



Published in final edited form as:

Proc IEEE Int Conf Comput Vis. 2021 October ; 2021: 13557–13567. doi:10.1109/iccv48922.2021.01332.

VidTr: Video Transformer Without Convolutions

Yanyi Zhang^{1,2,*}, Xinyu Li^{1,*}, Chunhui Liu¹, Bing Shuai¹, Yi Zhu¹, Biagio Brattoli¹, Hao Chen¹, Ivan Marsic², Joseph Tighe¹

¹ Amazon Web Service

² Rutgers University

Abstract

We introduce Video Transformer (VidTr) with separable-attention for video classification. Comparing with commonly used 3D networks, VidTr is able to aggregate spatio-temporal information via stacked attentions and provide better performance with higher efficiency. We first introduce the vanilla video transformer and show that transformer module is able to perform spatio-temporal modeling from raw pixels, but with heavy memory usage. We then present VidTr which reduces the memory cost by 3.3× while keeping the same performance. To further optimize the model, we propose the standard deviation based topK pooling for attention ($\text{pool}_{\text{topK_std}}$), which reduces the computation by dropping non-informative features along temporal dimension. VidTr achieves state-of-the-art performance on five commonly used datasets with lower computational requirement, showing both the efficiency and effectiveness of our design. Finally, error analysis and visualization show that VidTr is especially good at predicting actions that require long-term temporal reasoning.

1. Introduction

We introduce Video Transformer (VidTr) with separable-attention, one of the first transformer-based video action classification architecture that performs global spatio-temporal feature aggregation. Convolution-based architectures have dominated the video classification literature in recent years [19, 32, 55], and although successful, the convolution-based approaches have two drawbacks: 1. they have limited receptive field on each layer and 2. information is slowly aggregated through stacked convolution layers, which is inefficient and might be ineffective [31, 55]. Attention is a potential candidate to overcome these limitations as it has a large receptive field which can be leveraged for spatio-temporal modeling. Previous works use attention to modeling long-range spatio-temporal features in videos but still rely on convolutional backbones [31, 55]. Inspired by recent successful applications of transformers on NLP [12, 52] and computer vision [14, 47], we propose a transformer-based video network that directly applies attentions on raw video pixels for video classification, aiming at higher efficiency and better performance.

yzaws@amazon.com .
*Equally Contributed.

We first introduce a vanilla video transformer that directly learns spatio-temporal features from raw-pixel inputs via vision transformer [14], showing that it is possible to perform pixel-level spatio-temporal modeling. However, as discussed in [56], the transformer has $\mathcal{O}(n^2)$ complexity with respect to the sequence length. The vanilla video transformer is memory consuming, as training on a 16-frame clip (224×224) with only batch size of 1 requires more than 16GB GPU memory, which makes it infeasible on most commercial devices. Inspired by the R(2+1)D convolution that breaks down 3D convolution kernel to a spatial kernel and a temporal kernel [50], we further introduce our separable-attention, which performs spatial and temporal attention separately. This reduces the memory consumption by $3.3\times$ with no drop in accuracy. We can further reduce the memory and computational requirements of our system by exploiting the fact that a large portion of many videos have redundant information temporally. This notion has been explored in the context of convolutional networks to reduce computation previously [32]. We build on this intuition and propose a standard deviation based topK pooling operation (*topK_std* pooling), which reduces the sequence length and encourages the transformer network to focus on representative frames.

We evaluated our VidTr on 6 most commonly used datasets, including Kinetics 400/700, Charades, Something-something V2, UCF-101 and HMDB-51. Our model achieved state-of-the-art (SOTA) or comparable performance on five datasets with lower computational requirements and latency compared to previous SOTA approaches. Our error analysis and ablation experiments show that the VidTr works significantly better than I3D on activities that requires longer temporal reasoning (e.g. making a cake vs. eating a cake), which aligns well with our intuition. This also inspires us to ensemble the VidTr with the I3D convolutional network as features from global and local modeling methods should be complementary. We show that simply combining the VidTr with a I3D50 model (8 frames input) via ensemble can lead to roughly a 2% performance improvement on Kinetics 400. We further illustrate how and why the VidTr works by visualizing the separable-attention using attention rollout [1], and show that the spatial-attention is able to focus on informative patches while temporal attention is able to reduce the duplicated/non-informative temporal instances. Our contributions are:

1. **Video transformer:** We propose to efficiently and effectively aggregate spatio-temporal information with stacked attentions as opposed to convolution based approaches. We introduce vanilla video transformer as proof of concept with SOTA comparable performance on video classification.
2. **VidTr:** We introduce VidTr and its permutations, including the VidTr with SOTA performance and the compact-VidTr with significantly reduced computational costs using the proposed standard deviation based pooling method.
3. **Results and model weights:** We provide detailed results and analysis on 6 commonly used datasets which can be used as reference for future research. Our pre-trained model can be used for many down-streaming tasks.

2. Related Work

2.1. Action Classification

The early research on video based action recognition relies on 2D convolutions [28]. The LSTM [25] was later proposed to model the image feature based on ConvNet features [30, 51, 63]. However, the combination of ConvNet and LSTM did not lead to significantly better performance. Instead of relying on RNNs, the segment based method TSN [54] and its permutations [22, 35, 64] were proposed with good performance.

Although 2D network was proved successful, the spatio-temporal modeling was still separated. Using 3D convolution for spatio-temporal modeling was initially proposed in [26] and further extended to the C3D network [48]. However, training 3D convnet from scratch was hard, initializing the 3D convnet weights by inflate from 2D networks was initially proposed in I3D [7] and soon proved applicable with different type of 2D network [10, 24, 58]. The I3D was used as backbone for many following work including two-stream network [19, 55], the networks with focus on temporal modeling [31, 32, 59], and the 3D networks with refined 3D convolution kernels [27, 33, 39, 44].

The 3D networks are proved effective but often not efficient, the 3D networks with better performance often requires larger kernels or deeper structures. The recent research demonstrates that depth convolution significantly reduce the computation [49], but depth convolution also increase the network inference latency. TSM [37] and TAM [17] proposed a more efficient backbone for temporal modeling, however, such design couldn't achieve SOTA performance on Kinetics dataset. The neural architecture search was proposed for action recognition [18, 43] recently with competitive performance, however, the high latency and limited generalizability remain to be improved.

The previous methods heavily rely on convolution to aggregate features spatio-temporally, which is not efficient. A few previous work tried to perform global spatio-temporal modeling [31, 55] but still limited by the convolution backbone. The proposed VidTr is fundamentally different from previous works based on convolutions, the VidTr doesn't require heavily stacked convolutions [59] for feature aggregation but efficiently learn feature globally via attention from first layer. Besides, the VidTr don't rely on sliding convolutions and depth convolutions, which runs at less FLOPs and lower latency compared with 3D convolutions [18, 59].

2.2. Vision Transformer

The transformers [52] was previously proposed for NLP tasks [13] and recently adopted for computer vision tasks. The transformers were roughly used in three different ways in previous works: 1. To bridge the gap between different modalities, e.g. video captioning [65], video retrieval [20] and dialog system [36]. 2. To aggregate convolutional features for down-streaming tasks, e.g. object detection [5, 11], pose estimation [61], semantic segmentation [15] and action recognition [21]. 3. To perform feature learning on raw pixels, e.g. most recently image classification [14, 47].

Action recognition with self-attention on convolution features [21] is proved successful, however, convolution also generates local feature and gives redundant computations. Different from [21] and inspired by very recent work on applying transformer on raw pixels [14, 47], we pioneer the work on aggregating spatio-temporal feature from raw videos without relying on convolution features. Different from very recent work [41] that extract spatial feature with vision transformer on every video frames and then aggregate feature with attention, our proposed method jointly learns spatio-temporal feature with lower computational cost and higher performance. Our work differs from the concurrent work [4], we present a split attention with better performance without requiring larger video resolution nor extra long clip length. Some more recent work [2, 4, 4, 16, 40, 42] further studied the multi-scale and different attention factorization methods.

3. Video Transformer

We introduce the Video Transformer starting with the vanilla video transformer (section 3.1) which illustrates our idea of video action recognition without convolutions. We then present VidTr by first introducing separable-attention (section 3.2), and then the attention pooling to drop non-representative information temporally (section 3.2).

3.1. Vanilla Video Transformer

Following previous efforts in NLP [13] and image classification [14], we adopt the transformer [52] encoder structure for action recognition that operates on raw pixels. Given a video clip $V \in \mathbb{R}^{C \times T \times W \times H}$, where T denotes the clip length, W and H denote the video frame width and height, and C denotes the number of channel, we first convert V to a sequence of $s \times s$ spatial patches, and apply a linear embedding to each patch, namely $S \in \mathbb{R}^{T \frac{H}{s} \frac{W}{s} \times C'}$, where C' is the channel dimension after the linear embedding. We add a 1D learnable positional embedding [13, 14] to S and following previous work [13, 14], append a class token as well, whose purpose is to aggregate features from the whole sequence for classification. This results in $S' \in \mathbb{R}^{(\frac{TWH}{s^2} + 1) \times C'}$, where $S'_0 \in \mathbb{R}^{1 \times C'}$ is the attached class token. S' is feed into our transformer encoder structure detailed next.

As Figure 1 middle shows, we expand the previous successful ViT transformer architecture for 3D feature learning. Specifically, we stack 12 encoder layers, with each encoder layer consisting of an 8-head self-attention layer and two dense layers with 768 and 3072 hidden units. Different from transformers for 2D images, each attention layer learns a spatio-temporal affinity map $Attn \in \mathbb{R}^{(\frac{TWH}{s^2} + 1) \times (\frac{TWH}{s^2} + 1)}$.

3.2. VidTr

In Table 2 we show that this simple formulation is capable of learning 3D motion features on a sequence of local patches. However, as explained in [3], the affinity attention matrix $Attn \in \mathbb{R}^{(\frac{TWH}{s^2} + 1) \times (\frac{TWH}{s^2} + 1)}$ needs to be stored in memory for back propagating, and thus the memory consumption is quadratically related to the sequence length. We can see that the vanilla video transformer increases memory usage for the affinity map from

$\mathcal{O}(W^2H^2)$ to $\mathcal{O}(T^2W^2H^2)$, leading to $T^2\times$ memory usage for training, which makes it impractical on most available GPU devices. We now address this inefficiency with a separable attention architecture.

3.2.1 Separable-Attention—To address these memory constraints, we introduce a multi-head separable-attention (MSA) by decoupling the 3D self-attention to a spatial attention MSA_s and a temporal attention MSA_t (Figure 1):

$$MSA(S) = MSA_s(MSA_t(S)) \quad (1)$$

Different from the vanilla video transformer that applies 1D sequential modeling on S , we decouple S to a 2D sequence $\hat{S} \in \mathbb{R}^{(T+1) \times (\frac{WH}{s^2} + 1) \times C'}$ with positional embedding and two types of class tokens that append additional tokens along the spatial and temporal dimensions. Here, the spatial class tokens gather information from spatial patches in a single frame using spatial attention, and the temporal class tokens gather information from patches across frames (at same location) using temporal attention. Then the intersection of the spatial and temporal class tokens $\hat{S}^{(0,0,\cdot)}$ is used for the final classification. To decouple 1D self-attention functions on 2D sequential features \hat{S} , we first operate on each spatial location (i) independently, applying temporal attention as:

$$\hat{S}_t^{(:,i,\cdot)} = MSA_t(k = q = v = \hat{S}^{(:,i,\cdot)}) \quad (2)$$

$$= pool(Attn_t) \cdot v_t \quad (3)$$

$$= pool(\text{Softmax}(q_t \cdot k_t^T)) \cdot v_t \quad (4)$$

where $\hat{S}_t \in \mathbb{R}^{(\tau+1) \times (\frac{WH}{s^2} + 1) \times C}$ is the output of MSA_t , $pool$ denotes the down-sampling method to reduce temporal dimension (from T to τ , $\tau = T$ when no down-sampling is performed) that will be described later, q_t , k_t and v_t denote key, query, and value features after applying independent linear functions (LN) on \hat{S} :

$$q_t = \text{LN}_{q_t}(\hat{S}^{(:,i,\cdot)}); k_t = \text{LN}_{k_t}(\hat{S}^{(:,i,\cdot)}); v_t = \text{LN}_{v_t}(\hat{S}^{(:,i,\cdot)}) \quad (5)$$

Moreover, $Attn_t \in \mathbb{R}^{(\tau+1) \times (T+1)}$ represent a temporal attention obtained from matrix multiplication between q_t and k_t . Following MSA_t , we apply a similar 1D sequential self-attention MSA_s on spatial dimension:

$$\hat{S}_{st}^{(i,\cdot,\cdot)} = MSA_s(k = q = v = \hat{S}_t^{(i,\cdot,\cdot)}) \quad (6)$$

$$= Attn_s \cdot v_s \quad (7)$$

$$= \text{Softmax}(q_s \cdot k_s^T) \cdot v_s \quad (8)$$

where $\hat{S}_{st} \in \mathbb{R}^{(\tau+1) \times (\frac{WH}{s^2} + 1) \times C}$ is the output of MSA_s , q_s , k_s , and v_s denotes key, query, and value features after applying independent linear functions on \hat{S}_t .

$\text{Attn}_s \in \mathbb{R}^{(\frac{WH}{s^2} + 1) \times (\frac{WH}{s^2} + 1)}$ represent a spatial-wise affinity map. We do not apply down-sampling on the spatial attention as we saw a significant performance drop in our preliminary experiments.

Our spatio-temporal split attention decreased the memory usage of the transformer layer by reducing the affinity matrix from $\mathcal{O}(T^2W^2H^2)$ to $\mathcal{O}(\tau^2 + W^2H^2)$. This allows us to explore longer temporal sequence lengths that were infeasible on modern hardware with the vanilla transformer.

3.2.2 Temporal Down-sampling method—Video content usually contains redundant information [31], with multiple frames depicting near identical content over time. We introduce compact VidTr (C-VidTr) by applying temporal down-sampling within our transformer architecture to remove some of this redundancy. We study different temporal down-sampling methods (*pool* in Eq. 3) including temporal average pooling and 1D convolutions with stride 2, which reduce the temporal dimension by half (details in Table 5d).

A limitation of these pooling the methods is that they uniformly aggregate information across time but often in video clips the informative frames are not uniformly distributed. We adopted the idea of non-uniform temporal feature aggregation from previous work [31]. Different from previous work [31] that directly down-sample the query using average pooling, we found that in our proposed network, the temporal attention highly activates on a small set of temporal features when the clip is informative, while the attention equally distributed over the length of the clip when the clip carries little additional semantic information. Building on this intuition, we propose a topK based pooling (*topK_std* pooling) that orders instances by the standard deviation of each row in the attention matrix:

$$\text{pool}_{\text{topK_std}}(\text{Attn}_t^{(1:,:,:)}) = \text{Attn}_t^{(\text{topK}(\sigma(\text{Attn}_t^{(1:,:,:)}), :), :)} \quad (9)$$

where $\sigma \in \mathbb{R}^T$ is row-wise standard deviation of $\text{Attn}_t^{(1:,:,:)}$ as:

$$\sigma^{(i)} = \frac{1}{T} \sqrt{\sum_{i=1}^T (\text{Attn}_t^{(1:,:,:)}) - \mu^2} \quad (10)$$

$$\mu^{(i)} = \frac{1}{T} \sum_{i=1}^T \text{Attn}_t^{(i,:)} \quad (11)$$

where $\mu \in \mathbb{R}^T$ is the mean of $Attn_t^{(1:\cdot, \cdot)}$. Note that the *topK_std* pooling was applied to the affinity matrix excludes the token ($Attn_t^{(0:\cdot, \cdot)}$) as we always preserve token for information aggregation. Our experiments show that *topK_std* pooling gives better performance than average pooling or convolution. The *topK_std* pooling can be intuitively understood as selecting the frames with strong localized attention and removing frames with uniform attention.

3.3. Implementation Details

Model Instantiating: Based on the input clip length and sample rate, we introduce three base VidTr models (VidTr-S, VidTr-M and VidTr-L). By applying the different pooling strategies we introduce two compact VidTr permutations (C-VidTr-S, and C-VidTr-M). To normalize the feature space, we apply layer normalization before and after the residual connection of each transformer layer and adopt the GELU activation as suggested in [14]. Detailed configurations can be found in Table 1. We empirically determined the configuration for different clip length to produce a set of models from low FLOPs and low latency to high accuracy (details in Ablations).

During **training** we initialize our model weights from ViT-B [14]. To avoid over fitting, we adopted the commonly used augmentation strategies including random crop, random horizontal flip (except for Something-something dataset). We trained the model using 64 Tesla V100 GPUs, with batch size of 6 per-GPU (for VidTr-S) and weight decay of $1e-5$. We adopted SGD as the optimizer but found the Adam optimizer also gives us the same performance. We trained our network for 50 epochs in total with initial learning rate of 0.01, and reduced it by 10 times after epochs 25 and 40. It takes about 12 hours for VidTr-S model to converge, the training process also scales well with fewer GPUs (e.g. 8 GPUs for 4 days). During **inference** we adopted the commonly used 30-crop evaluation for VidTr and compact VidTr, with 10 uniformly sampled temporal segments and 3 uniformly sampled spatial crop on each temporal segment [55]. It is worth mentioning that we can further boost the inference speed of compact VidTr by adopting a single pass inference mechanism, this is because the attention mechanism captures global information more effectively than 3D convolution. We do this by training a model with frames sampled in TSN [54] style, and uniformly sampling N frames in inference (details in supplemental materials).

4. Experimental Results

4.1. Datasets

We evaluate our method on six of the most widely used datasets. **Kinetics 400** [8] and **Kinetics 700** [6] consists of approximately 240K/650K training videos and 20K/35K validation videos trimmed to 10 seconds from 400/700 human action categories. We report top-1 and top-5 classification accuracy on the validation sets. **Something-Something V2** [23] dataset consists of 174 actions and contains 168.9K training videos and 24.7K evaluation videos. We report top-1 accuracy following previous works [37] evaluation setup. **Charades** [45] has 9.8k training videos and 1.8k validation videos spanning about 30 seconds on average. Charades contains 157 multi-label classes with longer activities,

performance is measured in mean Average Precision (mAP). **UCF-101** [46] and **HMDB-51** [29] are two smaller datasets. UCF-101 contains 13320 videos with an average length of 180 frames per video and 101 action categories. The HMDB-51 has 6,766 videos and 51 action categories. We report the top-1 classification on the validation videos based on split 1 for both dataset.

4.2. Kinetics 400 Results

4.2.1 Comparison To SOTA—We report results on the validation set of Kinetics 400 in Table 2, including the top-1 and top-5 accuracy, GFLOPs (Giga Floating-Point Operations) and latency (ms) required to compute results on one view.

As shown in Table 2, the VidTr achieved the SOTA performance compared to previous I3D based SOTA architectures with lower GFLOPs and latency. The VidTr significantly outperform previous SOTA methods at roughly same computational budget, e.g. at 200 GFLOPs, the VidTr-M outperform I3D50 by 3.6%, NL50 by 2.1%, and TPN50 by 0.9%. At similar accuracy levels, VidTr is significantly more computationally efficient than other works, e.g. at 78% top-1 accuracy, the VidTr-S has 6× fewer FLOPs than NL-101, 2× fewer FLOPs than TPN and 12% fewer FLOPs than Slowfast-101. We also see that our VidTr outperforms I3D based networks at higher sample rate (e.g. $s = 8$, TPN achieved 76.1% top-1 accuracy), this denotes, the global attention learns temporal information more effectively than 3D convolutions. X3D-XXL from architecture search is the only network that outperforms our VidTr. We plan to use architecture search techniques for attention based architecture in future work.

4.2.2 Compact VidTr—We evaluate the effectiveness of our compact VidTr with the proposed temporal down-sampling method (Table 1). The results (Table 3) show that the proposed down-sampling strategy removes roughly 56% of the computation required by VidTr with only 2% performance drop in accuracy. The compact VidTr complete the VidTr family from small models (only 39GFLOPs) to high performance models (up to 79.1% accuracy). Compared with previous SOTA compact models [34, 39], our compact VidTr achieves better or similar performance with lower FLOPs and latency, including: TEA (+0.6% with 16% fewer FLOPs) and TEINet (+0.5% with 11% fewer FLOPs).

4.2.3 Error and Ensemble Analysis—We compare the errors made by VidTr-S and the I3D50 network to better understand the local networks’ (I3D) and global networks’ (VidTr) behavior. We provide the top-5 activities that our VidTr-S gain most significant improvement over the I3D50. We find that our VidTr-S outperformed the I3D on the activities that requires long-term video contexts to be recognized. For example, our VidTr-S outperformed the I3D50 on “making a cake” by 26% in accuracy. The I3D50 overfits to “cakes” and often recognize making a cake as eating a cake. We also analyze the top-5 activities where I3D does better than our VidTr-S (Table 4). Our VidTr-S performs poorly on the activities that need to capture fast and local motions. For example, our VidTr-S performs 21% worse in accuracy on “shaking head”.

Inspired by the findings in our error analysis, we ensembled our VidTr with a light weight I3D50 network by averaging the output values between the two networks. The results (Table

2) show that the the I3D model and transformer model complements each other and the ensemble model roughly lead to 2% performance improvement on Kinetics 400 with limited additional FLOPs (37G). The performance gained by ensembling the VidTr with I3D is significantly better than the improvement by combine two 3D networks (Table 2).

4.2.4 Ablations—We perform all ablation experiments with our VidTr-S model on Kinetics 400. We used $8 \times 224 \times 224$ input with a frame sample rate of 8, and 30-view evaluation.

Patching strategies: We first compare the cubic patch (4×16^2), where the video is represented as a sequence of spatio-temporal patches, with the square patch (1×16^2), where the video is represented as a sequence of spatial patches. Our results (Table 5a) show that the model using cubic patches with longer temporal size has fewer FLOPs but results in significant performance drop (73.1 vs. 75.5). The model using square patches significantly outperform all cubic patch based models, likely because the linear embedding is not enough to represent the shot-term temporal association in the cubic. We further compared the performance of using different patch sizes (1×16^2 vs. 1×32^2), using 32^2 patches lead to 4× decreasing of the sequence length, which decreases memory consumption of the affinity matrices by 16×, however, using 16^2 patches significantly outperform the model using 32^2 patches (77.7 vs. 71.2). We did not evaluate the model using smaller patching sizes (e.g., 8×8) because of the high memory consumption.

Attention Factorization: We compare different factorization for attention design, including spatial modeling only (WH), jointly spatio-temporal modeling module (WHT, vanilla-Tr), spatio-temporal separable-attention (WH + T, VidTr), and axial separable-attention (W + H + T). We first evaluate an spatio-only transformer. We average the class token for each input frame for our final output. Our results (Table 5b) show that the spatio-only transformer requires less memory but has worse performance compare with spatio-temporal attention models. This shows that temporal modeling is critical for attention based architectures. The joint spatio-temporal transformer significantly outperforms the spatio-only transformer but requires a restrictive amount of memory (T^2 times for the affinity matrices). Our VidTr using spatio-temporal separable-attention requires 3.3× less memory with no accuracy drop. We further evaluate the axial separable-attention (W + H + T), which requires the least memory. The results (Table 5b) show that the axial separable-attention has a significant performance drop likely due to breaking the X and Y spatial dimensions.

Sequence down-sampling comparison: We compare different down-sampling strategy including temporal average pooling, 1D temporal convolution and the proposed STD-based topK pooling method. The results (Table 5d) show that our proposed STD-based down-sampling method outperformed the temporal average pooling and the convolution-based down-sampling strategies that uniformly aggregate information over time.

Backbone generalization: We evaluate our VidTr initialized with different models, including T2T [62], ViT-B, and ViT-L. The results on Table 5c show that our VidTr achieves reasonable performance across all backbones. The VidTr using T2T as the backbone has the lowest FLOPs but also the lowest accuracy. The ViT-L-based VidTr achieve similar

performance with the Vit-B-based VidTr even with $3\times$ FLOPs. As showed in previous work [14], transformer-based network are more likely to over-fit and Kinetics-400 is relatively small for Vit-L-based VidTr.

Where to down-sample: Finally we study where to perform temporal down-sampling. We perform temporal down-sampling at different layers (Table 5e). Our results (Table 5e) show that starting to perform down-sampling after the first encoder layer has the best trade-off between the performance and FLOPs. Starting to perform down-sampling at very beginning leads to the fewest FLOPs but has a significant performance drop (72.9 vs. 74.9). Performing down-sampling later only has slight performance improvement but requires higher FLOPs. We then analyze how many layers to skip between two down-sample layers. Based on the results in Table 5f, skipping one layer between two down-sample operations has the best trade-off. Performing down-sampling on consecutive layers (0 skip layers) has lowest FLOPs but the performance decreases (73.9 vs. 74.9). Skipping more layers did not show significant performance improvement but does have higher FLOPs.

4.2.5 Run-time Analysis—We further analyzed the trade-off between latency, FLOPs and accuracy. We note that the VidTr achieved the best balance between these factors (Figure 2). The VidTr-S achieve similar performance but significantly fewer FLOPs compare with I3D101-NL ($5\times$ fewer FLOPs), Slowfast101 8×8 (12% fewer FLOPs), TPN101 ($2\times$ fewer FLOPs), and CorrNet50 ($20\times$ fewer FLOPs). Note that the X3D has very low FLOPs but high latency due to the use of depth convolution. Our experiments show that the X3D-L has about $3.6\times$ higher latency comparing with VidTr-S (Figure 2).

4.3. More Results

Kinetics-700 Results: Our experiments show a consistent performance trend on Kinetics 700 (Table 6). The VidTr-S significantly outperformed the baseline I3D model (+9%), the VidTr-M achieved the performance comparable to Slowfast101 8×8 and the VidTr-L is comparable to previous SOTA slowfast101-nonlocal. There is a small performance gap between our model and Slowfast-NL [19], because Slowfast is pre-trained on both Kinetics 400 and 600 while we only pre-trained on Kinetics 400. Previous findings that VidTr and I3D are being complementary is consistent on Kinetics 700, ensemble VidTr-L with I3D leads to +0.6% performance boost.

Charades Results: We compare our VidTr with previous SOTA models on Charades. Our VidTr-L outperformed previous SOTA methods LFB and NUTA101, and achieved the performance comparable to Slowfast101-NL (Table 6). The results on Charades demonstrates that our VidTr generalizes well to multi-label activity datasets. Our VidTr performs worse than the current SOTA networks (X3D-XL) on Charades likely due to overfitting. As discussed in previous work [14], the transformer-based networks overfit easier than convolution-based models, and Charades is relatively small. We observed a similar finding with our ensemble, ensembling our VidTr with a I3D network (40.3 mAP) achieved SOTA performance.

Something-something V2 Results: We observe that the VidTr does not work well on the something-something dataset (Table 6), likely because pure transformer based approaches do not model local motion as well as convolutions. This aligns with our observation in our error analysis. Further improving local motion modeling ability is an area of future work.

UCF and HMDB Results: Finally we train our VidTr on two small dataset UCF-101 and HMDB-51 to test if VidTr generalizes to smaller datasets. The VidTr achieved SOTA comparable performance with 6 epochs of training (96.6% on UCF and 74.4% on HMDB), showing that the model generalize well on small dataset (Table 6).

5. Visualization and Understanding VidTr

We first visualized the VidTr's separable-attention with attention roll-out method [1] (Figure 3a). We find that the spatial attention is able to focus on informative regions and temporal attention is able to skip the duplicated/non-representative information temporally. We then visualized the attention at 4th, 8th and 12th layer of VidTr (Figure 3b), we found the spatial attention is stronger on deeper layers. The attention does not capture meaningful temporal instances at early stages because the temporal feature relies on the spatial information to determine informative temporal instances. Finally we compared the I3D activation map and rollout attention from VidTr (Figure 3c). The I3D misclassified the catching fish as sailing, as the I3D attention focused on the people sitting behind and water. The VidTr is able to make the correct prediction and the attention showed that the VidTr is able to focus on the action related regions across time.

6. Conclusion

In this paper, we present video transformer with separable-attention, an novel stacked attention based architecture for video action recognition. Our experimental results show that the proposed VidTr achieves state-of-the-art or comparable performance on five public action recognition datasets. The experiments and error analysis show that the VidTr is especially good at modeling the actions that requires long-term reasoning. Further combining the advantage of VidTr and convolution for better local-global action modeling [38, 57] and adopt self-supervised training [9] on large-scaled data will be our future work.

Acknowledgements.

We thank NSF grant IIS-1763827 for supporting Yanyi Zhang to study as a Ph.D. student at Rutgers University.

References

- [1]. Abnar Samira and Zuidema Willem. Quantifying attention flow in transformers. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 4190–4197, 2020. 2, 8
- [2]. Arnab Anurag, Dehghani Mostafa, Heigold Georg, Sun Chen, Lu i Mario, and Schmid Cordelia. Vivit: A video vision transformer. arXiv preprint arXiv:2103.15691, 2021. 2
- [3]. Beltagy Iz, Peters Matthew E, and Cohan Arman. Longformer: The long-document transformer. arXiv preprint arXiv:2004.05150, 2020. 3

- [4]. Bertasius Gedas, Wang Heng, and Torresani Lorenzo. Is space-time attention all you need for video understanding? arXiv preprint arXiv:2102.05095, 2021. 2
- [5]. Carion Nicolas, Massa Francisco, Synnaeve Gabriel, Usunier Nicolas, Kirillov Alexander, and Zagoruyko Sergey. End-to-end object detection with transformers. In European Conference on Computer Vision, pages 213–229. Springer, 2020. 2
- [6]. Carreira Joao, Noland Eric, Hillier Chloe, and Zisserman Andrew. A short note on the kinetics-700 human action dataset. arXiv preprint arXiv:1907.06987, 2019. 5
- [7]. Carreira Joao and Zisserman Andrew. Quo vadis, action recognition? a new model and the kinetics dataset. In proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 6299–6308, 2017. 2, 7
- [8]. Carreira J and Zisserman A. Quo vadis, action recognition? a new model and the kinetics dataset. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 4724–4733, July 2017. 5
- [9]. Chen Xinlei, Xie Saining, and He Kaiming. An empirical study of training self-supervised visual transformers. arXiv preprint arXiv:2104.02057, 2021. 8
- [10]. Chen Yunpeng, Kalantidis Yannis, Li Jianshu, Yan Shuicheng, and Feng Jiashi. Multi-fiber networks for video recognition. In Proceedings of the european conference on computer vision (ECCV), pages 352–367, 2018. 2
- [11]. Dai Zhigang, Cai Bolun, Lin Yugeng, and Chen Junying. Up-detr: Unsupervised pre-training for object detection with transformers. arXiv preprint arXiv:2011.09094, 2020. 2
- [12]. Devlin Jacob, Chang Ming-Wei, Lee Kenton, and Toutanova Kristina. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805, 2018. 1
- [13]. Devlin Jacob, Chang Ming-Wei, Lee Kenton, and Toutanova Kristina. Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186, 2019. 2, 3
- [14]. Dosovitskiy Alexey, Beyer Lucas, Kolesnikov Alexander, Weissenborn Dirk, Zhai Xiaohua, Unterthiner Thomas, Dehghani Mostafa, Minderer Matthias, Heigold Georg, Gelly Sylvain, et al. An image is worth 16×16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020. 1, 2, 3, 4, 6, 7, 8
- [15]. Duke Brendan, Ahmed Abdalla, Wolf Christian, Aarabi Parham, and Taylor Graham W. Sstvos: Sparse spatiotemporal transformers for video object segmentation. arXiv preprint arXiv:2101.08833, 2021. 2
- [16]. Fan Haoqi, Xiong Bo, Mangalam Karttikeya, Li Yanghao, Yan Zhicheng, Malik Jitendra, and Feichtenhofer Christoph. Multiscale vision transformers. arXiv preprint arXiv:2104.11227, 2021. 2
- [17]. Fan Quanfu, Chen Chun-Fu, Kuehne Hilde, Pistoia Marco, and Cox David. More is less: Learning efficient video representations by big-little network and depthwise temporal aggregation. arXiv preprint arXiv:1912.00869, 2019. 2
- [18]. Feichtenhofer Christoph. X3d: Expanding architectures for efficient video recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 203–213, 2020. 2, 5, 6, 7
- [19]. Feichtenhofer Christoph, Fan Haoqi, Malik Jitendra, and He Kaiming. Slowfast networks for video recognition. In Proceedings of the IEEE International Conference on Computer Vision, pages 6202–6211, 2019. 1, 2, 5, 7, 8
- [20]. Gabeur Valentin, Sun Chen, Alahari Karteek, and Schmid Cordelia. Multi-modal transformer for video retrieval. In European Conference on Computer Vision (ECCV), volume 5. Springer, 2020. 2
- [21]. Girdhar Rohit, Carreira Joao, Doersch Carl, and Zisserman Andrew. Video action transformer network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 244–253, 2019. 2
- [22]. Girdhar Rohit, Ramanan Deva, Gupta Abhinav, Sivic Josef, and Russell Bryan. ActionVLAD: Learning SpatioTemporal Aggregation for Action Classification. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017. 2

- [23]. Goyal Raghav, Kahou Samira Ebrahimi, Michalski Vincent, Materzynska Joanna, Westphal Susanne, Kim Heuna, Haenel Valentin, Freund Ingo, Yianilos Peter, Mueller-Freitag Moritz, et al. The "something something" video database for learning and evaluating visual common sense. In ICCV, volume 1, page 3, 2017. 5
- [24]. Hara Kensho, Kataoka Hirokatsu, and Satoh Yutaka. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, pages 6546–6555, 2018. 2
- [25]. Hochreiter Sepp and Schmidhuber Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2 [PubMed: 9377276]
- [26]. Ji Shuiwang, Xu Wei, Yang Ming, and Yu Kai. 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):221–231, 2012. 2
- [27]. Jiang Boyuan, Wang MengMeng, Gan Weihao, Wu Wei, and Yan Junjie. STM: SpatioTemporal and Motion Encoding for Action Recognition. In The IEEE International Conference on Computer Vision (ICCV), 2019. 2
- [28]. Karpathy Andrej, Toderici George, Shetty Sanketh, Leung Thomas, Sukthankar Rahul, and Fei-Fei Li. Large-scale video classification with convolutional neural networks. In Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, pages 1725–1732, 2014. 2
- [29]. Kuehne Hildegard, Jhuang Hueihan, Garrote Estíbaliz, Poggio Tomaso, and Serre Thomas. Hmdb: a large video database for human motion recognition. In 2011 International Conference on Computer Vision, pages 2556–2563. IEEE, 2011. 5
- [30]. Li Qing, Qiu Zhaofan, Yao Ting, Mei Tao, Rui Yong, and Luo Jiebo. Action recognition by learning deep multi-granular spatio-temporal video representation. In Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval, pages 159–166, 2016. 2
- [31]. Li Xinyu, Liu Chunhui, Shuai Bing, Zhu Yi, Chen Hao, and Tighe Joseph. Nuta: Non-uniform temporal aggregation for action recognition. *arXiv preprint arXiv:2012.08041*, 2020. 1, 2, 4
- [32]. Li Xinyu, Shuai Bing, and Tighe Joseph. Directional temporal modeling for action recognition. In European Conference on Computer Vision, pages 275–291. Springer, 2020. 1, 2, 5
- [33]. Li Yan, Ji Bin, Shi Xintian, Zhang Jianguo, Kang Bin, and Wang Limin. TEA: Temporal Excitation and Aggregation for Action Recognition. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2
- [34]. Li Yan, Ji Bin, Shi Xintian, Zhang Jianguo, Kang Bin, and Wang Limin. Tea: Temporal excitation and aggregation for action recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 909–918, 2020. 5, 6
- [35]. Li Yingwei, Li Weixin, Mahadevan Vijay, and Vasconcelos Nuno. VLAD3: Encoding Dynamics of Deep Features for Action Recognition. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016. 2
- [36]. Li Zekang, Li Zongjia, Zhang Jinchao, Feng Yang, Niu Cheng, and Zhou Jie. Bridging text and video: A universal multimodal transformer for video-audio scene-aware dialog. *arXiv preprint arXiv:2002.00163*, 2020. 2
- [37]. Lin Ji, Gan Chuang, and Han Song. Tsm: Temporal shift module for efficient video understanding. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 7083–7093, 2019. 2, 5, 6, 7
- [38]. Liu Ze, Lin Yutong, Cao Yue, Hu Han, Wei Yixuan, Zhang Zheng, Lin Stephen, and Guo Baining. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 8
- [39]. Liu Zhaoyang, Luo Donghao, Wang Yabiao, Wang Limin, Tai Ying, Wang Chengjie, Li Jilin, Huang Feiyue, and Lu Tong. TEINet: Towards an Efficient Architecture for Video Recognition. In The Conference on Artificial Intelligence (AAAI), 2020. 2, 5, 6, 7
- [40]. Liu Ze, Ning Jia, Cao Yue, Wei Yixuan, Zhang Zheng, Lin Stephen, and Hu Han. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 2
- [41]. Neimark Daniel, Bar Omri, Zohar Maya, and Asselmann Dotan. Video transformer network. *arXiv preprint arXiv:2102.00719*, 2021. 2

- [42]. Patrick Mandela, Campbell Dylan, Asano Yuki M, Metze Ishan Misra Florian, Feichtenhofer Christoph, Vedaldi Andrea, Henriques Jo, et al. Keeping your eye on the ball: Trajectory attention in video transformers. arXiv preprint arXiv:2106.05392, 2021. 2
- [43]. Piergiorganni AJ, Angelova Anelia, and Ryoo Michael S. Tiny video networks. arXiv preprint arXiv:1910.06961, 2019. 2
- [44]. Shao Hao, Qian Shengju, and Liu Yu. Temporal interlacing network. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, pages 11966–11973, 2020. 2
- [45]. Sigurdsson Gunnar A., Varol Gül, Wang Xiaolong, Laptev Ivan, Farhadi Ali, and Gupta Abhinav. Hollywood in homes: Crowdsourcing data collection for activity understanding. ArXiv e-prints, 2016. 5
- [46]. Soomro Khurram, Zamir Amir Roshan, and Shah M. A dataset of 101 human action classes from videos in the wild. Center for Research in Computer Vision, 2012. 5
- [47]. Touvron Hugo, Cord Matthieu, Douze Matthijs, Massa Francisco, Sablayrolles Alexandre, and Jégou Hervé. Training data-efficient image transformers & distillation through attention. arXiv preprint arXiv:2012.12877, 2020. 1, 2
- [48]. Tran Du, Bourdev Lubomir, Fergus Rob, Torresani Lorenzo, and Paluri Manohar. Learning spatiotemporal features with 3d convolutional networks. In Proceedings of the IEEE international conference on computer vision, pages 4489–4497, 2015. 2
- [49]. Tran Du, Wang Heng, Torresani Lorenzo, and Feiszli Matt. Video classification with channel-separated convolutional networks. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 5552–5561, 2019. 2, 7
- [50]. Tran Du, Wang Heng, Torresani Lorenzo, Ray Jamie, LeCun Yann, and Paluri Manohar. A closer look at spatiotemporal convolutions for action recognition. In Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, pages 6450–6459, 2018. 1
- [51]. Ullah Amin, Ahmad Jamil, Muhammad Khan, Sajjad Muhammad, and Baik Sung Wook. Action recognition in video sequences using deep bi-directional lstm with cnn features. IEEE access, 6:1155–1166, 2017. 2
- [52]. Vaswani Ashish, Shazeer Noam, Parmar Niki, Uszkoreit Jakob, Jones Llion, Aidan N Gomez Lukasz Kaiser, and Polosukhin Illia. Attention is all you need. arXiv preprint arXiv:1706.03762, 2017. 1, 2, 3
- [53]. Wang Heng, Tran Du, Torresani Lorenzo, and Feiszli Matt. Video modeling with correlation networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 352–361, 2020. 5
- [54]. Wang Limin, Xiong Yuanjun, Wang Zhe, Qiao Yu, Lin Dahua, Tang Xiaoou, and Gool Luc Van. Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. In The European Conference on Computer Vision (ECCV), 2016. 2, 5
- [55]. Wang Xiaolong, Girshick Ross, Gupta Abhinav, and He Kaiming. Non-local neural networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 7794–7803, 2018. 1, 2, 5
- [56]. Wu Chao-Yuan, Feichtenhofer Christoph, Fan Haoqi, He Kaiming, Krahenbuhl Philipp, and Girshick Ross. Long-term feature banks for detailed video understanding. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 284–293, 2019. 1
- [57]. Wu Haiping, Xiao Bin, Codella Noel, Liu Mengchen, Dai Xiyang, Yuan Lu, and Zhang Lei. Cvt: Introducing convolutions to vision transformers. arXiv preprint arXiv:2103.15808, 2021. 8
- [58]. Xie Saining, Girshick Ross, Dollár Piotr, Tu Zhuowen, and He Kaiming. Aggregated residual transformations for deep neural networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1492–1500, 2017. 2
- [59]. Yang Ceyuan, Xu Yinghao, Shi Jianping, Dai Bo, and Zhou Bolei. Temporal Pyramid Network for Action Recognition. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2, 7
- [60]. Yang Ceyuan, Xu Yinghao, Shi Jianping, Dai Bo, and Zhou Bolei. Temporal pyramid network for action recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 591–600, 2020. 5

- [61]. Yang Sen, Quan Zhibin, Nie Mu, and Yang Wankou. Transpose: Towards explainable human pose estimation by transformer. arXiv preprint arXiv:2012.14214, 2020. 2
- [62]. Yuan Li, Chen Yunpeng, Wang Tao, Yu Weihao, Shi Yujun, Tay Francis EH, Feng Jiashi, and Yan Shuicheng. Tokens-to-token vit: Training vision transformers from scratch on imagenet. arXiv preprint arXiv:2101.11986, 2021. 6, 7
- [63]. Ng Joe Yue-Hei, Hausknecht Matthew, Vijayanarasimhan Sudheendra, Vinyals Oriol, Monga Rajat, and Toderici George. Beyond short snippets: Deep networks for video classification. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4694–4702, 2015. 2
- [64]. Zhou Bolei, Andonian Alex, Oliva Aude, and Torralba Antonio. Temporal Relational Reasoning in Videos. In The European Conference on Computer Vision (ECCV), 2018. 2
- [65]. Zhou Luowei, Zhou Yingbo, Corso Jason J, Socher Richard, and Xiong Caiming. End-to-end dense video captioning with masked transformer. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 8739–8748, 2018. 2

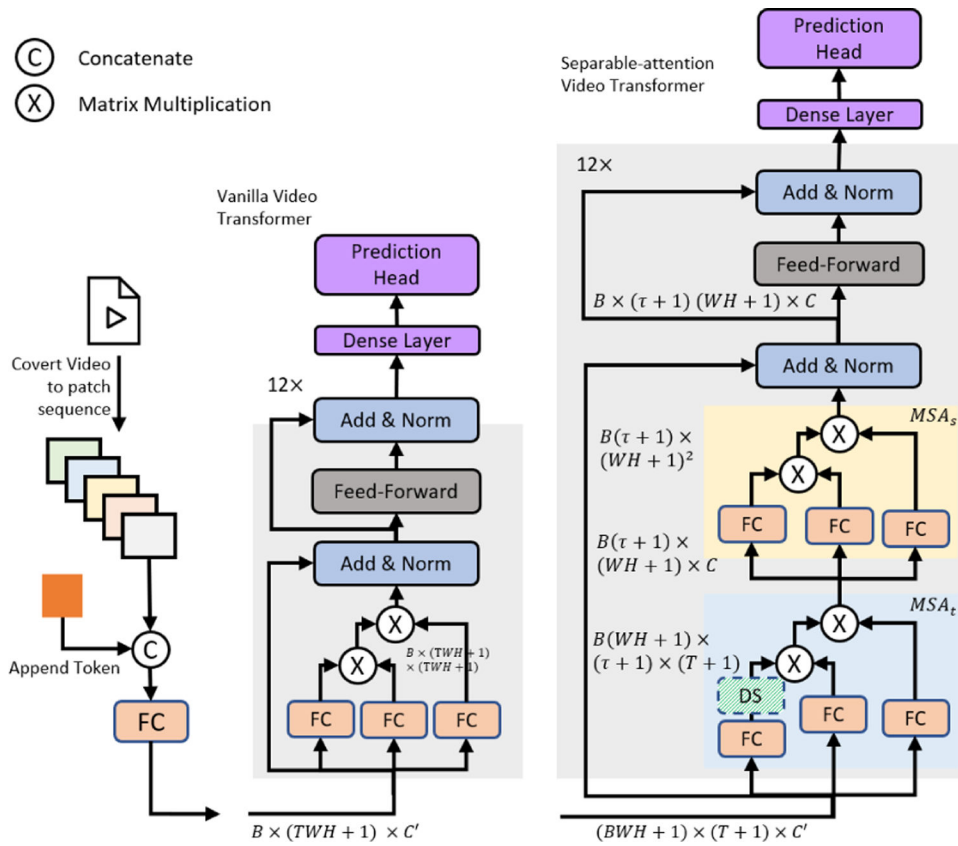


Figure 1: Spatio-temporal separable-attention video transformer (VidTr). The model takes pixels patches as input and learns the spatial temporal feature via proposed separable-attention. The green shaded block denotes the down-sample module which can be inserted into VidTr for higher efficiency. τ denotes the temporal dimension after downsampling.






Figure 2:
The comparison between different models on accuracy, FLOPs and latency.



(a) The spatial and temporal attention in VidTr. The attention is able to focus on the informative frames and regions.



(b) The rollout attentions from different layers of VidTr.



(c) Comparison of I3D activations and VidTr attentions.

Figure 3: Visualization of spatial and temporal attention of VidTr and comparison with I3D activation.

Table 1:

Detailed configuration of different VidTr permutations. *clip_len* denotes the sampled clip length and *sr* stands for the sample rate. We uniformly sample *clip_len* frames out of $clip_len \times sr$ consecutive frames. The configurations are empirically selected, details in Ablations.

Model	<i>clip_len</i>	<i>sr</i>	Down-sample Layer	τ
VidTr-S	8	8	-	-
VidTr-M	16	4	-	-
VidTr-L	32	2	-	-
C-VidTr-S	8	8	[1,2,4]	[6,4,2]
C-VidTr-M	16	4	[1,2,4]	[8,4,2]

Table 2:

Results on Kinetics-400 dataset. We report top-1 accuracy(%) on the validation set. The ‘Input’ column indicates what frames of the 64 frame clip are actually sent to the network. $n \times s$ input indicates we feed n frames to the network sampled every s frames. Lat. stands for the latency on single crop.

Model	Input	GFLOPs	Lat.	top-1	top-5
I3D50 [60]	32×2	167	74.4	75.0	92.2
I3D101 [60]	32×2	342	118.3	77.4	92.7
NL50 [55]	32×2	282	53.3	76.5	92.6
NL101 [55]	32×2	544	134.1	77.7	93.3
TEA50 [34]	16×2	70	-	76.1	92.5
TEINet [39]	16×2	66	49.5	76.2	92.5
CIDC [32]	32×2	101	82.3	75.5	92.1
SF50 8×8 [19]	$(32+8) \times 2$	66	49.3	77.0	92.6
SF101 8×8 [19]	$(32+8) \times 2$	106	71.9	77.5	92.3
SF101 16×8 [19]	$(64+16) \times 2$	213	124.3	78.9	93.5
TPN50 [60]	32×2	199	89.3	77.7	93.3
TPN101 [60]	32×2	374	133.4	78.9	93.9
CorrNet50 [53]	32×2	115	-	77.2	N/A
CorrNet101 [53]	32×2	187	-	78.5	N/A
X3D-XXL [18]	16×5	196	-	80.4	94.6
Vanilla-Tr	8×8	89	32.8	77.5	93.2
VidTr-S	8×8	89	36.2	77.7	93.3
VidTr-M	16×4	179	61.1	78.6	93.5
VidTr-L	32×2	351	110.2	79.1	93.9
En-I3D-50-101	32×2	509	192.7	77.7	93.2
En-I3D-TPN-101	32×2	541	207.8	79.1	94.0
En-VidTr-S	8×8	130	73.2	79.4	94.0
En-VidTr-M	16×4	220	98.1	79.7	94.2
En-VidTr-L	32×2	392	147.2	80.5	94.6

Table 3:

Comparison of VidTr to other fast networks. We present the number of views used for evaluation and FLOPs required for each view. The latency denotes the total time required to get the reported top-1 score.¹

Model	Input	Res.	GFLOPs	Latency(ms)	top-1
TSM [37]	8/TSN	256	69	29	74.7
TEA [34]	16×4	256	70	-	76.1
3DEff-B4 [18]	16×5	224	7	-	72.4
TEINet [39]	16×4	256	33	36	74.9
X3D-M [18]	16×5	224	5	40.9	74.6
X3D-L [18]	16×5	312	19	59.4	76.8
C-VidTr-S	8×8	224	39	17.5	75.7
C-VidTr-M	16×4	224	59	26.1	76.7

¹we measure latency of X3D using the authors' code and fast depth convolution patch: <https://github.com/facebookresearch/SlowFast/blob/master/projects/x3d/README.md>, which only has models for X3D-M and X3D-L and not the XL and XXL variants

Table 4:

Quantitative analysis on Kinetics-400 dataset. The performance gain is defined as the disparity of the top-1 accuracy between VidTr network and that of I3D.

Top 5 (+)	Acc. gain	Top 5 (-)	Acc. gain
making a cake	+26.0%	shaking head	-21.7%
catching fish	+21.2%	dunking basketball	-20.8%
catching baseball	+20.8%	lunge	-19.9%
stretching arm	+19.1%	playing guitar	-19.9%
spraying	+18.0 %	tap dancing	-16.3%

(a) Top 5 classes that VidTr works better than I3D.

(b) Top 5 classes that I3D works better than VidTr.

Table 5:

Ablation studies on Kinetics 400 dataset. We use an VidTr-S backbone with 8 frames input for (a,b) and C-VidTr-S for (c,d). The evaluation is performed on 30 views with 8 frame input unless specified. FP. stands for FLOPs.

Model	FP.	top-1	Model	Mem.	top-1
Cubic (4×16^2)	23G	73.1	WH	2.1GB	74.7
Cubic (2×16^2)	45G	75.5	WHT	7.6GB	77.5
Square (1×16^2)	89G	77.7	WH + T	2.3GB	77.7
Square (1×32^2)	21G	71.2	W + H + T.	1.5GB	72.3

(a) Comparison between different patching strategies. (b) Comparison between different factorization.

Init. from	FP.	top-1	Configurations	top-1	top-5
T2T [62]	34G	76.3	Temp. Avg. Pool.	74.9	91.6
ViT-B [14]	89G	77.7	1D Conv. [62]	75.4	92.3
ViT-L [14]	358	77.5	STD Pool.	75.7	92.2

(c) Comparison between different backbones. (d) Comparison between different down-sample methods.

Layer	τ	FP.	top-1	Layer	τ	FP.	top-1
[0, 2]	[4, 2]	26G	72.9	[1, 2]	[4, 2]	30G	73.9
[1,3]	[4, 2]	32G	74.9	[1,3]	[4, 2]	32G	74.9
[2, 4]	[4, 2]	47G	74.9	[1,4]	[4, 2]	33G	75.0
[6, 8]	[4, 2]	60G	75.3	[1,5]	[4, 2]	34G	75.2

(e) Compact VidTr down-sampling twice at layer k and $k + 2$. (f) Compact VidTr down-sampling twice starting from layer 1 and skipping different number of layers.

Table 6:

Results on Kinetics-700 dataset (K700), Charades dataset (Chad), something-something-V2 dataset (SS), UCF-101 and HMDB (HM) dataset. The evaluation metrics are mean average precision (mAP) in percentage for Charades (32×4 input is used), top-1 accuracy for Kinetics 700, something-something-V2 (TSN styled dataloader is used), UCF and HMDB.

Model	Input	K700	Chad	SS	UCF	HM
I3D [7]	32×2	58.7	32.9	50.0	95.1	74.3
TSM [37]	8(TSN)	-	-	59.3	94.5	70.7
I3D101 [59]	32 × 4	40.3	-	-	-	-
CSN152 [49]	32 × 2	70.1	-	-	-	-
TEINet[39]	16 (TSN)	-	-	62.1	96.7	73.3
SF101 [19]	64×2	70.2	-	60.9	-	-
SF101-NL [19]	64×2	70.6	45.2	-	-	-
X3D-XL [18]	16 × 5	-	47.1	-	-	-
VidTr-M	16 × 4	69.5	-	61.9	96.6	74.4
VidTr-L	32 × 2	70.2	43.5	63.0	96.7	74.4
En-VidTr-L	32 × 2	70.8	47.3	-	-	-