

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

10-2021

Token shift transformer for video classification

ZHANG Hao

Yanbin. HAO

Chong-wah NGO

Singapore Management University, cwngo@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

ZHANG Hao; HAO, Yanbin.; and NGO, Chong-wah. Token shift transformer for video classification. (2021). *Proceedings of the 29th ACM International Conference on Multimedia, October 20-24*. 917-925.

Available at: https://ink.library.smu.edu.sg/sis_research/6807

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Token Shift Transformer for Video Classification

Hao Zhang

City University of Hong Kong
Hong Kong SAR, China
zhanghaoinf@gmail.com

Yanbin Hao*

University of Science and Technology
of China
Hefei, China
haoyanbin@hotmail.com

Chong-Wah Ngo

Singapore Management University
Singapore
cwngo@smu.edu.sg

ABSTRACT

Transformer achieves remarkable successes in understanding 1 and 2-dimensional signals (e.g., NLP and Image Content Understanding). As a potential alternative to convolutional neural networks, it shares merits of strong interpretability, high discriminative power on hyper-scale data, and flexibility in processing varying length inputs. However, its encoders naturally contain computational intensive operations such as pair-wise self-attention, incurring heavy computational burden when being applied on the complex 3-dimensional video signals.

This paper presents **Token Shift Module** (i.e., TokShift), a novel, zero-parameter, zero-FLOPs operator, for modeling temporal relations within each transformer encoder. Specifically, the TokShift barely temporally shifts partial [Class] token features back-and-forth across adjacent frames. Then, we densely plug the module into each encoder of a plain 2D vision transformer for learning 3D video representation. It is worth noticing that our TokShift transformer is a pure convolutional-free video transformer pilot with computational efficiency for video understanding. Experiments on standard benchmarks verify its robustness, effectiveness, and efficiency. Particularly, with input clips of 8/12 frames, the TokShift transformer achieves SOTA precision: 79.83%/80.40% on the Kinetics-400, 66.56% on EGTEA-Gaze+, and 96.80% on UCF-101 datasets, comparable or better than existing SOTA convolutional counterparts. Our code is open-sourced in: <https://github.com/VideoNetworks/TokShift-Transformer>.

CCS CONCEPTS

• **Computing methodologies** → **Activity recognition and understanding**.

KEYWORDS

Video classification; Transformer; Shift; Self-attention

ACM Reference Format:

Hao Zhang, Yanbin Hao, and Chong-Wah Ngo. 2021. Token Shift Transformer for Video Classification. In *Proceedings of the 29th ACM International*

*Yanbin Hao is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475272>

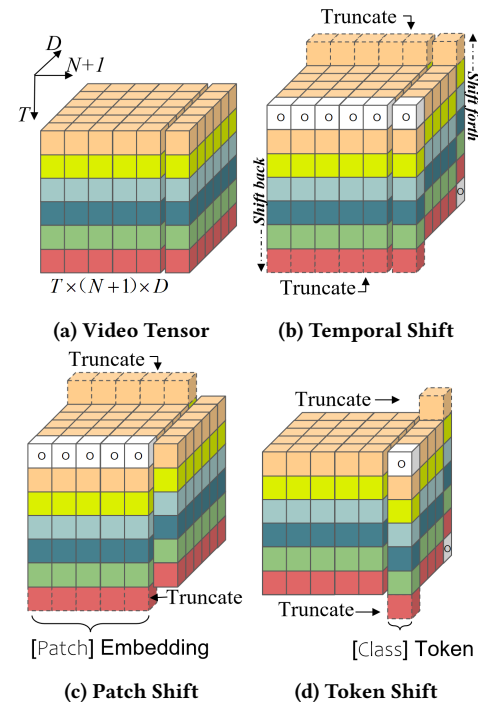


Figure 1: Types of Shift for video transformer. A video embedding contains two types of “words”: N [Patch] + 1 [Class] Token. Hence, the Shift can be applied on: (a) Neither, (b) Both, (c) Patch, and (d) [Class] Token along the temporal axis. “○” indicates padding zeros.

Conference on Multimedia (MM '21), October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3474085.3475272>

1 INTRODUCTION

“If a picture is worth a thousand words, is a video worth a million?” [40]. This quote basically predicts that image and video can be potentially interpreted as linguistic sentences, except that videos contain richer information than images. The recent progress of extending linguistic-style, convolutional-free transformers [9] on visual content understanding successfully verifies the quote’s prior half, whereas the latter half for video remains an open hypothesis. This paper studies the efficient and effective way of applying similar transformers for video understanding and answer the quote’s prophecy for videos.

As a potential alternative for convolutional neural networks, the transformer achieves remarkable progress in both NLP [1, 2, 8, 17, 31] and vision tasks [3, 9, 28, 32, 35]. Particularly, language and vision transformers show several merits over their CNN counterparts, such as good interpretability (e.g., attention to highlight core parts),

unsaturated discriminatory power scalable with hyper-scale data (e.g., million-scale language corpus[8, 31] or billion-scale images[9]), and flexibility in processing varying length inputs [1, 17].

However, unlike signal understanding on 1 and 2-dimensional data (i.e., NLP & static vision), the application of transformers on 3-dimensional video signals is challenging. Specifically, each encoder of a transformer naturally contains heavy computations such as pair-wise self-attention; meanwhile, a video has a longer sequential representation ($\mathbf{z}_v \in \mathbb{R}^{T \times (N+1) \times D}$) than an image ($\mathbf{z}_i \in \mathbb{R}^{(N+1) \times D}$) due to an extra temporal axis. Consequently, directly applying general transformers on flattened spatio-temporal video sequences will introduce an exponential explosion of computations (e.g., $T^2 \cdot \frac{(N+1)^2}{2}$ pair-wise distance calculations¹) in the training and inference phases.

To tackle this, we propose **Token Shift Module** (i.e., TokShift), a novel zero-parameter, zero-FLOPs operator, for modeling temporal relations within each video encoder. Specifically, the TokShift module barely temporally shifts partial [Class] token features back-and-forth (Figure 1d) across frames. Then, we densely plug the TokShift module into each encoder of a plain 2D vision transformer for learning 3D video representation.

Our TokShift is partially inspired by the success of spatial [6, 37, 41] and temporal [10, 20] shift operators for efficiency optimization on CNNs, but bears its own uniqueness. Specifically, on CNNs, spatial/temporal shift is uniformly applied across all spatial receptive fields on a feature-map. Follow this imitation, a copycat of the Temporal Shift Module for transformers is shown in Figure 1(b), where the shift is imposed on all patches and [Class] token features. However, we experimentally verify that, rather than temporally shifting features of all reception fields, just shift [Class] token feature is sufficient. As in Figure 1(b-1d), our TokShift introduces minimum modifications on original feature \mathbf{z}_v , and gets the most improvement among all shift variants. More importantly, equipped with the TokShift, a video transformer only needs to handle $T \cdot \frac{(N+1)^2}{2}$ pair-wise calculations. Finally, we verify that the shift operator is generalizable on transformers as it does on CNNs.

We try to intuitively explain the efficacy of the TokShift by drawing an analogy between *textual* description and *visual* feature vector of each frame. Specifically, the text describing a frame usually follow a similar linguistic form: “verb” + “noun” (e.g., “walking the dog” in Figure 4)), where “verb” is dynamically correlated with neighboring frames and “noun” are static. Correspondingly, the TokShift exchanges partial visual features of a frame (in the form of a token vector) temporally back-and-forth for motion capturing while keeping the rest for static semantics modeling.

We conduct extensive experiments of the TokShift-xfmr on standard benchmarks, such as Kinetics-400 [4], EGTEA-Gaze+ [19] and UCF-101 [26] datasets. Experiments demonstrate that with the TokShift, video transformers could perform comparable or better than the best 3D-CNNs, reaching SOTA performance. Our contributions are summarized below.

- **Transformer for video classification.** Our TokShift-xfmr is an efficient pure convolutional-free transformer pilot applied for video classification. We repeat the successes of

transformers on 3D videos as on 1/2D texts/images, promoting applications of transformers on different domains.

- **Efficient TokShift for transformer encoder.** We specifically design a novel zero-parameter, zero-FLOPs TokShift operator for the encoder. By merely manipulating global frame representation (i.e., [Class] token vector) via shift operations, we can reduce computations of pair-wise attention in dealing with sequential flattening video representations within transformers.
- **State-of-the-art performances.** We test the TokShift-xfmr on several standard benchmarks. Experiments show that the TokShift-xfmr achieves SOTA top-1 accuracy 79.83%/80.40% (8/12f) on Kinetics-400, 66.56% (8f) on EGTEA-Gaze+, and 96.80% (8f) on UCF-101 datasets, comparable or better than the best 3D-CNNs available (TSM, SlowFast-R101-NL, X3D-XXL).

2 RELATED WORKS

The TokShift-xfmr is closely relevant to research areas, including 3D-CNNs for video classification, shift strategies for convolutional optimization, and 1/2D transformers for NLP/vision understanding; we separately elaborate the related literature for each area as below.

3D-CNNs for video classifications. Convolutional neural networks stepped into a prosperous age in the past few years [5, 16, 18, 24, 25, 27, 36, 38, 42–45]. Consequently, the 3D-CNN [4, 11, 12, 14, 15, 21–23, 30, 39] has become a de-facto standard for video content understanding. Specifically, C3D [30] firstly inflates the convolutional kernel from 2D to 3D, to facilitate temporal modeling. For motion capturing, the two-stream network [23] fuses two CNNs trained on optical-flow with RGB modalities. Furthermore, I3D [4] deepens 3D-CNNs by inflating 3D kernels on inception-net, accompanied with a large-scale Kinetics dataset. To reduce computational overhead, Pseudo-3D [21] decomposes a 3D Spatio-temporal convolution into Spatial-2D + Temporal-1D form. This decomposition greatly balances computations, parameters, and performances, hence gain its popularity in later 3D-CNNs. S3D-G [39] scales channel elements with attention yielded by global feature. Non-local networks [33] introduces self-attention on top of CNNs; furthermore, CBA-QSA CNN [15] extends self-attention with compact bilinear mapping for fine-grained action classification. SlowFast [12] and X3D [11] are currently best two 3D-CNNs. The former designs dual-paths CNNs, where each path receives input clip of slow/fast sampling rate; the latter presents an efficient strategy to search for optimal hyperparameters (e.g., spatial/temporal resolutions, channels, depth, etc.) on a template network (i.e., X2D).

Spatial/temporal shift for efficient CNNs. Mobile computing demands more efficiency than cloud computing. Thereby, the Shift [6, 10, 20, 37, 41], a zero-parameter, zero-FLOPs operator for local feature aggregation, is proposed to reduce the complexities of CNNs. Specifically, Shift + 1D/2D convolution can lossily approximate a 2D/3D convolution. For example, in [37], Wu replaces spatial convolution of 3×3 kernel with Shift + 1×1 kernel for vision understanding. To further improve efficiency, ShiftAddNet [41] introduces Shift + Add, eliminating all multiplications. For videos, Lin [20] constructs TSM with Temporal Shift + spatial convolutions;

¹ $T, N, 1, D$ denote time-stamps, spatial-patches, [Class] token and feature-dim

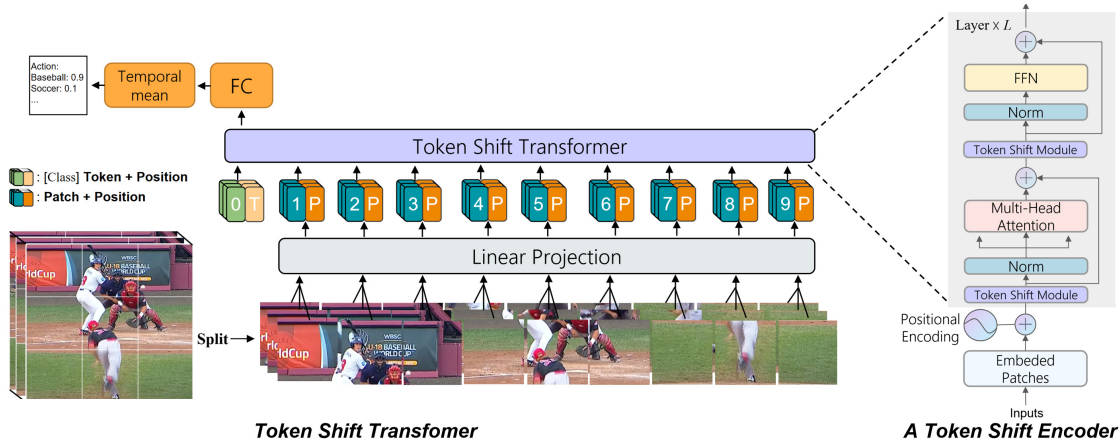


Figure 2: Token Shift Transformer: the input video is divided into multiple patch clips, according to the rigid spatial grids. Patch clips are linear-projected, added with positional embeddings (E_{pos}), and then concatenated with a [Class] token tensor (c_0) to form video embedding (z_0). The z_0 is fed into L layers of identical TokShift encoders for learning video representation.

RubiksNet [10] further removes spatial convolutions and introduces a learnable spatio-temporal shift operator.

Language & vision transformer. Firstly proposed in [31], the transformer relies on residual attention & feedforward (FFN) in feature learning and outperforms CNNs to become the de-facto standard in NLP tasks [2, 8]. To tackle inputs of different/long length, Reformer [17] and Longformer [1] introduce optimizations on attention calculations. Recently, researchers repeat this success in vision tasks, including image classification [9, 28], object detection [3], and segmentations [35]. Specifically, Dosovitskiy [9] proposes the first pure convolutional-free vision transformer (i.e., ViT) for image classification. To tackle the generalization problem caused by insufficient data, they equip the pre-training phase with a hyper-scale internal dataset (JFT-300M). In [28], DeiT bypasses the necessity of pre-training transformer on hyper-scale data through distillation. They introduce the token distillation for transformers to teach “student” transformer from various “teachers” (i.e., convnet or deeper transformer). For object detection and segmentation, as the transformer contains positional encodings under the encoder/decoder mechanism, it can be naturally extended to regress coordinates [3] and masks [35]. Overall, transformers are strong alternatives to CNNs and share merits like good interpretability, unsaturated discriminative power and flexibility on input length.

3 METHOD

An overview of the Token Shift transformer (i.e., TokShift-xfmr) is presented in Figure (2). Our TokShift-xfmr follows a general pipeline of transformers, except for working on video inputs and containing extra temporal processing modules (i.e., the TokShift).

To fit with transformer, a video $\mathbf{v} \in \mathbb{R}^{T \times H \times W \times 3}$ is firstly reshaped into sequential tensor $\hat{\mathbf{v}} \in \mathbb{R}^{T \times N \times d}$, where $T, H/W, P$ separately represent clip-length, spatial-resolutions and patch-size, and $N = \frac{H \times W}{P^2}$, $d = P^2 \times 3$ denote the number of patches and the number of RGB pixels in a patch respectively. Then, each i -th patch \mathbf{x}_0^i in $\hat{\mathbf{v}}$ is linearly projected into embedding space by \mathbf{E} and further added with the corresponding spatial positional embeddings \mathbf{E}_{pos} . Similar as BERT [8] and ViT [9], an extra [Class] token tensor $\mathbf{c}_0 \in \mathbb{R}^{T \times D}$

is concatenated with patches’ embedding to represent per-frame global contents. Finally, the video $\hat{\mathbf{v}}$ is embedded as \mathbf{z}_0 (Equation (1) and Figure (1a)).

$$\begin{aligned} \mathbf{z}_0 &= [\mathbf{c}_0; \mathbf{x}_0^1 \mathbf{E}; \mathbf{x}_0^2 \mathbf{E}; \dots; \mathbf{x}_0^i \mathbf{E}; \dots; \mathbf{x}_0^N \mathbf{E}] + \mathbf{E}_{pos} \quad (1) \\ \mathbf{x}_0^i &\in \mathbb{R}^{T \times d}, \quad i = 1, 2, \dots, N \\ \mathbf{E} &\in \mathbb{R}^{d \times D}, \quad \mathbf{E}_{pos} \in \mathbb{R}^{(N+1) \times D} \\ \mathbf{z}_0 &\in \mathbb{R}^{T \times (N+1) \times D} \end{aligned}$$

The TokShift-xfmr (Figure (2)) contains L replicated, identical encoders. An encoder consists of the TokShift module, Layer-Norm (LN), Multi-Head Self-attention (MSA), and Feed-Forward Network (FFN). The workflow for connecting them is shown by Figure (2) or Equations (2-4 & 8-10). We will describe each module separately as below.

The TokShift module (Equation (2) & (8)) serves to in-placed manipulate global token $\mathbf{c}/\hat{\mathbf{c}}_{l-1}$ along temporal axis, introducing temporal interactions before feeding video embedding $\mathbf{z}/\hat{\mathbf{z}}_{l-1}$ through MSA/FFN. Hereby, the token $\mathbf{c}/\hat{\mathbf{c}}_{l-1}$ is the special [Class] word of $\mathbf{z}/\hat{\mathbf{z}}_{l-1}$. We will present details of the TokShift and its variants in section 3.1.

$$\mathbf{c}_{l-1} = TokShift(\mathbf{c}_{l-1}) \quad (2)$$

$$\mathbf{z}_{l-1} = LN(\mathbf{z}_{l-1}) \quad (3)$$

$$\hat{\mathbf{z}}_{l-1} = MSA(\mathbf{z}_{l-1}) + \mathbf{z}_{l-1}, \quad (4)$$

$$\mathbf{z}/\hat{\mathbf{z}}_{l-1} \in \mathbb{R}^{T \times (N+1) \times D}$$

$$\mathbf{c}/\hat{\mathbf{c}}_{l-1} \in \mathbb{R}^{T \times D}$$

The MSA (same for LN & FFN) acts in the same way as in 2D ViT [9]. Specifically, its functions independently process each frame. Pick $\mathbf{z}_{l-1}^{(t)} \in \mathbb{R}^{(N+1) \times D}$ at time-stamp t of \mathbf{z}_{l-1} for example, functions of MSA are performed within each frame (Equations (5)-(7)).

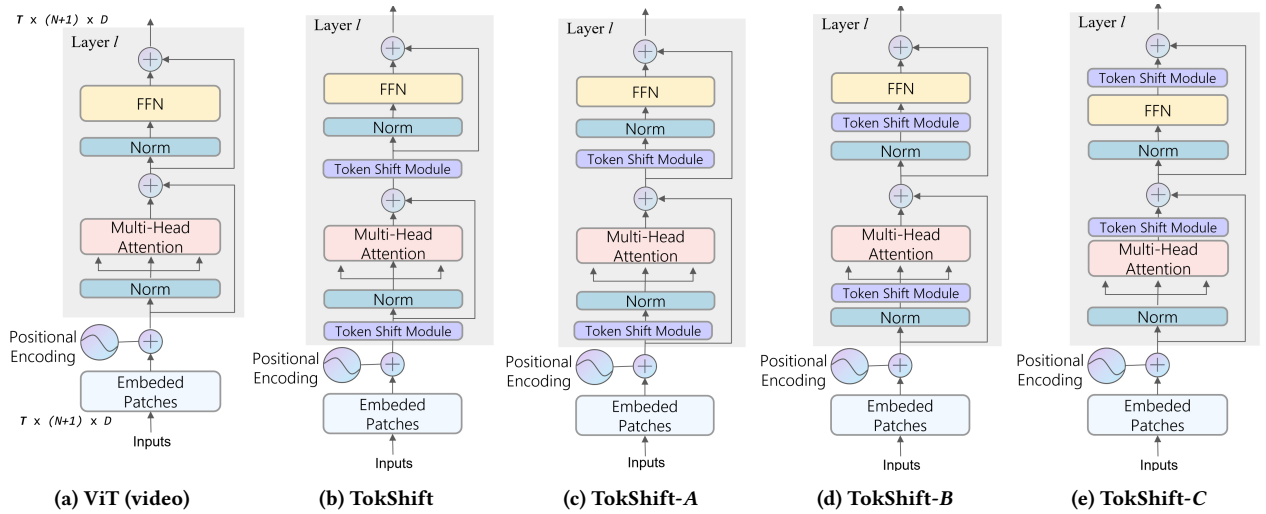


Figure 3: Integrations of TokShift module into ViT Encoder: the TokShift can be plugged in multiple positions of a vision encoder: “prior residual” (TokShift), “prior layer-norm” (TokShift-A), “prior MSA/FFN” (TokShift-B), and “post MSA/FFN” (TokShift-C). Notably, a fronter position will affect the subsequent modules inside residual blocks more.

$$MSA(\mathbf{z}_{l-1}^{(t)}) = \text{Concat} \left[\text{head}_1^{(t)}, \text{head}_2^{(t)}, \dots, \text{head}_M^{(t)} \right] \quad (5)$$

$$\text{head}_i^{(t)} = \text{Softmax} \left(\frac{\mathbf{Q}^{(t)} \mathbf{K}^{(t)}}{\sqrt{D}} \right) \mathbf{V}^{(t)} \quad (6)$$

$$\mathbf{Q}^{(t)}, \mathbf{K}^{(t)}, \mathbf{V}^{(t)} = \mathbf{z}_{l-1}^{(t)} \times \mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v \quad (7)$$

$$\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v \in \mathbb{R}^{D \times D}$$

The FFN contains two linear projections and a GELU activation, serving to project each frame’s feature independently.

$$\hat{\mathbf{c}}_{l-1} = \text{TokShift}(\hat{\mathbf{c}}_{l-1}) \quad (8)$$

$$\hat{\mathbf{z}}_{l-1} = \text{LN}(\hat{\mathbf{c}}_{l-1}) \quad (9)$$

$$\mathbf{z}_l = \text{FFN}(\hat{\mathbf{z}}_{l-1}) + \hat{\mathbf{z}}_{l-1}, \quad l = 1, 2, \dots, L \quad (10)$$

Finally, a classification layer works on outputs (i.e., \mathbf{c}_L) of the last encoder. Specifically, the video label is obtained by averaging frame-level predictions (Equation (11)), where FC represents fully-connected layer of shape “ $D \times \text{Categories}$ ”.

$$\mathbf{y} = \frac{1}{T} \sum_{i=1}^T FC(\mathbf{c}_L), \quad \mathbf{c}_L \in \mathbb{R}^{T \times D} \quad (11)$$

3.1 Shift Variants and the TokShift

We propose and compare several shift variants, including the TokShift module, customized for the video transformer paradigm.

Token Shift merely manipulates [Class] token tensors within each encoder (Figure (1d)). Our motivation follows a principle that the current frame’s partial contents are exchanged with prior/post time-stamps for dynamic motion while the rest are kept for static semantics. Coincidentally, the extra global [Class] token tensor, aggregated from local features via weighted-sum (i.e., attention), is appropriate for implementing the principle.

Given a token tensor $\mathbf{c}_l \in \mathbb{R}^{T \times D}$ representing the global sequence of a clip, the TokShift works as follows: \mathbf{c}_l is firstly split into 3 groups along channel dimension (Equation (12)).

$$\mathbf{c}_l = [\mathbf{s}_a, \mathbf{s}_b, \mathbf{s}_c] \quad (12)$$

$$\mathbf{s}_a, \mathbf{s}_b, \mathbf{s}_c \in \mathbb{R}^{T \times a}, \mathbb{R}^{T \times b}, \mathbb{R}^{T \times c}$$

$$a + b + c = D$$

Then, channels of splits $\mathbf{s}_a/\mathbf{s}_c$ are temporally shifted with prior/post time-stamps. For the split \mathbf{s}_b , its content remains unchanged (Equations (13)-(15)). The percentages of shifted channels are determined by $\frac{a}{D}$ & $\frac{c}{D}$.

$$\mathbf{s}_a(t) = \mathbf{s}_a(t-1) \quad (13)$$

$$\mathbf{s}_b(t) = \mathbf{s}_b(t) \quad (14)$$

$$\mathbf{s}_c(t) = \mathbf{s}_c(t+1) \quad (15)$$

$$t = 1, 2, \dots, T$$

Non Shift module can be obtained by replacing “TokShift()” with “Identity()” function in Equation (2) & (8). This modification removes all temporal interactions and processes each frame independently. We name this simple 2D vision transformer for videos as ViT (video) for simplicity (i.e., Figure (1a) & (3a)).

Temporal Shift extension for the encoder is obtained by imitating spatial/temporal shifts [20, 37, 41] in CNN. Recall that a CNN feature-map of image/video is a tensor collecting features from all the receptive fields (spatial or spatio-temporal locations). A shift operator then exchanges partial channels from neighboring receptive fields for each location on the feature-map. Similarly, in Figure (1b), temporal shift for encoder exchanges partial channels from prior/post time-stamps for all “ N patches” and “1 [Class] token”.

Patch Shift is a variant of the temporal shift. Since the [Class] token is a global representation, we exempt it from shift operation and only process patch embeddings (Figure (1c)).

We compare the four shift variants on a standard benchmark and observe that the TokShift performs the best among them, showing the potential of capitalizing [*Class*] tokens for temporal modeling in video transformers. We claim that TokShift introduces the minimum modifications on \mathbf{z}_l but grants the maximum benefits.

3.2 Integration into ViT Encoder

There are multiple candidate positions in an encoder to implant the TokShift module. The position is essential since it determines the degree of motions involved in representation learning. Specifically, from TokShift to TokShift-*A/B/C* (Figure (3b)-(3e)), the influence of motions decreases as fewer modules work upon shifted features.

A general encoder contains two residual blocks: one for self-attention, another for feature embedding (Figure (3a)). In each block, the TokShift can be placed at “prior residual”, “prior layer-norm”, “prior MSA/FFN” and “post MSA/FFN”. We experimentally verify that placing TokShift at a front position (i.e., “prior residual”) brings the best effect. This finding is different from TSM on CNN, where the optimal position is “post residual”.

4 EXPERIMENTS

We conduct extensive experiments on three standard benchmarks for video classification and adopt Top-1/5 accuracy (%) as evaluation metrics. We also report the model parameters and GFLOPs to quantify computations.

4.1 Datasets

We adopt one large-scale (Kinetics-400) and two small-scale (EGTEA-Gaze+ & UCF-101) datasets as evaluation benchmarks. Below presents their brief descriptions.

Kinetics-400 [4] serves as a standard large-scale benchmark for video classification. Its clips are truncated into 10 seconds duration with humans’ annotations. The training/validation set separately contains $\sim 246\text{k}/20\text{k}$ clips covering 400 video categories.

EGTEA Gaze+ [19] is a First-Person-View video dataset covering 106 fine-grained daily action categories. We select split-1 to evaluate our model, and this split contains 8,299/2,022 training/validating clips, with an average clip length of 3.1 seconds.

UCF-101 [26] contains 101 action categories, such as “human-human/object” interactions, sports and etc. We also select split-1 for performance reporting. This split contains 9,537/3,783 training/validating clips with average duration of 5.8 seconds.

4.2 Implementations

We implement a paradigm of shift-transformers, including the ViT (video), TemporalShift/PatchShift/TokShift-xfmr, on top of 2D ViT [9] with PyTorch. We also support backbones of various types or depths.

All experiments share the same settings unless particularly specified. For ablations with higher resolutions, more frames, or heavier backbones, we follow the “linear scaling rule” [13] to adjust lr according to batch-size.

Training. Each clip, randomly cropped into 224×224 (256/384 for high resolutions), contains 8/16 frames with a temporal step of 32. We apply training augmentations before cropping to reduce overfitting, including random resize (with the short side in [244,

330] and maintain aspect ratio), random brightness, saturation, gamma, hue, and horizontal-flip. We train the model for 18 epochs with batch-size 21 per GPU. Base lr is set to 0.1, is decayed by 0.1 at epoch [10, 15] during training. The *Base-16* ViT [9] (contains 12 encoders) is adopted as the backbone and the shifted proportion is set to $\frac{a}{D}, \frac{c}{D} = \frac{1}{4}$. All experiments are run on $2/8 \times V100$ GPUs.

Inference adopts the same testing strategies as [12]. Specifically, we uniformly sample 10 clips from a testing video. Each clip are resized to short-side=224 (256/384), then cropped into three 224×224 sub-clips (from “left”, “center”, “right” parts). A video prediction is obtained by averaging scores of 30 sub-clips.

4.3 Ablation Study

Our ablation studies the impacts of the shift-types, integrations, hyperparameters of video transformers on the Kinetics-400 dataset.

Non Shift vs TokShift. We compare the Non Shift and TokShift under various sampling steps. Specifically, we fix clip size (T) to 8 frames, but change sampling steps (S) to cover variable temporal duration.

Acc1 Xfmr	$T \times S$			
	8×8	8×16	8×32	8×64
ViT (video)[9]	76.17	76.32	76.02	75.73
TokShift	76.90 (0.73 \uparrow)	76.81 (0.49 \uparrow)	77.28 (1.26 \uparrow)	76.60 (0.87 \uparrow)

Table 1: NonShift vs TokShift under different sampling strategies (“T/S” refers to frames/sampling-step; Accuracy-1).

As in Tab. 1, the TokShift-xfmr outperforms ViT (video) under all settings. We observe that the performance of TokShift varies with different temporal steps. A smaller step makes TokShift not effective in learning temporal evolution. On the other hand, a large step of 64 frames results in a highly discontinuous frame sequence, affecting learning effectiveness. A temporal step of 32 frames appears to be a suitable choice.

Integration. We study the impacts of implanting the TokShift at different positions of an encoder. As mentioned in section 3.2, a frontier position will affect the later modules inside residual blocks more.

Model	Words Shifted	GFLOPs \times Views	Acc1 (%)	Acc5 (%)
ViT (video) [9]	None	134.7×30	76.02	92.52
TokShift-A	Token	134.7×30	76.85	93.10
TokShift-B	Token	134.7×30	77.21	92.81
TokShift-C	Token	134.7×30	77.00	92.92
TokShift	Token	134.7×30	77.28	92.91

Table 2: Comparisons of the TokShift residing in different positions (GFLOPs, Accuracy-1/5).

Table 2 lists comparisons of the TokShift residing in different positions. We observe that: (1). All models share the same GFLOPs, which is consistent with our zero-FLOPs claims on the TokShift; (2) Regardless of positions, the TokShift stably improves over the ViT (video) baseline; (3). Placing the TokShift at “prior residual”, a front position that affects all later modules, presents the best top-1 accuracy.

Words Shifted. We assess contributions of the token-specific shift by comparing all shift variants in section 3.1.

Shift Type	Words Shifted	$GFLOPs \times Views$	Acc1 (%)	Acc5 (%)
ViT (video) [9]	None	134.7×30	76.02	92.52
TemporalShift	Token + Patches	134.7×30	72.88	91.24
PachShift	Patches	134.7×30	73.08	91.17
TokShift	Token	134.7×30	77.28	92.91

Table 3: Comparisons of shifting different visual “Words” (GFLOPs, Accuracy-1/5).

In Table 3, the TokShift exhibits the best top-1 accuracy. Additionally, consideration of patches in shift (TemporalShift & PatchShift) reduces overall performance. The reason lies in that patches of the same spatial grid across times incurs visual misalignment when rigidly dividing a moving object. We verify this hypothesis by measuring the mean cosine similarity between the patch and its temporal neighbor features on kinetics-400 val set with Temporal/Patch Shift transformer and observe distances scores of 0.577/0.570. As for [Class] token part, since it reflects global frame-level contents without alignment concerns, the distance between neighboring temporal tokens is 0.95. Consequently, for the TokShift or CNN-Shift, the drawback is bypassed by global alignment or sliding window.

Proportion of shifted channels. This hyperparameter controls the ratio of dynamic/static information in a video feature. We evaluate $\frac{a}{D} \cdot \frac{c}{D}$ in range $[1/4, 1/8, 1/16]$ and present results in Table 4. We experimentally find that 1/4 is an optimal value, indicating that half channels are shifted (1/4 back + 1/4 forth) while the rest half remain unchanged.

Model	Channels Shifted ($\frac{a}{D} + \frac{c}{D}$)	Acc1 (%)	Acc5 (%)
TokShift	1/4 + 1/4	77.28	92.91
TokShift	1/8 + 1/8	77.18	92.95
TokShift	1/16 + 1/16	76.93	92.82

Table 4: Comparisons of the TokShift regarding to proportions of shifted channels (Accuracy-1/5).

Word count is an essential factor in transformer learning since more words indicate more details. As for videos, the number of visual words is positively correlated with two factors: spatial resolutions and temporal frames. We study the impacts of word count on the TokShift-xfmr. Table 5 lists the detailed performances under different word counts. Here, MR/HR represents middle/high resolutions.

Model	Res $H \times W$	Words/Frame $N + 1$	#Frames T	# Words $T \cdot (N + 1)$	Acc1 (%)
TokShift	224×224	$14^2 + 1$	6	1,182	76.72
TokShift	224×224	$14^2 + 1$	8	1,576	77.28
TokShift	224×224	$14^2 + 1$	10	19,70	77.56
TokShift	224×224	$14^2 + 1$	16	3,152	78.18
TokShift (MR)	256×256	$16^2 + 1$	8	2,056	77.68
TokShift (HR)	384×384	$24^2 + 1$	8	4,616	78.14

Table 5: Comparisons of the TokShift according to visual word counts (Accuracy-1).

We observe that: (1). Increasing word counts by either introducing more frames or spatial grids will lead to improvement (e.g., 76.72→78.18 or 77.68→78.14); (2). Increasing temporal resolution is more economic than spatial resolution. The former performs

Model	Backbone	Res $(H \times W)$	# Words $T \cdot (N + 1)$	Acc1 (%)
TokShift (HR)	Base-16	384×384	$8 \cdot (24^2 + 1)$	78.14
TokShift-Large (HR)	Large-16	384×384	$8 \cdot (24^2 + 1)$	79.83
TokShift-Large (HR)	Large-16	384×384	$12 \cdot (24^2 + 1)$	80.40
TokShift (MR)	Base-16	256×256	$8 \cdot (16^2 + 1)$	77.68
TokShift-Hybrid (MR)	R50+Base-16	256×256	$8 \cdot (16^2 + 1)$	77.55
TokShift-Hybrid (MR)	R50+Base-16	256×256	$16 \cdot (16^2 + 1)$	78.34

Table 6: Comparisons of the TokShift on different transformer backbones (Accuracy-1).

comparable (78.18 vs 78.14) with fewer words than the latter (3,152 vs 4,616).

On various backbones. We test the TokShift on transformers of various depths and types. Specifically, we further evaluate the TokShift-xfmr on *Large-16* (24 encoders) and *Hybrid-16* (12 encoders) ViTs. Compared to the Base-16, the Large-16 doubles the number of encoders. Hybrid-16 shares the same number of encoders but replaces the linear projection layer (\mathbf{E} in Equation (1)) with three residual blocks of ResNet50. Notably, for the TokShift-Large (HR), we set $T = 12$ due to the memory limit.

Table 6 lists their comparisons with base backbones. The TokShift-xfmr performs better with deeper layers (Base-16 vs Large-16). Whereas hybrid model achieves comparable performance as convolution free transformers, indicating that transformer is independent of convolutions. Additionally, performances further increase when using more frames (79.83→80.40; 77.55→78.34)

4.4 Comparison with the State-of-the-Art

We compare our TokShift-xfmrs with current SOTAs on Kinetics-400 datasets in Table 7. Since there is no prior work that purely utilizes transformer, we mainly select 3D-CNNs, such as I3D, Non-Local, TSM, SlowFast, X3D as SOTAs.

Compared to 3D-CNNs, transformers exhibit prominent performances. Particularly, the spatial-only transformer, i.e., ViT (video), achieves comparable or better performance (76.02) than strong 3D-CNNs, such as I3D (71.1), S3D-G (74.7) and TSM (76.3). Moreover, a comparison of our most slimed TokShift-xfmr (77.28) with TSM (76.3) indicates that shift gains extra benefits on transformers than CNNs. Thirdly, the TokShift-xfmr (MR) achieves comparable performance (77.68) with Non-Local R101 networks (77.7) under the same spatial resolutions (255) but using fewer frames (8 vs 128). This is reasonable since both contain the attentions. Finally, we compare TokShift-xfmr with the current best 3D-CNNs: SlowFast and X3D. The two 3D-CNNs are particularly optimized by a dual-paths or efficient network parameters searching mechanism. Since both SlowFast & X3D introduce large structural changes over corresponding 2D nets, their authors prefer to train on videos with sufficient epochs (256), rather than initialize from ImageNet pre-trained weights. We show that our TokShift-Large-xfmr (80.40) performs better than SlowFast (79.80) and is comparable with X3D-XXL (80.4), where models are under their best settings. Notably, we only use 12 frames, whereas SlowFast/X3D-XXL uses 32/16+ frames.

In all, a pure transformer can perform comparable or better than 3D-CNNs for video classification. We additionally verify that temporal modeling can be imposed just on global [Class] tokens rather than whole video embeddings. Besides, the common computational limitation of the transformer (appeared in NLP/vision task) also

Model	Backbone	Pretrain	Inference Res ($H \times W$)	# Frames/Clip T	GFLOPs×Views	Params	Accuracy-1 (%)	Accuracy-5 (%)
I3D [4] from [11]	InceptionV1	ImageNet	224 × 224	250	108× NA	12M	71.1	90.3
Two-Stream I3D [4] from [11]	InceptionV1	ImageNet	224 × 224	500	216× NA	25M	75.7	92.0
S3D-G [39]	InceptionV1	ImageNet	224 × 224	250	71.3×NA	11.5M	74.7	93.4
Two-Stream S3D-G [39]	InceptionV1	ImageNet	224 × 224	500	142.6×NA	11.5M	77.2	93.0
Non-Local R50 [33] from [11]	ResNet50	ImageNet	256 × 256	128	282 × 30	35.3M	76.5	92.6
Non-Local R101[33] from [11]	ResNet101	ImageNet	256 × 256	128	359 × 30	54.3M	77.7	93.3
TSM [20]	ResNet50	ImageNet	256 × 256	8	33 × 10	24.3M	74.1	91.2
TSM [20]	ResNet50	ImageNet	256 × 256	16	65 × 10	24.3M	74.7	-
TSM [20]	ResNext101	ImageNet	256 × 256	8	NA×10	-	76.3	-
SlowFast 4 × 16 [12]	ResNet50	None	256 × 256	32	36.1 × 30	34.4M	75.6	92.1
SlowFast 8 × 8 [12]	ResNet50	None	256 × 256	32	65.7 × 30	-	77.0	92.6
SlowFast 8 × 8 [12]	ResNet101	None	256 × 256	32	106 × 30	53.7M	77.9	93.2
SlowFast 8 × 8 [12]	ResNet101+NL	None	256 × 256	32	116 × 30	59.9M	78.7	93.5
SlowFast 16 × 8 [12]	ResNet101+NL	None	256 × 256	32	234 × 30	59.9M	79.8	93.9
X3D-M [11]	X2D [11]	None	256 × 256	16	6.2 × 30	3.8M	76.0	92.3
X3D-L [11]	X2D [11]	None	356 × 356	16	24.8 × 30	6.1M	77.5	92.9
X3D-XL[11]	X2D [11]	None	356 × 356	16	48.4 × 30	11M	79.1	93.9
X3D-XXL [11]	X2D [11]	None	NA	NA	194.1 × 30	20.3M	80.4	94.6
ViT (Video)[9]	Base-16	ImageNet-21k	224 × 224	8	134.7 × 30	85.9M	76.02	92.52
TokShift	Base-16	ImageNet-21k	224 × 224	8	134.7 × 30	85.9M	77.28	92.91
TokShift (MR)	Base-16	ImageNet-21k	256 × 256	8	175.8 × 30	85.9M	77.68	93.55
TokShift (HR)	Base-16	ImageNet-21k	384 × 384	8	394.7 × 30	85.9M	78.14	93.91
TokShift	Base-16	ImageNet-21k	224 × 224	16	269.5 × 30	85.9M	78.18	93.78
TokShift-Large (HR)	Large-16	ImageNet-21k	384 × 384	8	1397.6 × 30	303.4M	79.83	94.39
TokShift-Large (HR)	Large-16	ImageNet-21k	384 × 384	12	2096.4 × 30	303.4M	80.40	94.45

Table 7: Comparison to state-of-the-arts on Kinetics-400 Val.

exists for videos. Though the TokShift introduces zero parameters & FLOPs, the 2D transformer alone consumes more computations than 3D-CNNs. The reasons lie in: (1). Pair-wise distance calculation in attention is computationally expensive; (2). 3D-CNNs have been computationally optimized (TSM, P3D, etc.). Nevertheless, computation optimization for the transformer is already very hot for NLP [1, 7] and image understanding [29]. As expected, computation optimization for long-length videos also requires future exploration.

4.5 Attention Visualization

An advantage of transformer over CNN is interpretability. Specifically, the self-attention emphasizes vital visual clues with high attentions. Hence, we can visualize the attention map to infer what the transformer values.

We adopt a public script² to visualize the attention map from the last encoder. Maps from multiple heads are averaged to a heat map. We sample six clips, including daily and sports activities, from the Kinetics-400 and show their maps in Figure (4). We keep half (4 of 8) frames for space-saving. Notably, most action categories follow a form of “verb” + “noun” (e.g., “flying kites”), meeting our intuitive explanations. Our TokShift-xfmr learns to value core parts such as “kites” in “flying kites”, “boy + dog” in “walking a dog” and etc. Moreover, for “flying kites”, the transformer even highlights all kites in the sky, indicating its potentials in counting applications.

4.6 Fine-tune on small-scale datasets

We study the impacts of pre-trained weights on small-scale datasets. Specifically, we fine-tune models on EGTEA Gaze+ and UCF-101 datasets with pre-weights from “None”, “ImageNet-21k” and “Kinetics-400”. The sampling step is reduced to 8, as both small datasets have a shorter duration (3/5s) than Kinetics-400 (10s). Training schemes

²https://github.com/jeonsworld/ViT-pytorch/blob/main/visualize_attention_map.ipynb

Model	Modality	Pretrain	Res ($H \times W$)	# Frames T	EGAZ+ Acc1 (%)
TSM [20]	RGB	Kinetics-400	224 × 224	8	63.45
SAP [34]	RGB	Kinetics-400	256 × 256	64	64.10
ViT (Video) [9]	RGB	ImageNet-21k	224 × 224	8	62.59
TokShift	RGB	None	224 × 224	8	28.90
TokShift	RGB	ImageNet-21k	224 × 224	8	62.85
TokShift*	RGB	ImageNet-21k	224 × 224	8	59.24
TokShift	RGB	Kinetics-400	224 × 224	8	63.69
TokShift*	RGB	Kinetics-400	224 × 224	8	64.82
TokShift	OptFlow	Kinetics-400	224 × 224	8	48.81
TokShift*-En	OptFlow+RGB	Kinetics-400	224 × 224	8	65.08
TokShift* (HR)	RGB	Kinetics-400	384 × 384	8	65.77
TokShift-Large* (HR)	RGB	Kinetics-400	384 × 384	8	66.56

Table 8: Impacts of various pretrained weights on EGTEA-GAZE++ Split-1 dataset (“*” means freeze layer-norm).

are also optimized (EGTEA: 18 epochs, decay at [10,15]) and UCF: 25 epochs, decay at [10, 20]).

Table (8) & (9) list their performances. Firstly, pre-training with Kinetics-400 outperforms the rest (i.e., **64.82** vs 62.85 vs 28.90 on EGTEA Gaze+; **95.35** vs 91.65 vs 91.60 on UCF-101). We project features of 60 categories from EGTEA Gaze+ with large improvements by t-SNE and find that the larger an upstream dataset is, the more scattered a downstream features are (Figure (5)). Also, we study the impacts of freezing norm. Freezing “layer-norms” is beneficial for fine-tuning from Kinetics-400 weights, whereas harmful for ImageNet-21k weights. Because the Kinetics-400 shares the same modality (Video) with EGTEA Gaze+, while ImageNet-21k is different (Image). We verify the efficacy of two-stream TokShift-xfmr by adding the optical-flow (64.82→65.08). Finally, we compare the TokShift with strong 3D-CNNs (I3D, P3D & TSM), and verify its efficacy on small-scale datasets.

5 CONCLUSION

We propose a zero-parameter, zero-FLOPs TokShift operator for constructing a pure convolutional-free video transformer. Specifically, our TokShift-xfmr alleviates intensive pair-wise distance




Figure 4: Visualization of attention on sample clips from the Kinetics-400 dataset. The odd row presents original frames, and the even presents corresponding attention maps.

Model	Pretrain	Res ($H \times W$)	# Frames T	UCF101 Acc1 (%)
I3D [4]	Kinetics-400	224×224	250	84.50
P3D [21]	Kinetics-400	224×224	16	84.20
Two-Stream I3D [4]	Kinetics-400	224×224	500	93.40
TSM [20]	Kinetics-400	256×256	8	95.90
ViT (Video) [9]	ImageNet-21k	256×256	8	91.46
TokShift	None	256×256	8	91.60
TokShift	ImageNet-21k	256×256	8	91.65
TokShift*	Kinetics-400	256×256	8	95.35
TokShift* (HR)	Kinetics-400	384×384	8	96.14
TokShift-Large* (HR)	Kinetics-400	384×384	8	96.80

Table 9: Impacts of various pretrained weights on UCF-101 Split dataset (“*” means freeze layer-norm).

calculations of the Spatio-temporal attention, maintains the same complexity as a common 2D ViT, while achieving better or comparable performance (80.40%) as 3D-CNN SOTAs. More importantly, TokShift firstly conducts temporal modeling on the global token of video transformers. Since global [Class] token is aggregated by weighted-sum all spatial patches’ embeddings and reflects global visual content. Hence, modeling temporal interactions across frames can be done through the tokens. Besides, we still face an inherent




Figure 5: Visualizing feature distributions with different pretrained weights via t-SNE.

computational burden of attention, within the 2D ViT, especially in processing long-length video. We leave computational optimization of video transformer as future works. Finally, with the help of our TokShift-xfmr, we can partially answer the prophecy of the quote on video: “A 10 seconds video is worthy of 3,152 visual words”.

6 ACKNOWLEDGMENTS

I thank Dr. Jingjing Chen and Dr. Bin Zhu for inspiring discussion and encouragement.

REFERENCES

- [1] Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150* (2020).
- [2] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. (2020). arXiv:2005.14165 [cs.CL]
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*. Springer, 213–229.
- [4] Joao Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6299–6308.
- [5] Jingjing Chen and Chong-Wah Ngo. 2016. Deep-based ingredient recognition for cooking recipe retrieval. In *Proceedings of the 24th ACM international conference on Multimedia*. 32–41.
- [6] Weijie Chen, Di Xie, Yuan Zhang, and Shiliang Pu. 2019. All you need is a few shifts: Designing efficient convolutional neural networks for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7241–7250.
- [7] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. 2020. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794* (2020).
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [10] Linxi Fan, Shyamal Buch, Guanzhi Wang, Ryan Cao, Yuke Zhu, Juan Carlos Niebles, and Li Fei-Fei. 2020. RubiksNet: Learnable 3D-Shift for Efficient Video Action Recognition. In *European Conference on Computer Vision*. Springer, 505–521.
- [11] Christoph Feichtenhofer. 2020. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 203–213.
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slow-fast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6202–6211.
- [13] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677* (2017).
- [14] Yanbin Hao, Zi-Niu Liu, Hao Zhang, Bin Zhu, Jingjing Chen, Yu-Gang Jiang, and Chong-Wah Ngo. 2020. Person-level Action Recognition in Complex Events via TSD-TSM Networks. In *Proceedings of the 28th ACM International Conference on Multimedia*. 4699–4702.
- [15] Yanbin Hao, Hao Zhang, Chong-Wah Ngo, Qiang Liu, and Xiaojun Hu. 2020. Compact Bilinear Augmented Query Structured Attention for Sport Highlights Classification. In *Proceedings of the 28th ACM International Conference on Multimedia*. 628–636.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [17] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451* (2020).
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012), 1097–1105.
- [19] Yin Li, Miao Liu, Rehg, and James M. 2018. In the eye of beholder: Joint learning of gaze and actions in first person video. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 619–635.
- [20] Ji Lin, Chuang Gan, and Song Han. 2019. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7083–7093.
- [21] Zhaofan Qiu, Ting Yao, and Tao Mei. 2017. Learning spatio-temporal representation with pseudo-3d residual networks. In *proceedings of the IEEE International Conference on Computer Vision*. 5533–5541.
- [22] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. 2019. Learning spatio-temporal representation with local and global diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12056–12065.
- [23] Karen Simonyan and Andrew Zisserman. 2014. Two-stream convolutional networks for action recognition in videos. *arXiv preprint arXiv:1406.2199* (2014).
- [24] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [25] Xue Song, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. 2021. Spatial-temporal Graphs for Cross-modal Text2Video Retrieval. *IEEE Transactions on Multimedia* (2021).
- [26] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012).
- [27] Zixuan Su, Xindi Shang, Jingjing Chen, Yu-Gang Jiang, Zhiyong Qiu, and Tat-Seng Chua. 2020. Video Relation Detection via Multiple Hypothesis Association. In *Proceedings of the 28th ACM International Conference on Multimedia*. 3127–3135.
- [28] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. 2020. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877* (2020).
- [29] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. 2020. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877* (2020).
- [30] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*. 4489–4497.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017).
- [32] Junke Wang, Zuxuan Wu, Jingjing Chen, and Yu-Gang Jiang. 2021. M2TR: Multi-modal Multi-scale Transformers for Deepfake Detection. *arXiv preprint arXiv:2104.09770* (2021).
- [33] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. 2018. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7794–7803.
- [34] Xiaohan Wang, Yu Wu, Linchao Zhu, and Yi Yang. 2020. Symbiotic attention with privileged information for egocentric action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 12249–12256.
- [35] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. 2020. End-to-End Video Instance Segmentation with Transformers. *arXiv preprint arXiv:2011.14503* (2020).
- [36] Zhipeng Wei, Jingjing Chen, Xingxing Wei, Linxi Jiang, Tat-Seng Chua, Fengfeng Zhou, and Yu-Gang Jiang. 2020. Heuristic black-box adversarial attacks on video recognition models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 12338–12345.
- [37] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, and Kurt Keutzer. 2018. Shift: A zero flop, zero parameter alternative to spatial convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9127–9135.
- [38] Jiaxin Wu and Chong-Wah Ngo. 2020. Interpretable embedding for ad-hoc video search. In *Proceedings of the 28th ACM International Conference on Multimedia*. 3357–3366.
- [39] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 305–321.
- [40] Aman Yadav, Michael M Phillips, Mary A Lundeborg, Matthew J Koehler, Katherine Hilden, and Kathryn H Dirkin. 2011. If a picture is worth a thousand words is video worth a million? Differences in affective and cognitive processing of video and text cases. *Journal of Computing in Higher Education* 23, 1 (2011), 15–37.
- [41] Haoran You, Xiaohan Chen, Yongan Zhang, Chaojian Li, Sicheng Li, Zihao Liu, Zhiyong Wang, and Yingyan Lin. 2020. Shiftaddnet: A hardware-inspired deep network. *arXiv preprint arXiv:2010.12785* (2020).
- [42] Hao Zhang and Chong-Wah Ngo. 2016. Object pooling for multimedia event detection and evidence localization. *ITE Transactions on Media Technology and Applications* 4, 3 (2016), 218–226.
- [43] Hao Zhang and Chong-Wah Ngo. 2018. A fine granularity object-level representation for event detection and recounting. *IEEE Transactions on Multimedia* 21, 6 (2018), 1450–1463.
- [44] Shihao Zhao, Xingjun Ma, Xiang Zheng, James Bailey, Jingjing Chen, and Yu-Gang Jiang. 2020. Clean-label backdoor attacks on video recognition models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14443–14452.
- [45] Bin Zhu, Chong-Wah Ngo, and Jing-jing Chen. 2020. Cross-domain Cross-modal Food Transfer. In *Proceedings of the 28th ACM International Conference on Multimedia*. 3762–3770.