

TWO-STREAM TEMPORAL TRANSFORMER FOR VIDEO ACTION CLASSIFICATION

Nattapong Kurpukdee and Adrian G. Bors

Department of Computer Science,
University of York, York YO10 5DD, UK

nattapong.kurpukdee@york.ac.uk, adrian.bors@york.ac.uk

ABSTRACT

Motion representation plays an important role in video understanding and has many applications including action recognition, robot and autonomous guidance or others. Lately, transformer networks, through their self-attention mechanism capabilities, have proved their efficiency in many applications. In this study, we introduce a new two-stream transformer video classifier, which extracts spatio-temporal information from content and optical flow representing movement information. The proposed model identifies self-attention features across the joint optical flow and temporal frame domain and represents their relationships within the transformer encoder mechanism. The experimental results show that our proposed methodology provides excellent classification results on three well-known video datasets of human activities.

Index Terms— Video Transformer, Optical Flow, Two-Stream video processing, Video Action Classification.

1. INTRODUCTION

Video action classification is an active research area with many applications related to improving the society well-being and security. So far video processing and understanding relied on applying Convolutional Neural Networks (CNNs) to various tasks. Good video classification results have been obtained by multi-stream video processing networks such as the Slow-Fast model [8] where RestNet [14] was used as the backbone for extracting multi-stream data, as well as the Busy-Quiet model [17]. However, recently transformers have emerged as a powerful deep learning network relying to their self-attention mechanism. Current state-of-the-art transformers [6, 24, 39] achieve better results than CNNs on standard video benchmark databases such as UCF101 [31] and HMDB51 [21], Kinetic-400 dataset [19], and Something-Something V2 [13]. Multi-view transformers was lately have shown their efficiency in [39].

Video action recognition, when compared to image classification, raises many additional challenges posed by the uncertainty with respect to the video capturing conditions, because of the variability of the environment or of the capturing parameters. Moreover, video classification also requires

additional training resources including training examples and computational resources. The traditional two-stream architectures [2, 11, 17, 27, 32, 40], consider training separate network branches decomposing the video into two different streams. Lately, neural networks are increasingly being used for optical flow estimation. The performance of the neural optical flow model is excellent when compared with the output of hand-crafted methods as shown in several studies [5, 34].

This research study develops a model for efficiently extracting and using representation features for video classification. Transformers are shown to be able to intrinsically detect the similarities in data. A transformer-based architecture fusing scene and movement features from video through a self-attention mechanism by a video transformer, is proposed in this research study for video classification. The transformer self-attention mechanism is shown to jointly model efficiently video features from the content defined by sets of frames and the movement from the corresponding optical flow.

The rest of the paper is organized as follows. In Section 2 we provide an overview of the related work on video action recognition using deep network architectures. In Section 3 we introduce the proposed two-stream video transformer processing approach. The video action classification performance is evaluated in Section 4 followed by the discussion of the results and comparisons with other models. The conclusions are drawn in Section 5.

2. RELATED WORK

Traditional video classification research has been successful at developing video descriptors encoding both appearance and motion information [2, 11, 27, 30, 32, 33, 38, 40]. Prior approaches [30] employ two-stream convolution neural networks for action recognition in videos. They show that their network combines single spatial features and temporal optical flow features then using class score fusion for video recognition. Non-local video operators have been studied in [16]. In [11] presents the study of fusion features between the single spatial and temporal optical flow feature. Moreover, in [32] the effectiveness by using the optical flow for motion representation was shown. With the optical flow guidance, they can achieve the video action classification. Traditional opti-

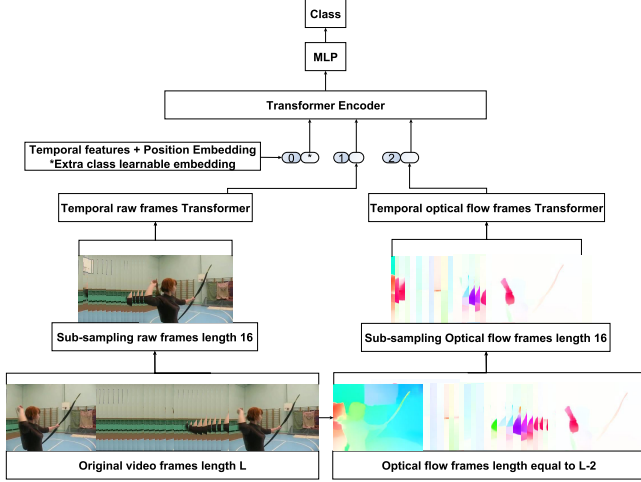


Fig. 1: The two-stream temporal transformer. The raw frames and optical flow are fed as inputs into the transformer neural network. The extracted features are then processed by the transformer encoder. Finally, a Multi-Layer Perceptron (MLP) is used to classify the video classes.

cal flow such as [15] is used to extract motion features and CNNs neural network is a neural network for learning both spatial and temporal features. MoNet [33] also provides the optical flow prediction without the heavy optical flow computation and improves the video classification performance. The authors of [29] also applied optical flow features together with the spatial features and the Long Short-Term Memory (LSTM) CNN [44] were used to achieve video action recognition. In addition, the Slow-Fast model [8] approach employs local appearance and movement information using various features, which are then fused to produce a global video-level descriptor.

Currently, the transformer model [36] is increasingly employed in various video classification tasks [6, 24, 26, 39]. The recent successes of the transformer approach for video classification in the multi-scale vision transformer version 1 [6] and the corresponding improved processing time version 2 [24] by using a transformer for encoding video patches into a deep scale data representation have opened an entire new area of applications for deep networks. The multi-view transformer [39] uses different frame frequencies for different input views. In addition, [12] also uses a two-stream spatio-temporal convolutional network employing a spatial-based transformer as well as a temporal transformer which feeds into a separate MLP. Then, both MLPs are merged for the final video classification output. However, no attempts have been made through using a transformer for identifying self-attention similarities in the content and movement and for optimally fusing motion information and temporal features for characterizing video information. Moreover, a neural network model such as the FlowFormer [18] and RAFT

[34], have shown significant improvements in optical flow estimation over classical hand-crafted movement estimation algorithms.

3. METHODOLOGY

In this study, we aim to classify videos by using two data streams representing both static scene and movement representations as inputs. The scene content is indicated by a series of image frames, obtained by temporally sub-sampling the videos, thus reducing the information being processed, while the movement is represented by the corresponding optical flow. We propose to use an efficient transformer backbone which enables the extraction of video features from selected frames and their corresponding optical flows. The transformer encoder finds self-similarities in such features spanning both fused content and movement latent spaces and fuses such features eventually enabling data characteristics recognition. We present the overview of our proposed model in Fig 1.

3.1. Optical flow as movement representation

First, we consider a neural network for predicting the optical flow instead of using a classical handcraft extraction. In general, the flow will indicate the displacement of every single pixel, or of groups of pixels, between a pair of images. Given a pair of consecutive RGB images, I_1, I_2 , the model will estimate a dense displacement field (f^1, f^2) which maps each pixel (u, v) in I_1 to its corresponding coordinates in I_2 as in the following :

$$(u', v') = (u + f^1(u), v + f^2(v)) \quad (1)$$

In the first stage, of preparing the training data, we sample a sequence of image frames from each given video sequence. Two consecutive frames will be used as a pair for predicting the flow. We consider RGB images for visualizing the optical flow, where the colour indicates the direction of the local movement vector. Examples of various optical flow representations are illustrated in Fig 2b,d,f for the corresponding frames from Fig 2a,c,e, respectively. This demonstrated the clear movement of objects. Consequently, we can apply the process of the movement data in the same way as the sampled video frames.

3.2. Video Transformer

In this section, we describe the video transformer operation, which is used to model the self-attention mechanisms within an increasing spatio-temporal resolution representation of video frames and corresponding movement flow. Both RGB and optical flow inputs are of size $T \times H \times W \times C$, where T corresponds to the number of frames H corresponds to the height W corresponds to the width and C corresponds

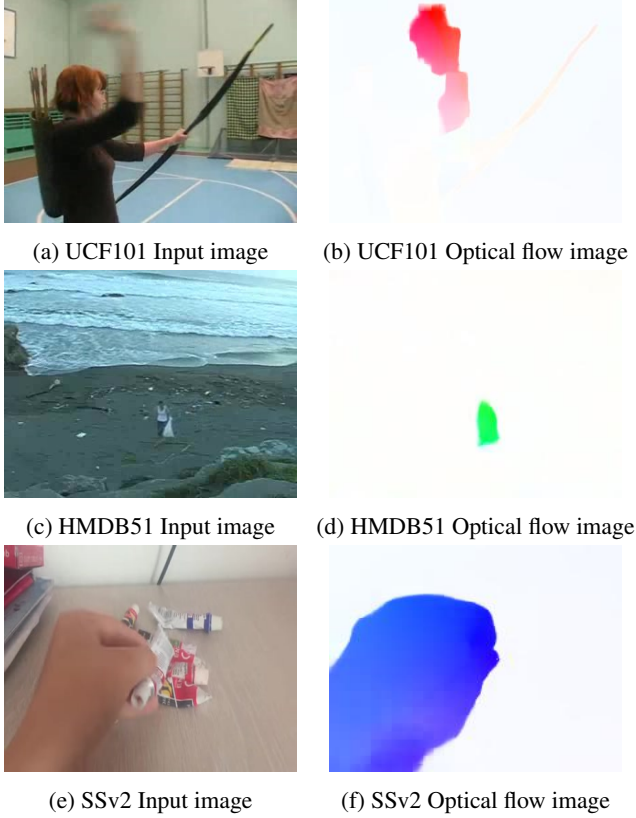


Fig. 2: Examples of frames and their corresponding optical flow prediction, where the optical flow is estimated by the RAFT model [34] on videos from UCF101, HMDB51, and SSv2 dataset. Each direction in the flow is mapped to a RGB color, according to the vector orientation.

to the number of channels. After flattening each image and corresponding optical flow. Segmented spatio-temporal tensor defined patches $\mathbf{X} \in \mathbb{R}^{T \times H \times W \times C}$ are then processed by the transformer. Following the attention mechanism, the transformer projects the input \mathbf{X} into the query tensor $\mathbf{Q} \in \mathbb{R}^{T \times H \times W \times C}$, key tensor $\mathbf{K} \in \mathbb{R}^{T \times H \times W \times C}$, and value tensor $\mathbf{V} \in \mathbb{R}^{T \times H \times W \times C}$. The input of the transformer consists of query and keys of dimension d_k , and values of dimension d_v . Then compute the dot products of the query with all keys, divide each by $\sqrt{d_k}$, and apply a softmax function to obtain the weights on the values as described in :

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} \right) \mathbf{V}. \quad (2)$$

The self-attention mechanism of the transformer is deployed in the spatio-temporal representation of the video space, using pooling operations in order to reduce the latent tensor dimension instead of considering reducing the input pixel tensor dimension. Reducing the latent tensor dimension speeds up the video processing. Finally, the transformer encodes each input into a certain feature vector. Various types

of transformers are considered in the experimental results including the Multi-scale vision transformer (MVITv2) [24] and the video Swin transformer [26], each achieving remarkable successes according to their original research studies.

3.3. Two-Stream Fusion and Classification

Unlike in other applications, such as when using CNN networks, the transformer encoder enables the self-attention mechanism on fused features instead of simply concatenating them into a single vector. The two-stream architecture encodes RGB frame and optical flow frame into output features vectors \mathbf{O}_r and \mathbf{O}_f , respectively. Each output vector considers a D_o dimensional output tensor \mathbf{O} of length L_o , $\mathbf{O}_r \in \mathbb{R}^{L_o \times D_o}$, $\mathbf{O}_f \in \mathbb{R}^{L_o \times D_o}$. The output vector with the extra class learnable vector $\mathbf{X}_{class} \in \mathbb{R}^{L_o \times D_o}$ is mapped with the positional embedding \mathbf{E}_{pos} for the transformer input like in the ViT [4] and described in the following :

$$\mathbf{X} = [\mathbf{X}_{class}; \mathbf{O}_r^1 \mathbf{E}; \mathbf{O}_f^2 \mathbf{E};] + \mathbf{E}_{pos}. \quad (3)$$

Where $\mathbf{E} \in \mathbb{R}^{L_o \times D_o}$, $\mathbf{E}_{pos} \in \mathbb{R}^{(N+1) \times D_o}$, and N is the number of inputs equal to 2. Considers a D dimensional input tensor \mathbf{X} of sequence length L , $\mathbf{X} \in \mathbb{R}^{L \times D}$. Instead of performing a single attention function with D -dimensional keys, values, and queries, we apply the Multi-Head Attention (MHA) in order to linearly project the queries, keys, and values h times with different, learned linear projections to d_k , d_k and d_v dimensions, respectively. Following the MHA, we project the input \mathbf{X} into the query tensor $\mathbf{Q} \in \mathbb{R}^{L \times D}$, key tensor $\mathbf{K} \in \mathbb{R}^{L \times D}$, and value tensor $\mathbf{V} \in \mathbb{R}^{L \times D}$. On each of the projected versions of query, keys, and values we then perform the attention function in parallel, yielding d_v output values. These are concatenated and once again projected, resulting in the final values given by :

$$\begin{aligned} \text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{Concat}(\text{head}_1, \dots, \text{head}_h) \mathbf{W}^o \\ \text{head}_i &= \text{Attn}(\mathbf{Q}\mathbf{W}_i^Q, \mathbf{K}\mathbf{W}_i^K, \mathbf{V}\mathbf{W}_i^V) \end{aligned} \quad (4)$$

where $\mathbf{W}_i^Q \in \mathbb{R}^{L \times D}$, $\mathbf{W}_i^K \in \mathbb{R}^{L \times D}$, $\mathbf{W}_i^V \in \mathbb{R}^{L \times D}$, and $\mathbf{W}^o \in \mathbb{R}^{hd_v \times D}$. In this work, we employ $h = 8$ parallel attention layers, or heads. For each of these, we use $d_k = d_v = D/h = 96$. For the classification module, we apply the MLP network consisting of three layers starting with a normalization layer followed by a dropout layer. Then the final layer is a fully-connected linear layer for the classifier.

4. EXPERIMENTAL RESULTS

In this section, we first discuss the details of the datasets UCF101, HMDB51, and Something-Something V2 used in the experiments. Then we discuss the implementation details followed by the experimental results and their discussion.

4.1. Datasets and Tasks

We evaluate our proposed framework on UCF101 [31] and HMDB51 [21] which are standard datasets for action recognition tasks. The UCF101 dataset consists of 13,320 videos from 101 classes and provides three splits of training, validation, and testing datasets. We randomly select 10% of each training set as the validation set. The HMDB51 dataset consists of 6,766 video examples from 51 action classes and also provides three splits for training and testing datasets. We average the classification accuracy over all three splits to evaluate the video classification results on either UCF101 or HMDB51 dataset. Meanwhile, Something-Something V2 dataset is a large-scale dataset, which contains 220,847 video examples in total from 174 action classes. We consider 168,913 videos for training and 27,157 for testing, while 24,777 videos represent the validation set. Something-Something V2 is characterized by a diversity of movement information, including movements that are rather specific to specialized activities. The information about the three databases used in the experiments is provided in Table 1.

Datasets	Video Per Task		
	Train	Val	Test
HMDB51 Fold-1,2,3	3,570	1,666	1,530
UCF101 Fold-1	7,629	1,908	3,783
UCF101 Fold-2	7,668	1,918	3,734
UCF101 Fold-3	7,699	1,925	3,696
Something-Something V2	168,913	24,777	27,157

Table 1: Data Statistics. We consider 7 splits for video classification tests.

4.2. Implementation Details

We consider the RAFT optical flow model [34] for predicting the movement between two images from a sequence of video frames. We treat the data of the optical flow as the same as the RGB frame. We map the optical flow into RGB representation by using the `flow_to_image` function provided by the `torchvision` utility. Moreover, we consider the Multiscale Vision Transformers for Classification and Detection (MViTv1) [6], Improved Multiscale Vision Transformers for Classification and Detection (MViTv2) [24] or the Video Swin Transformer (Swin) [26], as the backbone for temporal feature extraction. These video transformers are pre-trained on Kinetic-400 dataset [19]. For all transformer architectures, for the sake of using lower computational resources, we decided to use the smallest model architecture called MViTV1-B, MViTV2-S, and Swin-S. Because these transformer models are designed for a small dataset their operation processing is rather quick. These three models use as input videos of 16 frames of size $224 \times 224 \times 3$. The video input

and optical flow input are compressed at the final layer of the channel encoder to 768 channel features.

For the data pre-processing, we follow the protocol for each transformer. For the MViT transformer, we consider a random crop video of 224×224 pixels. Then the pixel values are re-scaled to $[0, 1]$. After that, the values are normalized using the mean= $[0.45, 0.45, 0.45]$ and standard deviation as $\text{std}=[0.225, 0.225, 0.225]$. For the video Swin transformer, we follow the pre-processing by normalized using mean= $[0.485, 0.456, 0.406]$ and $\text{std}=[0.229, 0.224, 0.225]$. We implement our models using PyTorch with a single GPU Tesla V100 32G on a POWER9 architecture. For each dataset, we train a neural network for a maximum of 200 epochs with a small batch size of 8. We optimize the model using Adam [20] with a learning rate of 0.0002 and dropout at 0.5. A single global transformer encoder will set the number of heads at 8 while the dimension of the feedforward network model is set at 1024. For training, we consider the Cross-entropy loss function. In addition, to avoid over-fitting the model, we apply the early stopping strategy to monitor the validation loss rate. If the validation loss rate does not improve for 10 epochs we would stop the training and start the evaluation.

4.3. Video classification results

In this section, we present the experimental results of our proposed video action classification. For the evaluation results, we consider the Top-1 Accuracy (Top-1 ACC) representing the classification accuracy of the model. We present the results on UCF101, HMDB51, and Something-Something V2 which are the standard action recognition datasets.

Baselines We retrain and evaluate three widely used state-of-the-art methods as baselines including MViTv1 [6], MViTv2 [24] and Swin-S [26] in the same environment without any data augmentation. We compare these three baseline training scenarios with our proposed method. Moreover, we also compare with other two-stream based video classification models, which splits the given data in the same manner.

4.3.1. Video classification performance on UCF101 and HMDB51

First, we report the results on the UCF101 and HMDB51 datasets as these two databases are extensively used in video action classification benchmarks. The average results across three splits when considering training on UCF101 and HMDB51 are provided in Table 2, where we compare with baselines and state-of-the-art video classification models. In Table 2 we group the methods used in the tests according to the database used for pre-training the model. We found that our proposed method shows a significant improvement when compared to the baselines and equivalent to the state-of-the-art models on the UCF101 dataset. Moreover, the results on the HMDB51 dataset show that our model has the best

accuracy when compared with other methods. Our model can achieve up to 10.9% increase when considering UCF101, and up to 25.92% increase when considering HMDB51, compared to the nearest baseline.

Classifier	Pre-trained	Average Top-1 Acc	
		UCF101	HMDB51
Two-Stream CNNs [30]	ImageNet	88.00%	59.40%
OFF [32]	-	96.00%	74.20%
Two-Stream CNNs [11]	ImageNet	93.50%	69.20%
Two-Stream I3D [2]	-	93.40%	66.40%
Two-Stream I3D [2]	Imagenet+Kinetics 400	98.00%	80.70%
Two-Stream+LSTM [40]	-	88.60%	-
Two-Stream C3D [27]	-	91.40%	-
Two-Stream TSN [38]	-	94.00%	68.50%
Three-Stream TSN [38]	-	94.20%	69.40%
TDD+iDT [37]	-	91.50%	65.90%
LTC+iDT [35]	-	91.70%	64.80%
ST-ResNet + IDT [10]	-	94.60%	70.30%
MViTv1-B finetune (our baseline)[6]	Kinetics-400	89.66%	66.75%
MViTv2-S finetune (our baseline) [24]	Kinetics-400	92.11%	73.59%
Swin-S finetune (our baseline)[26]	Kinetics-400	82.64%	57.47%
Our (MViTv2-S based)	Kinetics-400	93.54%	83.39%

Table 2: Video classification results on UCF-101 and HMDB-51, representing the average over 3 splits.

4.3.2. Video classification performance on SSv2

Second, we have assessed the proposed methodology on the Something-Something V2 as this dataset is far more complex than the others used in video action classification. The results are provided in Table 3 where we compare with baselines and state-of-the-art models. The methods are grouped according to the database used for pre-training the model. We can observe that the proposed method provides a significant improvement when compared to the baselines and can achieve up to 6.82% increase in the classification performance when compared to the baseline without any data augmentation.

Classifier	Pre-trained	Top-1 Acc
		SSv2
TimeSformer-HR [1]	ImageNet-21K	62.5%
SlowFast R101, 8x8 [9]	Kinetics-400	63.1%
TSM-RGB [25]	ImageNet-1K+K400	63.3%
MSNet [22]	ImageNet-1K	64.7%
TEA [23]	ImageNet-1K	65.1%
ViT-B-TimeSformer [1]	ImageNet-21K	62.5%
ViT-B	ImageNet-21K	63.5%
MViTv1-B, 16x4 [6]	Kinetics-400	64.7%
MViTv2-S, 16x4 [24]	Kinetics-400	68.2%
Swin-B [26]	Kinetics-400	69.6%
bLResNet, 32x2 [7]	SSv2	67.1%
MViTv1-B finetune (our baseline)[6]	Kinetics-400	50.31%
MViTv2-S finetune (our baseline) [24]	Kinetics-400	53.36%
Swin-S finetune (our baseline)[26]	Kinetics-400	49.56%
Our (MViTv2-S based)	Kinetics-400	56.38%

Table 3: Video classification performance comparison to baseline and the state-of-the-art methods on Something-Something V2.

4.4. Discussions

4.4.1. Data augmentation

To observe that the performance of the proposed method relies on limited resources, such as employing the smallest video processing architecture available without using any data augmentation strategy in both the baseline model and our two-stream model. The results show that our proposed method can achieve better video accuracy when compared with the baseline in all datasets. However, in future studies, our proposed model will require a data augmentation strategy when considering learning complex data as same as described in the state-of-the-art techniques [8, 6, 24, 26] such as flipping, Mixup [42] with CutMix [41], Random Erasing [3], and Random Augment [43].

4.4.2. Optical flow

We produce the optical-flow prediction based on a neural network model instead of hand-crafted prediction. The model trained on FlyingChairs [5] + FlyingThings3D [28] dataset. Both datasets are synthetic and not realistic, actually far in their appearance from human activity video. A good optical flow prediction can be shown in Fig 2 with a clear moving object direction and position when mapped into RGB color. This is clear that the optical flow model learned from the synthetic data can perform well on realistic activity data. However, we found some weaknesses in the optical flow prediction such as producing noisy flow images in some cases. Thus, a better movement estimation could improve the classification results. Moreover, the computational cost for estimation optical flow by neural network model is rather high and we have to balance the accuracy of the optical flow and the processing computation required.

5. CONCLUSION

In this paper, we propose a two-stream transformer architecture for video classification. A multi-scale video transformer is used to extract self-attention features from combining sampled image frames and corresponding optical flow. For the two-stream fusion, we apply a pre-trained transformer encoder to learn the relation between the original frame and the optical flow frame. By using a powerful transformer architecture the proposed approach can harness the power of defining well the inter-relationships among the spatio-temporal features from the original image frame and optical flow frame. The experimental results indicate significant improvement when compared with the baselines and other two-stream video processing models. We plan to use this architecture for video continual learning which has many potential applications.

6. REFERENCES

- [1] G. Bertasius, H. Wang, and L. Torresani. Is space-time attention all you need for video understanding? In *Int. Conf. on Machine Learning (ICML)*, volume PMLR 139, pages 813–824, 2021.
- [2] J. Carreira and A. Zisserman. Quo vadis, action recognition? A new model and the Kinetics dataset. In *Proc. IEEE/CVF on Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017.
- [3] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le. RandAugment: Practical automated data augmentation with a reduced search space. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 18613–18624, 2020.
- [4] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conf. on Learning Representations (ICLR)*, *arXiv preprint arXiv:2010.11929*, 2021.
- [5] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox. Flownet: Learning optical flow with convolutional networks. In *Proc. of the IEEE international conference on computer vision*, pages 2758–2766, 2015.
- [6] H. Fan, B. Xiong, K. Mangalam, Y. Li, Z. Yan, J. Malik, and C. Feichtenhofer. Multiscale vision transformers. In *Proc. of the IEEE/CVF International Conference on Computer Vision*, pages 6824–6835, 2021.
- [7] Q. Fan, C.-F. R. Chen, H. Kuehne, M. Pistoia, and D. Cox. More is less: Learning efficient video representations by big-little network and depthwise temporal aggregation. In *Advances in Neural Information Processing Systems*, volume 32, pages 2261–2270, 2019.
- [8] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slow-fast networks for video recognition. In *Proc. of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [9] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slow-fast networks for video recognition. In *Proc. of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019.
- [10] C. Feichtenhofer, A. Pinz, and R. P. Wildes. Spatiotemporal residual networks for video action recognition. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3476–3484, 2016.
- [11] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1933–1941, 2016.
- [12] E. Fish, J. Weinbren, and A. Gilbert. Two-stream transformer architecture for long form video understanding. In *British Machine Vision Conference (BMVC)*, 2022.
- [13] R. Goyal, S. E. Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haanel, I. Fruend, P. Yianilos, M. Mueller-Freitag, F. Hoppe, C. Thureau, I. Bax, and R. Memisevic. The “something something” video database for learning and evaluating visual common sense. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, pages 5843–5851, 2017.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. of the IEEE Conference on Computer vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [15] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981.
- [16] G. Huang and A. G. Bors. Region-based non-local operation for video classification. In *Proc. International Conference on Pattern Recognition (ICPR)*, pages 10010–10017, 2021.
- [17] G. Huang and A. G. Bors. BQN: Busy-Quiet net enabled by motion band-pass module for action recognition. *IEEE Trans. on Image Processing*, 31(7):4966–4979, 2022.
- [18] Z. Huang, X. Shi, C. Zhang, Q. Wang, K. C. Cheung, H. Qin, J. Dai, and H. Li. FlowFormer: a transformer architecture for optical flow. In *European Conf. on Computer Vision (ECCV)*, vol. LNCS 13677, pages 668–685, 2022.
- [19] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: A large video database for human motion recognition. In *Proc. IEEE Int. Conference on Computer Vision (ICCV)*, pages 2556–2563, 2011.
- [22] H. Kwon, M. Kim, S. Kwak, and M. Cho. Motion-Squeeze: neural motion feature learning for video understanding. In *Proc. European Conference on Computer Vision*, vol. LNCS 12361, pages 345–362, 2020.

- [23] Y. Li, B. Ji, X. Shi, J. Zhang, B. Kang, and L. Wang. TEA: Temporal excitation and aggregation for action recognition. In *Proc. of the IEEE/CVF conference on computer vision and pattern recognition*, pages 909–918, 2020.
- [24] Y. Li, C.-Y. Wu, H. Fan, K. Mangalam, B. Xiong, J. Malik, and C. Feichtenhofer. MViTv2: Improved multiscale vision transformers for classification and detection. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4794–4804, 2022.
- [25] J. Lin, C. Gan, and S. Han. TSM: Temporal shift module for efficient video understanding. In *Proc. of the IEEE/CVF International Conference on Computer Vision*, pages 7083–7093, 2019.
- [26] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu. Video swin transformer. In *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3192–3201, 2022.
- [27] T. Lu, S. Ai, Y. Jiang, Y. Xiong, and F. Min. Deep optical flow feature fusion based on 3D convolutional networks for video action recognition. In *Proc. IEEE SmartWorld*, pages 1077–1080, 2018.
- [28] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4040–4048, 2016.
- [29] A. Nebisoy and S. Malekzadeh. Video action recognition using spatio-temporal optical flow video frames. *arXiv preprint arXiv:2103.05101*, 2021.
- [30] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 27:568–576, 2014.
- [31] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [32] S. Sun, Z. Kuang, L. Sheng, W. Ouyang, and W. Zhang. Optical flow guided feature: A fast and robust motion representation for video action recognition. In *Proc. of the IEEE conference on computer vision and pattern recognition*, pages 1390–1399, 2018.
- [33] Y. Tang, L. Ma, and L. Zhou. Hallucinating optical flow features for video classification. *arXiv preprint arXiv:1905.11799*, 2019.
- [34] Z. Teed and J. Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *Proc. European Conference on Computer Vision*, vol. LNCS 12347, pages 402–419, 2020.
- [35] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 40(6):1510–1517, 2017.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [37] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4305–4314, 2015.
- [38] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *Proc. European Conference on Computer Vision (ECCV)*, vol. LNCS 9912, pages 20–36, 2016.
- [39] S. Yan, X. Xiong, A. Arnab, Z. Lu, M. Zhang, C. Sun, and C. Schmid. Multiview transformers for video recognition. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3333–3343, 2022.
- [40] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4694–4702, 2015.
- [41] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo. CutMix: Regularization strategy to train strong classifiers with localizable features. In *Proc. of IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6023–6032, 2019.
- [42] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, *arXiv preprint arXiv:1710.09412*, 2018.
- [43] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random erasing data augmentation. In *Proc. of the AAAI Conf. on Artificial Intelligence*, pages 13001–13008, 2020.
- [44] Y. Zhu, X. Li, C. Liu, M. Zolfaghari, Y. Xiong, C. Wu, Z. Zhang, J. Tighe, R. Manmatha, and M. Li. A comprehensive study of deep video action recognition. *arXiv preprint arXiv:2012.06567*, 2020.