworks to accommodate the costs and save parameters is non-trivial and has practical significance for the following reasons. First, prior works [15, 34] carefully analyzed the detailed structure of 2D convolutions in order to privatize the weights for different scale images. While our method does not touch the specific design of the spatial-temporal modeling components and shares their weights for inputs with different frames. This procedure not only enables our method to be easily applied to various architectures (2D/3D/Transformer models), but also enforces FFN to learn temporal frequency invariant representations. Second, it is, indeed, a common practice to conduct Separated Training (ST) in video recognition, which needs multiplying memory costs to store individual models, and the models are hard to resist the variance in temporal frequency which limits their applications in actual practice. While FFN provides a feasible solution to these challenges which significantly reduces the memory costs of storing multiple models and can be evaluated at different frames to adjust the cost with even higher accuracy compared to ST.

With the proposed framework, we can resolve Temporal Frequency Deviation and enable these methods to adjust their computation based on the current resource budget by sampling different frames, trimming the storage costs of ST remarkably. Moreover, we provide a naive solution that enables FFN to be evaluated at any frame and increases its flexibility during inference. Validation results prove that FFN outperforms ST even at frames that are not used in training. The contributions are summarized as follows:

- We reveal the phenomenon of *Temporal Frequency Deviation* that widely exists in video recognition. It is detailedly analyzed and practically inspires our study.
- We propose a general framework *Frame Flexible Network* (FFN) to resolve Temporal Frequency Deviation. We design Multi-Frequency Alignment (MFAL) to learn temporal frequency invariant representations and present Multi-Frequency Adaptation (MFAD) to further strengthen the representation abilities.
- Comprehensive empirical validations show that FFN, which only requires one-shot training, can adjust its computation by sampling different frames and outperform Separated Training (ST) at different frames on various architectures and datasets, reducing the memory costs of storing multiple models significantly.

2. Related Work

Video Recognition has been extensively explored in recent years and we can summarize the methods into three categories based on their architectures: 1) 2D networks: these methods [17, 18, 29, 30] utilize 2D CNNs as the backbone and specifically design temporal modeling module for spatial-temporal modeling. 2) 3D networks: a straightforward solution for video recognition is to utilize 3D convolutions [2, 6, 28] which naturally consider the temporal information in frame sequences. 3) Transformer networks: based on Vision Transformers [4, 19], many approaches [5, 16, 20] have been proposed recently for spatial-temporal learning and shown powerful performance.

Training-testing Discrepancy widely exists in many scenarios of deep learning. FixRes [27] discovers the deviation of image resolutions between training and testing. Based on this observation, there are methods [15, 34] being designed to train a universal network to fit the images at different resolutions and [35] further extended this idea to 3D CNNs. Slimmable Neural Networks [36, 37] train a shared network which can adjust its width to meet the resource constraints during inference. Different from these prior works, our work is motivated by Temporal Frequency Deviation in video recognition. This finding is essential as frame sampling is a necessary step for all methods and former procedures train the network with different frames individually which is parameter-inefficient and memory-consuming.

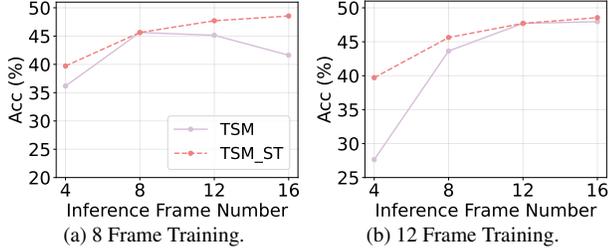


Figure 2. Nearby Alleviation phenomenon. TSM model is trained at 8 Frame and 12 Frame separately on Something-Something V1 dataset and will be evaluated at other frames.

Parameter-efficient Transfer Learning has aroused researchers’ attention in NLP because of the arising of large-scale pre-trained language models. An important research line is to design task-specific adapters [23, 24] to achieve parameter-efficient. Recently, the idea of adapters has been extended to vision tasks as well and shown favorable performance [22, 26, 39]. In this work, instead of focusing on tuning from large-scale pre-trained models, we present Multi-Frequency Adaptation (MFAD) to increase the representation abilities of sub-networks.

Dynamic Networks have been widely studied for efficient video recognition in recent years. Some methods [12, 32, 41] dynamically sample salient frames to reduce temporal redundancy for less cost, while others mainly focus on reducing spatial redundancy by adaptively processing frames with different resolutions [21] or cropping the most salient regions [31] for each frame. Note that these methods are designed to adaptively process every video (e.g., skip frames, crop patches) for efficiency and also requires repetitive training to obtain models with different computation. Our work aims to train a model which can be evaluated at different frames to adjust the costs and reduce the parameters of storing multiple models, while the mentioned dynamic networks do not solve this problem.

3. Temporal Frequency Deviation

Nearby Alleviation. We can observe Temporal Frequency Deviation phenomenon when the models are trained with high frame numbers but evaluated at fewer frames from Fig. 1. To step further, we train TSM [18] at 8/12 Frame and evaluate them at other frames. It is shown in Fig. 2 that there are performance gaps for both models if it is not evaluated with the same frame number which is used in training. Particularly, the discrepancies vary in terms of the value and the performance gap is smaller if the inference frame is close to the training frame number. We denote this phenomenon as Nearby Alleviation because Temporal Frequency Deviation is less severe at nearby frames.

Normalization Shifting. Prior works [36, 37] have studied the problem of normalization shifting in image classification. Specifically, when switching the widths of networks,

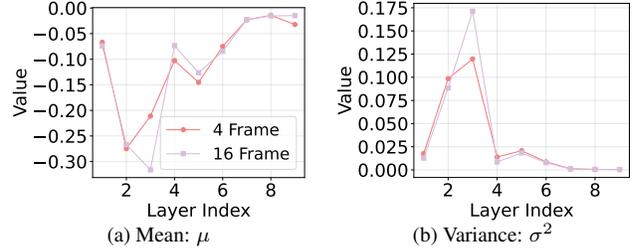


Figure 3. Batch Normalization statistics at various layers. TSM models are trained at 4 Frame and 16 Frame separately, and the statistics are calculated from the fourth stage of ResNet-50.

different numbers of channels will lead to different means and variances of the aggregated features, leading to inconsistency in feature aggregation.

While we do not consider the adjustment in model structure, the problem is whether the difference in frame numbers will cause normalization shifting. If we train the model with v^H which has high temporal frequency and evaluate it with low temporal frequency v^L , the input of Batch Normalization (BN) will be the intermediate feature x^L and the corresponding output is:

$$y^{L'} = \gamma^H \frac{x^L - \mu^H}{\sqrt{\sigma^{H^2} + \epsilon}} + \beta^H, \quad (1)$$

where μ^H, σ^{H^2} are calculated from the data distribution of v^H , and γ^H, β^H are learnt at the training process with v^H . We calculate the statistics of the models trained with v^L and v^H separately and show it in Fig. 3. We can observe a discrepancy of BN statistics at different frame numbers. Note that μ and σ^2 are data-dependent which means that the divergence lies in data intrinsically. Thus, we conjecture that the discrepancy of BN statistics at different frames is an essential factor which leads to Temporal Frequency Deviation. Layer Normalization (LN) [1] has been widely used in Transformer-based models and its statistics are calculated in a similar way with BN which is related to the data distribution. Therefore, we believe the discrepancy of LN statistics is also one of the reasons for Temporal Frequency Deviation on Transformer-based models.

4. Frame Flexible Network

In this section, we first present the training and inference paradigms of Frame Flexible Network (FFN). Then, we propose Multi-Frequency Alignment which is composed of Weight Sharing and Temporal Distillation to learn temporal frequency invariant representations. Further, we introduce Multi-Frequency Adaptation which fits the frequency invariant features to different sub-networks and further increases their representation abilities. Note that FFN is a general framework which can be built on different architectures (shown in Sec. 5.2) and we just take CNN based method as an example in this part for easier description.

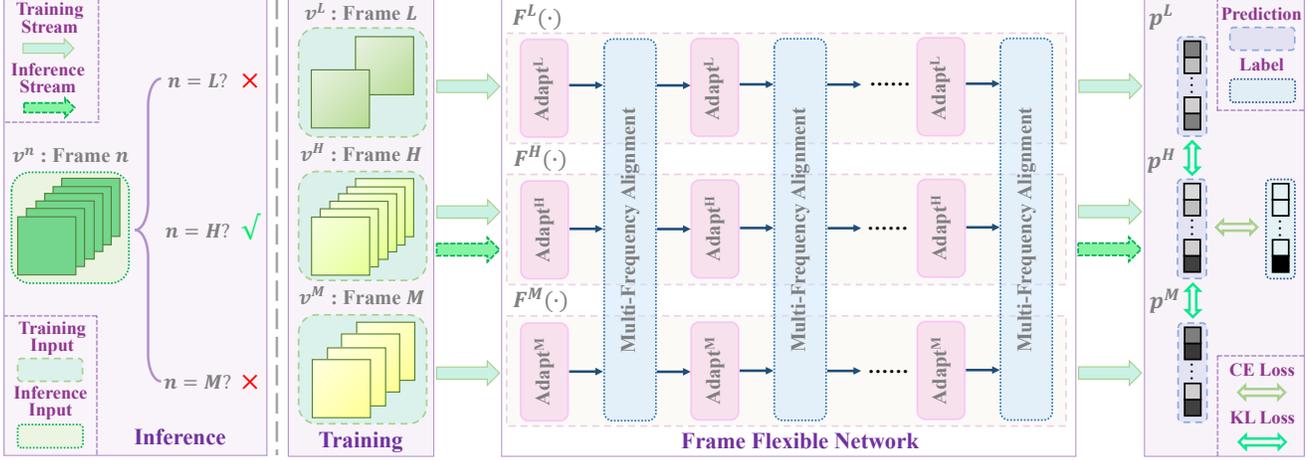


Figure 4. Illustration of Frame Flexible Network (FFN). During training, given inputs with different temporal frequency v^L , v^M and v^H , we propose Multi-Frequency Alignment which involves Weight Sharing and Temporal Distillation for temporal frequency invariant learning. Besides, we present Multi-Frequency Adaptation to fit the temporal invariant features to different sub-networks to further increase the representation abilities. During inference, we activate the sub-network which has the corresponding frame number with the input.

4.1. Framework

The goal of our work is to present a method which can be evaluated at multiple frames and exhibits similar or even better performance compared to Separated Training (ST). Based on the analysis in Sec. 3, Temporal Frequency Deviation will be less severe if the model is evaluated at nearby frames which are used in training. Therefore, we decide to import several sequences with different sampled frames to FFN shown in Fig. 4. Consider video v which is sampled at increasing frame numbers L , M and H , we can obtain v^L , v^M and v^H with temporal frequency of Low, Medium and High, respectively. These three sequences will be utilized at training phase to construct three sub-networks $F^L(\cdot)$, $F^M(\cdot)$ and $F^H(\cdot)$ accordingly. As for the inference paradigm, we will activate the sub-network which has the corresponding frame number with the input. In this manner, we build the computational stream that enables FFN to be evaluated with different frames during inference and adjust the computational costs accordingly.

4.2. Multi-Frequency Alignment

Prior resolution-adaptive networks [15, 34] carefully privatize the weights for 2D convolutions to learn the scale-aware representations for inputs with different resolutions. Recently, there are several works [25, 33] being proposed to maximize the mutual information of the same video at different temporal frequency for contrastive learning in video recognition. The core idea is that the same video instance with different speeds should share high similarity in terms of their discriminative semantics. Inspired by these works, we propose Multi-Frequency Alignment (MFAL) which leverages Weight Sharing and Temporal Distillation to efficiently expand the network and enforce the model to

learn temporal frequency invariant representations.

Weight Sharing. Given video v , we have v^L , v^M , and v^H with an increased temporal frequency and decreased action speed because of the difference in sampled frames. We share the weights of convolutions and classifier across the three sub-networks in order to find a group of parameters θ that mutually model the spatial-temporal relationships for inputs with different temporal frequency:

$$p^* = F^*(v^*; \theta), \quad (2)$$

where $*$ \in $\{L, M, H\}$ and p stands for the predictions. Compared to specialized convolutions, Weight Sharing is parameter-efficient as it only stores one set of weights which can be applied to different input frames. Moreover, it exhibits great potential for better performance (shown in Tab. 4) as it will enforce the model to learn temporal frequency invariant representations which implicitly provides the prior knowledge that the same video with different temporal frequency belongs to the same class, making the model robust to temporal frequency variance.

Temporal Distillation. In most cases, video recognition models trained with v^H have better performance as the network will have access to more information of the original video. Therefore, we consider p^H to be the most ‘accurate’ prediction among the three as v^H has the most sampled frames. Applying Cross-Entropy loss on p^H , we can update the parameters of $F^H(\cdot)$ by:

$$\mathcal{L}_{CE} = - \sum_{k=1}^K \hat{y}_k \log(p_k^H), \quad (3)$$

where \hat{y}_k is the one-hot label of class k and there are K classes in total. Directly calculating CE loss on p^L and p^M is a straightforward solution to update the parameters

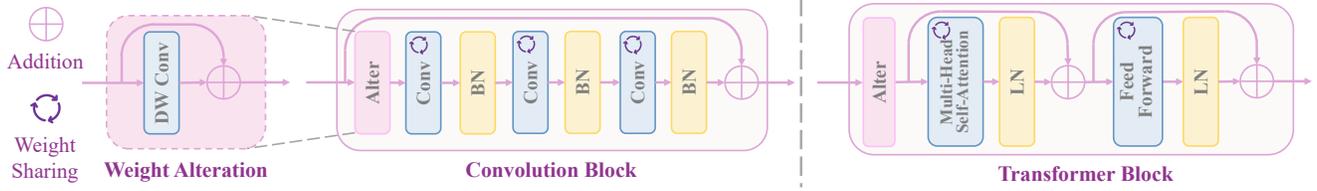


Figure 5. Specific designs of Weight Alteration, Convolution Block, and Transformer Block in Frame Flexible Network (FFN). Weight Alteration is a Depth-wise convolution layer with a residual structure and we insert it into each Convolution Block and Transformer Block.

in $F^L(\cdot)$ and $F^M(\cdot)$, but it will lead to some problems. Firstly, the weights of convolutions are shared across three sub-networks and the optimal parameters for v^L after optimization may not fit well to v^M and v^H . Moreover, optimizing CE loss of p^L and p^M will lead to less favorable parameters of convolutions compared to only calculating Eq. 3 as their inputs contain less information compared to v^H which may lead to inferior performance.

Consequently, we utilize KL divergence [14] loss to involve p^L and p^M in the computational graph and update the parameters of $F^L(\cdot)$ and $F^M(\cdot)$ using:

$$\mathcal{L}_{KL} = -\sum_{k=1}^K p_k^H \log\left(\frac{p_k^M}{p_k^H}\right) - \sum_{k=1}^K p_k^L \log\left(\frac{p_k^L}{p_k^H}\right). \quad (4)$$

As the weights of convolutions are shared across the three sub-networks, optimizing Eq. 4 will enforce the predictions of student (p^L and p^M) and teacher (p^H) networks to be as similar as possible and transfer the good knowledge from $F^H(\cdot)$ to $F^L(\cdot)$ and $F^M(\cdot)$. Considering the two losses in a uniform manner, we update the parameters of FFN by:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda \cdot \mathcal{L}_{KL}, \quad (5)$$

where λ is an introduced hyperparameter to balance the two terms and we simply let $\lambda = 1$ in our implementations without fine-tuning the hyperparameter.

Considering Weight Sharing and Temporal Distillation uniformly, \mathcal{L}_{CE} will provide inter-class supervisory information to enlarge the distance between videos belonging to different classes, and \mathcal{L}_{KL} will further add intra-instance knowledge to the network training, i.e., p^L , p^M and p^H should share high similarity with each other as temporal frequency variance will not change the class of the video. In this way, we not only enforce FFN to learn temporal frequency invariant representations, but also promise it to be easily applied to different structures as we do not touch the specific design of inner spatial-temporal modeling modules.

4.3. Multi-Frequency Adaptation

In the previous section, we propose MFAL to enforce FFN to learn temporal frequency invariant representations. Here, we present Multi-Frequency Adaptation (MFAD) to better fit the frequency invariant features to different sub-networks which further strengthen their representations.

According to our analysis in Sec. 3, normalization shifting is one of the reasons which leads to Temporal Frequency Deviation. Formally, we denote the intermediate features for v^L , v^M and v^H as x^L , x^M and x^H , respectively. Similar with [36, 37], we provide specialized normalization for different input sequences v^L , v^M and v^H :

$$y^* = \gamma^* \frac{x^* - \mu^*}{\sqrt{\sigma^{*2} + \epsilon}} + \beta^*, \quad (6)$$

where $* \in \{L, M, H\}$, and private normalization will learn its own γ and β and calculate the corresponding μ , σ^2 during training. Note that this procedure introduces negligible computation and parameters as normalization operation is a simple transformation and its parameters are often less than 1% of the model size.

Weight Alteration. Though Weight Sharing is necessary for MFAL, it may be difficult to find a set of parameters to display strong representation ability at all frames without further adaptation. Considering a shared convolution with weights W , the outputs of different sequences are:

$$y^* = W \otimes x^*, \quad (7)$$

where \otimes stands for convolution which applies the same transformation for inputs with different temporal frequency. We propose to alter the shared weights of each sub-network to diversify the parameters and strengthen their representation abilities through the transformation:

$$y^* = \phi^* \otimes W \otimes x^*, \quad (8)$$

which can also be written as:

$$y^* = W^* \otimes x^*, \quad W^* = \phi^* \otimes W, \quad (9)$$

where ϕ is a Depth-Wise convolution layer [3] at each Convolution Block which can convert the shared weights W into diversified weights W^* . In this way, we can increase the representation ability of FFN through a simple and efficient transformation. Given that video recognition methods often use pre-trained models, we include the residual structure [8] to avoid the added module breaking the original computational graph of pre-trained models and restore their behaviors. Similarly, we also include Weight Alteration in Transformer Block and we choose the inserted location following [22] shown in Fig. 5. Note that Depth-Wise convolution is lightweight and adding it will introduce negligible parameters and computation.

Method	Specification	Parameters	4 Frame (4F)		8 Frame (8F)		16 Frame (16F)	
			Top-1 Acc. (%)	GFLOPs	Top-1 Acc. (%)	GFLOPs	Top-1 Acc. (%)	GFLOPs
TSM [18]	-	25.6M	20.60	16.4	37.36	32.7	48.55	65.4
● TSM-Mixed	$\rho = 0.50$	25.6M	27.89	16.4	41.07	32.7	48.44	65.4
● TSM-Mixed	$\rho = 0.75$	25.6M	30.43	16.4	42.56	32.7	47.81	65.4
● TSM-Proportional	$\varrho = 0.50$	25.6M	37.56	16.4	44.82	32.7	45.37	65.4
● TSM-Proportional	$\varrho = 0.75$	25.6M	32.06	16.4	43.15	32.7	47.14	65.4
● TSM-Fine-tuning	16F→4F	25.6M	39.95	16.4	40.37	32.7	28.96	65.4
● TSM-Ensemble	-	25.6×3M	35.88	16.4×3	46.25	32.7×3	46.82	65.4×3
● TSM-ST	-	25.6×3M	39.71	16.4	45.63	32.7	48.55	65.4
TSM-FFN	-	25.7M	42.85 (2.90↑)	16.4	48.20 (1.95↑)	32.8	50.79 (2.24↑)	65.5

Table 1. Comparison with baseline methods on Something-Something V1 dataset. GFLOPs stands for the average computational cost to process a single video. The best results are bold-faced, the second best results are marked in color and the improvements are shown.

5. Experiments

In this part, we validate Frame Flexible Network (FFN) on various architectures and benchmarks. First, we provide several baseline solutions and compare them with FFN. Further, we apply our method to different methods and datasets to prove its generalization ability. Moreover, we provide a naive inference paradigm to enable FFN to be evaluated at any frame. Finally, we conduct detailed ablations and analyses to validate the effectiveness of our designs.

5.1. Experiment Settings

Datasets. We conduct experiments on four datasets, including: (1) Something-Something V1&V2 [7] include 98k and 194k videos, respectively. They contain strong temporal dependency and show the most significant Temporal Frequency Deviation phenomenon among all datasets. (2) Kinetics400 [11] is a large-scale dataset with 400 classes. (3) HMDB51 [13] is composed of 6,766 videos which can be categorized into 51 classes. We utilize the original three training/testing splits for training and evaluation.

Implementation Details. We uniformly sample 4/8/16 frames for v^L , v^M and v^H in all methods except for Slow-Fast [6] which samples 16/32/64 frames for fast pathway. For the baseline results, we train all methods with v^H and evaluate them at v^L , v^M and v^H . Separated Training (ST) denotes training the network at v^L , v^M , and v^H individually and evaluating them at the frame used in training.

Baseline Methods. In addition to Separated Training (ST) introduced before, we provide four more baseline methods for this problem: (1) **Mixed Sampling:** We sample 4 and 16 frames for v_i^L and v_i^H , respectively. Then we randomly choose 4 consecutive frames $v_i^{H'}$ from v_i^H and apply mixup [38] to integrate v_i^L into $v_i^{H'}$. The hyperparameter ρ decides the probability of whether to apply Mixed Sampling at each iteration. (2) **Proportional Sampling:** We let the network randomly sample 4 frames or 16 frames at each iteration as this pair has the most significant Temporal Frequency Deviation phenomenon. The hyperparameter ϱ

denotes the probability to sample 16 frames for every iteration. (3) **Fine-tuning:** We first train the model at 16 Frame and then fine-tune it at 4 Frame. (4) **Ensemble:** We make use of the models that are individually trained at 4, 8 and 16 Frame and averagely ensemble them to form a new model.

5.2. Main Results

Comparison with Baseline Methods. Tab. 1 shows that Proportional Sampling and Mixed Sampling help to alleviate Temporal Frequency Deviation as the performance at Frame 4/8 is better than the inference results of the model trained with standard protocol. Nevertheless, the increase is obtained at the cost of an accuracy drop at Frame 16. Then, we adjust the hyperparameter and the results show that both methods seem to provide a trade-off solution for this problem: if the performance at low frames is better, the results at high frame numbers will be worse.

Fine-tuning helps the baseline method to exhibit comparable performance with ST at 4 Frame, but at a loss of forgetting the knowledge at 16 Frame. Ensemble outperforms ST at Frame 8 with the cost of multiplying computation. Besides, its performance at 4/16 Frame is worse than ST which means that it still cannot effectively resolve Temporal Frequency Deviation problem. While FFN shows stronger results compared to ST and Ensemble at all frames with negligible added computation. Moreover, compared to ST and Ensemble which need repetitive training operations and multiplying storage costs, our method is trained for only one time, but can be evaluated at multiple frames, reducing the parameters of saving multiple models significantly which promises its applications on edge devices.

Performance Analysis across Architectures. We further validate FFN on different architectures in Fig. 6. We first build our method on TSM [18] which does not contain any parameters in the temporal modeling module. FFN exhibits advantages in performance at all frames compared to baseline TSM and ST. Then, we implement FFN on TEA [17] which involves convolutions and normalization in the temporal modeling module and our results also surpass ST at all

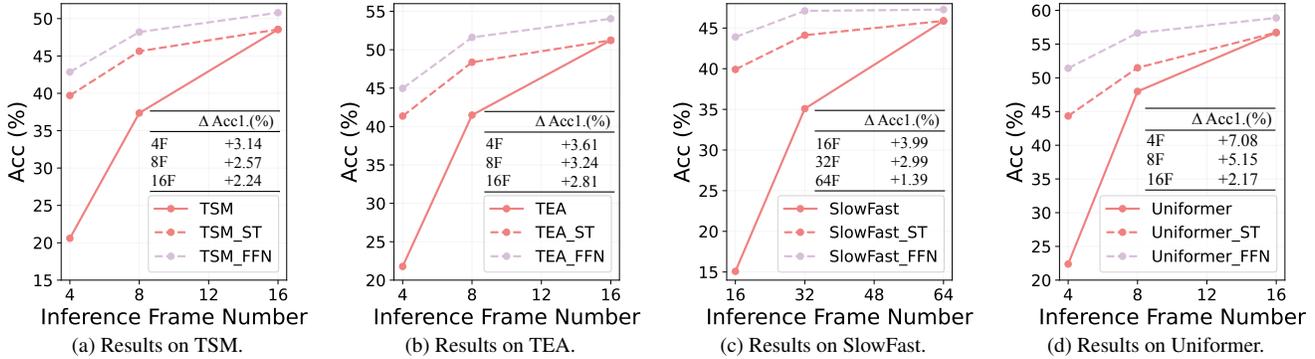


Figure 6. Validation results across different video recognition architectures on Something-Something V1 dataset, including 2D-network, 3D-network and Transformer-network. The improvements of FFN over ST are listed in the table.

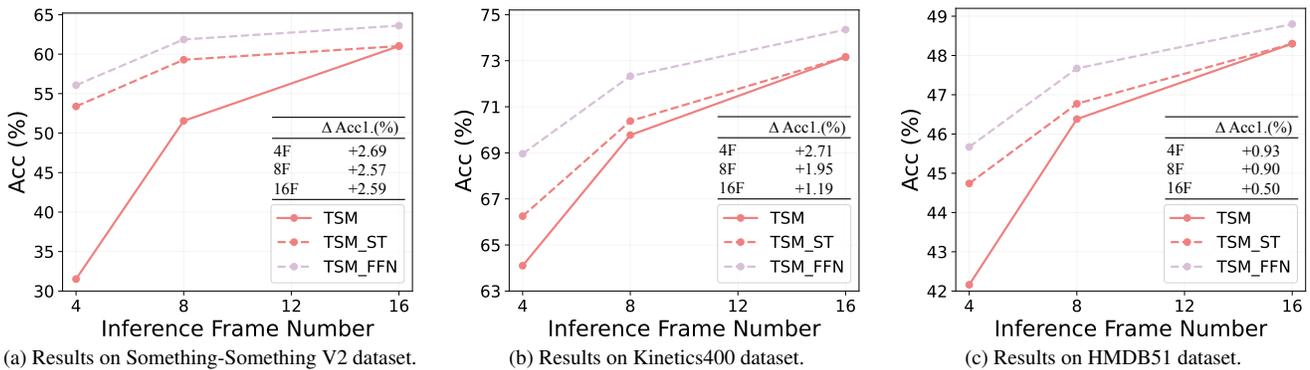


Figure 7. Validation results across various video recognition datasets. The improvements of FFN over ST are listed in the table.

frames. Moreover, we extend FFN to 3D-network: SlowFast [6] and Transformer-network: Uniformer [16]. The results exhibit similar improvements at all frame numbers compared to ST which validates the flexibility and generalization ability of our method. Note that FFN introduces less than 1% extra computation but reduces the memory costs of storing individual models by multiple times.

Performance Analysis across Datasets. In this part, we empirically evaluate FFN on various datasets in Fig. 7, including Something-Something V2, Kinetics400 and HMDB51. The first observation is that Temporal Frequency Deviation phenomenon is less obvious on Kinetics400 and HMDB51 as these two datasets contain less temporal information. Nevertheless, FFN continuously improves the accuracy of ST on these datasets as well. For example, there are 2.71/1.95/1.19% performance gains at Frame 4/8/16 on Kinetics400 which further demonstrate the generalization ability of our design.

5.3. Inference at Any Frame

We have proved that FFN can outperform ST at the frame numbers used in training, but the evaluation at other frames which are not included in training remains untouched. Motivated by Nearby Alleviation in Sec. 3, we provide a naive inference paradigm to enable FFN to be evaluated at any frame. Given a frame number n at inference phase, we will

calculate the frame difference with L , M and H , and activate the sub-network with the minimal difference for validation. If the frame difference is the same for two sub-networks, we will choose the one which corresponds higher frame number by default. In this manner, we can evaluate at other frame numbers which are not used in training.

Inbound Results. Fig. 8 shows that FFN outperforms ST at all frames within the range of 4-16 which are utilized in training. Though the improvement at 12 Frame is less obvious compared to other frames as it is the middle of 8/16 Frame which benefits the least from Nearby Alleviation.

Outbound Results. Moreover, we evaluate FFN at frames that are out of the bound of 4-16. One can observe that FFN even exhibits better performance compared to ST at Frame 2/18/20 which further demonstrates its generalization capacity at non-seen frames.

5.4. Ablation

Input Sequences Combinations. Shown in Tab. 2, we import different numbers of sequences to FFN and evaluate their performance at various frames. First, we can observe that FFN(2) outperforms ST at 4/16 Frame which are used in training, but its performance at 8/12 Frame is worse than ST because of the missing middle sequence in training. In contrast, both FFN(3) and FFN(4) obtain higher accuracy at all frames compared to ST which can be attributed to

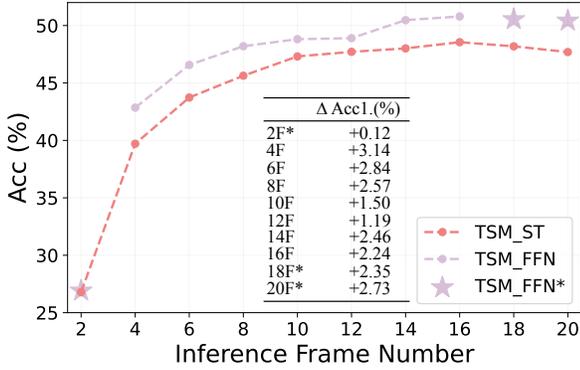


Figure 8. Validation results of FFN at various frame numbers on Something-Something V1. The improvements of FFN over ST are listed in the table. Outbound results are denoted with *.

Method	Sequences	Top-1 Acc. (%)			
		4F	8F	12F	16F
TSM-ST [18]	-	39.71	45.63	47.71	48.55
TSM-FFN(2)	4/16	41.69	37.93	48.10	49.79
TSM-FFN(3)	4/8/16	42.85	48.20	48.90	50.79
TSM-FFN(4)	4/8/12/16	43.40	48.66	49.77	50.63

Table 2. Experiments of different input sequences combinations on Something-Something V1. The best results are bold-faced.

the utilization of the middle sequence in training so that the Temporal Frequency Deviation at nearby frames can be mitigated by Nearby Alleviation. FFN(4) obtains better results at Frame 4/8/12 because of the added sequence, but it will cost more time and resources during training.

Frame Numbers. We sample more frames to v^H in this section and import four sequences with 4/8/16/24 frames to FFN, respectively. The first observation from Tab. 3 is that the performance of TSM-ST (24F) is even a little bit lower than TSM-ST (16F) which can be attributed to relatively simple temporal modeling measure of TSM. However, FFN still obtains better performance compared with ST at all frames and achieves the highest accuracy at 24 Frame, owing to the design of Temporal Distillation.

Design Choices. We conduct ablation to verify the effectiveness of our designs in Tab. 4. First, we build FFN with shared normalization and one can observe an obvious performance drop at 4/16 Frame due to the shift in normalization statistics. Then, we remove Weight Alteration in the convolution block and it exhibits worse performance at all frames which proves the strength of Multi-Frequency Adaptation (MFAD) as it increases the representation abilities of the sub-networks at corresponding frames. Further, we optimize FFN by calculating CE loss on the predictions of all sub-networks respectively and do not utilize KL divergence loss for optimization. The results are clearly worse than ST which suggests that Temporal Distillation is a necessary component for Multi-Frequency Alignment (MFAL). Finally, we specialize convolutions (w/o WS) and note that

Method	Parameters	Top-1 Acc. (%)			
		4F	8F	16F	24F
TSM-ST [18]	25.6×4M	39.71	45.63	48.55	47.90
TSM-FFN	25.7M	41.28	46.72	49.79	49.95

Table 3. Validation results of FFN with more sampled frames on Something-Something V1. The best results are bold-faced.

Method	Parameters	Specification	Top-1 Acc. (%)		
			4F	8F	16F
TSM-ST [18]	25.6×3M	-	39.71	45.63	48.55
TSM-FFN	25.7M	w/o SN	35.79	46.80	44.62
TSM-FFN	25.6M	w/o WA	41.91	47.92	49.84
TSM-FFN	25.7M	w/o TD	39.61	45.65	48.07
TSM-FFN	25.6×3M	w/o WS	41.51	47.16	48.23
TSM-FFN	25.7M	-	42.85	48.20	50.79

Table 4. Ablation of design choices of FFN on Something-Something V1. SN, TD, WA, WS denotes Specialized Normalization, Temporal Distillation, Weight Alteration, Weight Sharing, respectively. The best results are bold-faced.

this operation will multiply the parameters by three times. We can observe both FFN and FFN(w/o WA) outperform specialized convolutions at all frames with fewer parameters which demonstrates the effectiveness of MFAL to learn temporal frequency invariant representations.

6. Conclusion and Limitations

In this paper, we reveal Temporal Frequency Deviation phenomenon and propose Frame Flexible Network (FFN) to address it. Specifically, we propose Multi-Frequency Alignment to learn temporal frequency invariant representations and present Multi-Frequency Adaptation to further strengthen the representation ability. Extensive experiments demonstrate that FFN, which only requires one-shot training, can be evaluated at multiple frames and outperforms Separated Training with significantly fewer parameters, making it favorable for applications on edge devices.

One limitation of FFN is that, it requires more GPU memory during training as we import several input sequences. Second, FFN introduces slightly extra computation because of Weight Alteration. In future work, we are interested in improving the training efficiency of FFN.

Acknowledgment

Research was sponsored by the DEVCOM Analysis Center and was accomplished under Cooperative Agreement Number W911NF-22-2-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 3
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2
- [3] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 2017. 5
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [5] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 2
- [6] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 2, 6, 7
- [7] 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*, 2017. 6
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [9] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1
- [10] 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. *arXiv preprint arXiv:1703.09844*, 2017. 1
- [11] 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] Bruno Korbar, Du Tran, and Lorenzo Torresani. Scsampler: Sampling salient clips from video for efficient action recognition. In *ICCV*, 2019. 3
- [13] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, 2011. 6
- [14] Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997. 5
- [15] Duo Li, Anbang Yao, and Qifeng Chen. Learning to learn parameterized classification networks for scalable input images. In *ECCV*, 2020. 1, 2, 4
- [16] Kunchang Li, Yali Wang, Junhao Zhang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unifying convolution and self-attention for visual recognition. *arXiv preprint arXiv:2201.09450*, 2022. 2, 7
- [17] Yan Li, Bin Ji, Xintian Shi, Jianguo Zhang, Bin Kang, and Limin Wang. Tea: Temporal excitation and aggregation for action recognition. In *CVPR*, 2020. 2, 6
- [18] Ji Lin, Chuhan Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 2, 3, 6, 8
- [19] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 2
- [20] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, 2022. 2
- [21] 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*, 2020. 3
- [22] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning for action recognition. *arXiv preprint arXiv:2206.13559*, 2022. 3, 5
- [23] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. *arXiv preprint arXiv:2005.00247*, 2020. 3
- [24] Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. Adapterhub: A framework for adapting transformers. *arXiv preprint arXiv:2007.07779*, 2020. 3
- [25] Ankit Singh, Omprakash Chakraborty, Ashutosh Varshney, Rameswar Panda, Rogerio Feris, Kate Saenko, and Abir Das. Semi-supervised action recognition with temporal contrastive learning. In *CVPR*, 2021. 4
- [26] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. VI-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *CVPR*, 2022. 3
- [27] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. *NeurIPS*, 2019. 2
- [28] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 2
- [29] Limin Wang, Zhan Tong, Bin Ji, and Gangshan Wu. Tdn: Temporal difference networks for efficient action recognition. In *CVPR*, 2021. 2
- [30] 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*, 2016. 2
- [31] Yulin Wang, Zhaoxi Chen, Haojun Jiang, Shiji Song, Yizeng Han, and Gao Huang. Adaptive focus for efficient video recognition. *arXiv preprint arXiv:2105.03245*, 2021. 3

- [32] 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*, 2020. 3
- [33] Ceyuan Yang, Yinghao Xu, Bo Dai, and Bolei Zhou. Video representation learning with visual tempo consistency. *arXiv preprint arXiv:2006.15489*, 2020. 4
- [34] Taojiannan Yang, Sijie Zhu, Chen Chen, Shen Yan, Mi Zhang, and Andrew Willis. Mutualnet: Adaptive convnet via mutual learning from network width and resolution. In *ECCV*, 2020. 1, 2, 4
- [35] Taojiannan Yang, Sijie Zhu, Matias Mendieta, Pu Wang, Ravikumar Balakrishnan, Minwoo Lee, Tao Han, Mubarak Shah, and Chen Chen. Mutualnet: Adaptive convnet via mutual learning from different model configurations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [36] Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *ICCV*, 2019. 2, 3, 5
- [37] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*, 2018. 1, 2, 3, 5
- [38] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 6
- [39] Jinnian Zhang, Houwen Peng, Kan Wu, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Minivit: Compressing vision transformers with weight multiplexing. In *CVPR*, 2022. 3
- [40] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018. 1
- [41] Yitian Zhang, Yue Bai, Huan Wang, Yi Xu, and Yun Fu. Look more but care less in video recognition. *arXiv preprint arXiv:2211.09992*, 2022. 3