
Temporal-attentive Covariance Pooling Networks for Video Recognition

Zilin Gao[†], Qilong Wang[‡], Bingbing Zhang[†], Qinghua Hu[‡], Peihua Li^{†*}

[†]School of Information and Communication Engineering, Dalian University of Technology

[‡]College of Intelligence and Computing, Tianjin University

gzl@mail.dlut.edu.cn, qlwang@tju.edu.cn, icyzhang@mail.dlut.edu.cn

huqinghua@tju.edu.cn, peihuali@dlut.edu.cn

Abstract

For video recognition task, a global representation summarizing the whole contents of the video snippets plays an important role for the final performance. However, existing video architectures usually generate it by using a simple, global average pooling (GAP) method, which has limited ability to capture complex dynamics of videos. For image recognition task, there exist evidences showing that covariance pooling has stronger representation ability than GAP. Unfortunately, such plain covariance pooling used in image recognition is an orderless representative, which cannot model spatio-temporal structure inherent in videos. Therefore, this paper proposes a Temporal-attentive Covariance Pooling (TCP), inserted at the end of deep architectures, to produce powerful video representations. Specifically, our TCP first develops a temporal attention module to adaptively calibrate spatio-temporal features for the succeeding covariance pooling, approximately producing attentive covariance representations. Then, a temporal covariance pooling performs temporal pooling of the attentive covariance representations to characterize both intra-frame correlations and inter-frame cross-correlations of the calibrated features. As such, the proposed TCP can capture complex temporal dynamics. Finally, a fast matrix power normalization is introduced to exploit geometry of covariance representations. Note that our TCP is model-agnostic and can be flexibly integrated into any video architectures, resulting in TCPNet for effective video recognition. The extensive experiments on six benchmarks (e.g., Kinetics, Something-Something V1 and Charades) using various video architectures show our TCPNet is clearly superior to its counterparts, while having strong generalization ability. *The source code is publicly available.*

1 Introduction

Video recognition aims to automatically analyze the contents of videos (e.g., events and actions), and has a wide range of applications, including intelligent surveillance, multimedia retrieval and recommendation. The great progress of deep learning, especially convolutional neural networks (CNNs) [49, 53, 22, 62], remarkably improves performance of video recognition. Since videos in the wild involve complex dynamics caused by appearance changes and variation of visual tempo, powerful video representations potentially provide high performance of video recognition. However, most of the existing video recognition architectures based on 3D CNNs [54, 44, 2, 11, 55, 10] or 2D CNNs [36, 28, 34, 40] usually generate the final video representations by leveraging a global average pooling (GAP). Such GAP simply computes the first-order statistics of convolution features in a temporal orderless manner, which discards richer statistical information inherent in spatio-temporal features and has limited ability to capture complex dynamics of videos.

*Corresponding author.

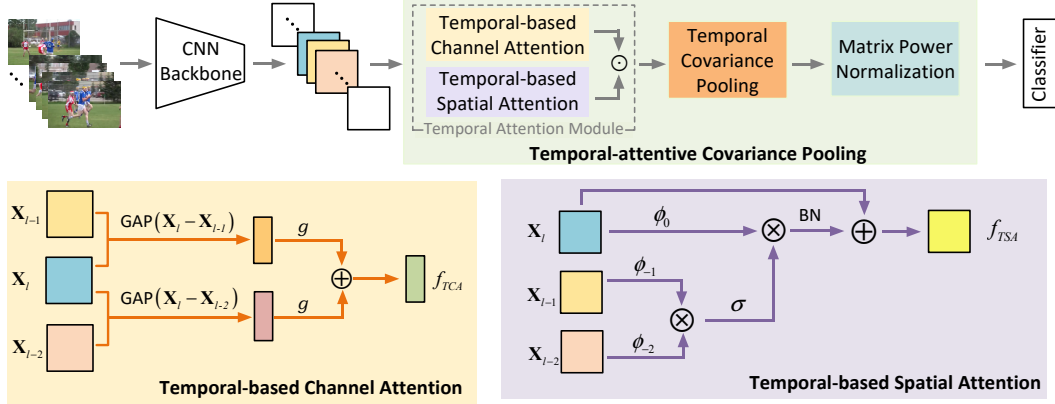


Figure 1: Overview of our TCPNet for video recognition. Its core is a Temporal-attentive Covariance Pooling (TCP) integrated at the end of deep recognition architectures to produce powerful video representations. Specifically, our TCP consists of a temporal attention module for adaptively calibrating spatio-temporal features, a temporal covariance pooling to characterize intra-frame correlations and inter-frame cross-correlations of the calibrated features in a temporal manner, and a fast matrix power normalization for use of geometry of covariances. Refer to Section 3.2 for details.

There are a few researches studying global high-order pooling for improving performance of video recognition [46, 30, 7]. Sanin et al. [46] propose spatio-temporal covariance descriptors as video representations, where covariance descriptors are built with hand-crafted features of candidate regions and a subset of the optimal descriptors are learned by LogitBoost classifiers. However, such a method is designed based on the idea of the classical shallow architectures, without making the most of the merits of deep learning paradigm. In [30], two tensor-based feature representations (namely SCK and DCK) are introduced for action recognition, where multiple subsequences and CNN classifier scores are used to compute tensor representations while a support vector machine is trained for classification (see Figures 2&4 and the third paragraph of Section 5.2). Under end-to-end deep architectures, Diba et al. [7] propose a temporal linear encoding (TLE) method for video classification. Specifically, TLE adopts a bilinear pooling [38] or compact bilinear pooling [15] to aggregate the features output by 2D and 3D CNNs as the entire video representations. TLE shows better performance than the original fully-connected (FC) layer, but it fails to fully exploit temporal information and geometry of covariance representations, limiting effectiveness of covariance pooling for video recognition. Recently, there also exist evidences showing that global covariance (second-order) pooling can significantly improve performance of deep CNNs in various *image-related* tasks, such as fine-grained visual recognition [38, 15], large-scale image classification [32, 31, 61] and visual question answering [13, 69]. However, these methods aim to generate image representations, without considering important temporal information inherent in videos. Meanwhile, naive usage of *image-related* covariance pooling (called *plain covariance pooling* in this work) technique for *video-related* tasks cannot make full use of the advantage of covariance pooling [64].

As discussed previously, though global covariance pooling shows the potential to improve recognition performance, aforementioned works suffer from the following drawbacks: (1) Shallow approach [46] fails to make the most of the merits of deep learning paradigm; (2) Plain covariance pooling methods [38, 15, 32, 31, 61] as well as TLE [7] cannot adequately model complex temporal dynamics, and in the meanwhile TLE considers no geometry of covariances. To overcome these issues, this paper proposes a Temporal-attentive Covariance Pooling (TCP), which can be integrated into existing deep architectures to produce powerful video representations. Different from the plain covariance pooling methods which perform orderless aggregation along temporal dimension, our TCP presents a temporal pooling of attentive covariance representations to capture complex temporal dynamics. Specifically, as illustrated in Figure 1, our TCP consists of a temporal attention module, a temporal covariance pooling, and a fast matrix power normalization of covariances. The temporal attention module is developed to adaptively calibrate spatio-temporal features for the subsequent covariance pooling. The covariance pooling of the calibrated features can be approximately regarded as producing attentive covariance representations, having the ability to handle complex dynamics. Then, temporal pooling of the attentive covariance representations (i.e., temporal covariance pooling) characterizes both intra-frame correlations and inter-frame cross-correlations of the calibrated features, which can fully harness the complex temporal information. The fast matrix power normalization [31] is finally

introduced to exploit geometry of covariance spaces. Note that our TCP is model-agnostic and can be flexibly integrated into existing 2D or 3D CNNs in a plug-and-play manner, resulting in an effective video recognition architecture (namely TCPNet).

The overview of our TCPNet is shown in Figure 1. To verify its effectiveness, extensive experiments are conducted on six video benchmarks (i.e., Mini-Kinetics-200 [67], Kinetics-400 [2], Something-Something V1 [19], Charades [47], UCF101 [50] and HMDB51 [23]) using various deep architectures (e.g., TSN [60], X3D [10] and TEA [34]). The contributions of our work are summarized as follows. (1) We propose an effective TCPNet for video recognition, where a model-agnostic Temporal-attentive Covariance Pooling (TCP) integrated at the end of existing deep architectures can produce powerful video representations. (2) To our best knowledge, our TCP makes the first attempt to characterize both intra-frame and inter-frame correlations of the calibrated spatio-temporal features in a temporal manner, which approximatively performs temporal pooling of the attentive covariance representations. The proposed TCP provides a compelling alternative to GAP for improving recognition performance of deep video architectures. (3) We perform a comprehensive study of global covariance pooling for deep video recognition architectures, and conduct extensive experiments on six challenging video benchmarks. The results show our TCPNet has strong generalization ability while performing favorably against state-of-the-arts.

2 Related Works

2.1 Deep Architectures for Video Recognition

Video data usually consists of appearance information and temporal motion. Two-stream architectures [48, 12] are specially developed for separately capturing the two kind of information. They obtain competitive performance, but suffer from huge computational cost due to offline computation of optical flow. Extending 2D convolutions to 3D ones, 3D CNNs achieve state-of-the-art performance by taking only RGB inputs [54, 2, 44, 67, 55, 11, 10], which can simultaneously learn appearance and temporal features. Furthermore, a family of attention models based on idea of non-local [62, 65, 3, 21] are proposed to further improve performance of 3D CNNs by adaptively calibrating spatio-temporal features. Recently, integration of temporal modeling with 2D CNNs achieves promising results in a more efficient way [36, 28, 9, 51, 34, 40, 41], where 2D convolutions extract appearance features while temporal modules (e.g., channel shift, temporal difference, and temporal convolution) are used to capture motion information. However, these video recognition architectures usually generate the final video representations by leveraging a global average pooling (GAP), which has the limited ability to capture complex dynamics of videos. Some works exploit classical encoding methods (i.e., VLAD [27] and Fisher vector [45]) to replace GAP for aggregating spatio-temporal features [18, 58]. Different from aforementioned methods, we propose a temporal-attentive covariance pooling (TCP) to replace GAP for producing powerful video representations in deep architectures.

2.2 Deep Covariance (Second-order) Pooling

Global covariance (second-order) pooling has been successfully applied to various image-related tasks, such as image classification [38, 26, 15, 32, 31, 16, 61] and visual question answering [13, 69]. These works show global covariance pooling is much better than the original GAP, and appropriate use of geometry of covariances plays a key role for achieving high performance [26, 32]. However, these works have no ability to capture temporal information inherent in video data, and naive usage of them cannot fully exploit advantages of covariance pooling [64]. For modeling videos based on deep second-order pooling, TLE [7] makes an attempt to use bilinear pooling [38] and CBP [15] for aggregating spatio-temporal features in an end-to-end manner. However, this method makes no full use of temporal information and geometry of covariance representations. Besides, some researches [17, 72, 35] integrate local second-order statistics into deep architectures for improving performance of video recognition. Specifically, AP [17] introduces second-order statistics by designing an element-wise multiplication of multiple-branch based on rank-1 approximation. Sharing similar idea, ABM [72] and TBN [35] apply element-wise multiplication to various frames for capturing motion information. Different from above works, our proposed TCP performs a global temporal pooling of attentive covariance representations, which can handle complex temporal dynamics and exploit geometry of covariances. Besides, our TCPNet has the ability to fully harness the merits of deep learning paradigm, compared to shallow approach [46].

3 Proposed Method

In this section, we first briefly introduce the plain covariance pooling for modeling videos. Then, we describe in detail our proposed temporal-attentive covariance pooling (TCP), which consists of a temporal attention module, a temporal covariance pooling, and a fast matrix power normalization. Finally, we instantiate our TCP with integration of existing deep video architectures, resulting in temporal-attentive covariance pooling networks (TCPNet).

3.1 Plain Covariance Pooling for Modeling Videos

For the existing video architectures, global average pooling (GAP) is widely used to generate the final video representations. Suppose we have a total of L feature tensors output by some deep architecture. Let H, W and C be spatial height, width and channel number of one feature tensor. We denote by $\mathbf{X}_l \in \mathbb{R}^{N \times C}$ the matrix obtained from the l -th feature tensor where $N = H \times W$ and each row $\mathbf{x}_{l,n}$ is a C -dimensional feature. GAP computes the final video representation $\mathbf{p}_{GAP} \in \mathbb{R}^{1 \times C}$ as

$$\mathbf{p}_{GAP} = \frac{1}{LN} \sum_{l=1}^L \sum_{n=1}^N \mathbf{x}_{l,n}, \quad (1)$$

where L indicates temporal resolution. As shown in Eqn. (1), from statistical view, GAP estimates first-order moment while higher-order statistics containing richer information are discarded.

Compared to GAP, covariance pooling can capture richer statistics by modeling correlation among the features. For video data, plain covariance pooling generates video representation $\mathbf{P}_{PCP} \in \mathbb{R}^{C \times C}$ by aggregating the features as

$$\mathbf{P}_{PCP} = \frac{1}{LN} \sum_{l=1}^L \sum_{n=1}^N \mathbf{x}_{l,n}^T \mathbf{x}_{l,n} = \frac{1}{L} \sum_{l=1}^L \frac{1}{N} \mathbf{X}_l^T \mathbf{X}_l, \quad (2)$$

where T indicates the matrix transpose. Although \mathbf{P}_{PCP} in Eqn. (2) has better representation ability than \mathbf{p}_{GAP} , both of them perform orderless aggregation along temporal dimension based on simple average operation, lacking the ability to handle complex temporal dynamics. In addition, naive use of \mathbf{P}_{PCP} only obtains small gains (see Section 4) as geometry of covariance spaces is neglected.

3.2 Temporal-attentive Covariance Pooling (TCP)

To overcome the drawbacks of plain covariance pooling, we propose a temporal-attentive covariance pooling (TCP) to capture richer statistics of spatio-temporal features and handle complex dynamics.

3.2.1 Temporal Covariance Pooling

We regard the video signal as a stationary stochastic process [39] and propose to model its statistics by the second-order moments in a learnable manner. Specifically, we perform a temporal-attentive covariance pooling \mathbf{P}_{TCP} for summarizing the statistics of the videos:

$$\mathbf{P}_{TCP} = \text{TP}_L(\{\mathbf{C}_1, \dots, \mathbf{C}_l, \dots, \mathbf{C}_L\}), \quad \mathbf{C}_l = \frac{1}{N} \mathbf{X}_l^T \mathbf{X}_l, \quad (3)$$

where $\text{TP}_L(\cdot)$ indicates a temporal pooling of frame-wise covariance representations $\{\mathbf{C}_l\}_{l=1, \dots, L}$. In this work, we achieve $\text{TP}_L(\cdot)$ with a temporal convolution of covariances followed by average pooling:

$$\text{TP}_L(\{\mathbf{C}_1, \dots, \mathbf{C}_l, \dots, \mathbf{C}_L\}) = \frac{1}{L} \sum_{l=1}^L \text{TC-COV}_\kappa(\mathbf{X}_l), \quad (4)$$

where $\text{TC-COV}_\kappa(\mathbf{X}_l)$ performs temporal convolution of covariance \mathbf{X}_l with kernel size of κ . By taking $\kappa = 3$ as an example, we achieve $\text{TC-COV}_{\kappa=3}(\mathbf{X}_l)$ by considering the temporal correlations of intra- and inter- frames in a sliding window manner. Thus, $\text{TC-COV}_{\kappa=3}(\mathbf{X}_l)$ is computed as

$$\begin{aligned} \text{TC-COV}_{\kappa=3}(\mathbf{X}_l) = & \underbrace{\mathbf{W}_{-1}^T \mathbf{X}_{l-1}^T \mathbf{X}_{l-1} \mathbf{W}_{-1} + \mathbf{W}_0^T \mathbf{X}_l^T \mathbf{X}_l \mathbf{W}_0 + \mathbf{W}_1^T \mathbf{X}_{l+1}^T \mathbf{X}_{l+1} \mathbf{W}_1}_{\text{intra-frame covariance pooling}} \quad (5) \\ & + \underbrace{\mathbf{W}_{-1}^T \mathbf{X}_{l-1}^T \mathbf{X}_l \mathbf{W}_0 + \dots + \mathbf{W}_0^T \mathbf{X}_l^T \mathbf{X}_{l-1} \mathbf{W}_{-1} + \dots + \mathbf{W}_1^T \mathbf{X}_{l+1}^T \mathbf{X}_l \mathbf{W}_0}_{\text{inter-frame cross-covariance pooling}} \end{aligned}$$

where $\{\mathbf{W}_*\}_{* \in \{-1,0,1\}}$ are a set of learnable parameters. For Eqn. (5), the first part of terms is a combination of frame-wise covariances learned based on geometry-aware projections [20], aiming to characterize intra-frame correlations. Meanwhile, the second part of terms is a combination for learnable cross-covariances between all pairs of different frames, aiming to capture inter-frame correlations. In this work, we model the distributions of stochastic process describing the video signals by second-order moments. Along this line, we can further exploit higher-order (e.g., third-order) moments [5] in a learnable manner for more accurate distribution modeling, which is left for future work.

3.2.2 Temporal Pooling of Attentive Covariances

Although Eqn. (5) performs a temporal convolution on covariances, both intra-frame covariances (e.g., $\mathbf{C}_{l,l} = \mathbf{X}_l^T \mathbf{X}_l$, $1/N$ is omitted for simplicity) and inter-frame cross-covariances (e.g., $\mathbf{C}_{l,l-1} = \mathbf{X}_l^T \mathbf{X}_{l-1}$) are computed in a specific manner, having limited ability to handle complex temporal dynamics. Therefore, we can calculate attentive covariances (or cross-covariances) to adaptively calibrate importance of covariance representations to account for dynamic changes. However, element-wise attention for covariances (e.g., $s \cdot \mathbf{C}_{l,l}$ or $\mathbf{S} \odot \mathbf{C}_{l,l}$) will break up their geometry. In this work, we approximatively compute attentive covariances by developing a temporal attention module to calibrate the features before temporal covariance pooling:

$$\text{TC-COV}_{\kappa=3}(\widehat{\mathbf{C}}_l) = \underbrace{\cdots + \mathbf{W}_0^T \widehat{\mathbf{C}}_{l,l} \mathbf{W}_0 + \cdots}_{\text{intra-frame covariance pooling}} + \underbrace{\cdots + \mathbf{W}_0^T \widehat{\mathbf{C}}_{l,l+1} \mathbf{W}_1 \cdots}_{\text{inter-frame cross-covariance pooling}}, \quad (6)$$

where $\widehat{\mathbf{C}}_{l,l} = \widehat{\mathbf{X}}_l^T \widehat{\mathbf{X}}_l$, and the features $\widehat{\mathbf{X}}_l$ can be adaptively calibrated by a temporal-based spatial attention f_{TSA} and a temporal-based channel attention f_{TCA} , i.e.,

$$\widehat{\mathbf{X}}_l = (\mathbf{X}_l \oplus f_{TSA}(\mathbf{X}_{l-2}, \mathbf{X}_{l-1}, \mathbf{X}_l)) \odot f_{TCA}(\mathbf{X}_{l-2}, \mathbf{X}_{l-1}, \mathbf{X}_l), \quad (7)$$

where \oplus and \odot represent element-wise addition and element-wise multiplication, respectively. Both f_{TSA} and f_{TCA} are functions of \mathbf{X}_l and its two preceding frames (i.e., \mathbf{X}_{l-2} and \mathbf{X}_{l-1}). Next, we will describe how to compute f_{TSA} and f_{TCA} .

Computation of f_{TSA} Given \mathbf{X}_{l-2} , \mathbf{X}_{l-1} and \mathbf{X}_l , we compute f_{TSA} as

$$f_{TSA}(\mathbf{X}_{l-2}, \mathbf{X}_{l-1}, \mathbf{X}_l) = BN(\sigma[\phi_{-1}(\mathbf{X}_{l-1})\phi_{-2}^T(\mathbf{X}_{l-2})] \otimes \phi_0(\mathbf{X}_l)), \quad (8)$$

where BN means batch normalization while σ is softmax operation. \otimes indicates matrix multiplication. ϕ_0 , ϕ_{-1} and ϕ_{-2} are linear transformations implemented by convolutions. As shown in Eqn. (8), we extract spatial importance of individual features of \mathbf{X}_l by considering relationship among three consecutive frames based on the idea of attention [56], where \mathbf{X}_{l-1} , \mathbf{X}_{l-2} and \mathbf{X}_l indicates queries, keys and values, respectively. In this way, we model long-range dependencies of spatial features while attending to the temporal relations based on the inherent smoothness between adjacent frames.

Computation of f_{TCA} We calibrate the feature channels of l -th frame by considering temporal difference. Specifically, we formulate f_{TCA} as

$$f_{TCA}(\mathbf{X}_{l-2}, \mathbf{X}_{l-1}, \mathbf{X}_l) = \frac{1}{2}g(\text{GAP}(\mathbf{X}_l - \mathbf{X}_{l-1})) + \frac{1}{2}g(\text{GAP}(\mathbf{X}_l - \mathbf{X}_{l-2})), \quad (9)$$

where GAP indicates global average pooling, and g is achieved by two FC layers followed by a Sigmoid function.

After some arrangement, our TCP in Eqn. (6) can be efficiently implemented by a temporal convolution of the calibrated features with kernel size of $1 \times 1 \times \kappa$ followed by a covariance pooling:

$$\mathbf{P}_{TCP} = \frac{1}{L} \sum_{l=1}^L [\text{TC}_{1 \times 1 \times \kappa}(\widehat{\mathbf{X}}_l)]^T [\text{TC}_{1 \times 1 \times \kappa}(\widehat{\mathbf{X}}_l)]. \quad (10)$$

where $\text{TC}_{1 \times 1 \times \kappa}(\widehat{\mathbf{X}}_l)$ performs temporal convolution on $\widehat{\mathbf{X}}_l$ with kernel size of $1 \times 1 \times \kappa$. Finally, let us compare GAP, plain covariance pooling and our TCP from the view of stochastic process. Since the video signals are high-dimensional, spanning temporal and spatial dimensions, their distribution is very complex. GAP (resp. plain covariance pooling) only considers mean and variances of *one single time*, while our TCP models two different times, capturing temporal dynamics. In addition, as the features possess spatial structure, various features (either nearby or long-range features) may function differently in the interplay of different times, but GAP and plain covariance pooling neglect this interplay. In contrast, our TCP can attend to this by using idea of attention.

3.2.3 Fast Matrix Power Normalization

Our TCP produces a covariance representation for each video. It is well known that the space of covariances is a Riemannian manifold [1, 43], which has geometrical structure. The appropriate use of geometry plays a key role for the effectiveness of covariance representations. Under deep architectures, recent works [32, 61] show matrix power normalization approximately yet effectively utilizes geometry of covariance space, and is superior to its counterparts (e.g., Log-Euclidean metrics [38] and element-wise power normalization [26]). Here we exploit a fast matrix power normalization method [31] for efficiency and effectiveness. Specifically, for each covariance representation \mathbf{P}_{TCP} output by our TCP, we compute its approximate matrix square root¹ as follows:

$$\text{Iteration: } \{ \mathbf{Q}_k = \frac{1}{2} \mathbf{Q}_{k-1} (3\mathbf{I} - \mathbf{R}_{k-1} \mathbf{Q}_{k-1}); \mathbf{R}_k = \frac{1}{2} (3\mathbf{I} - \mathbf{R}_{k-1} \mathbf{Q}_{k-1}) \mathbf{R}_{k-1} \}_{k=1, \dots, K}, \quad (11)$$

where $\mathbf{Q}_0 = \mathbf{P}_{TCP}$ and $\mathbf{R}_0 = \mathbf{I}$ is the identity matrix. After K iterations, we have $\mathbf{P}_{TCP}^{\frac{1}{2}} \approx \mathbf{Q}_K$. Finally, $\mathbf{P}_{TCP}^{\frac{1}{2}}$ is used as video representation for classification. In this paper, K is set to 3 throughout all the experiments. Note that Eq. (11) can be efficiently implemented on GPU due to involving of only matrix multiplications, and more details can refer to [31].

3.3 Temporal-attentive Covariance Pooling Networks (TCPNet)

In this subsection, we instantiate our TCP with existing deep video architectures, which yield powerful TCPNet for video recognition. As shown in Figure 1, given any specific video architecture (e.g., TSN [60] with 2D ResNet-50 [22] or 3D CNNs [10]), our TCPNet inserts the proposed TCP after the last convolution layer, which replaces the original GAP to produce the final video representations for classification. Note that we introduce $1 \times 1 \times 1$ convolutions between the last convolution layer and our TCP, which reduce dimension of convolution features and control size of video representations. Our TCP is lightweight: taking 8256-dimensional TCP (default setting) as an example, parameters and FLOPs of our TCP are 3.3M and 1.2G for 8 frames of 224×224 inputs, respectively. Compared to backbones of 2D ResNet-50 (24.3M and 33G) and 2D ResNet-152 (59.0M and 91.5G), our TCP brings about extra 13.6% (5.6%) parameters and 3.6% (1.3%) FLOPs, respectively.

4 Experiments

To evaluate our proposed method, we conduct experiments on six video benchmarks, including Mini-Kinetics-200 (Mini-K200) [67], Kinetics-400 (K-400) [2], Something-Something V1 (S-S V1) [19], Charades [47], UCF101 [50] and HMDB51 [23]. We first describe implementation details of our method, and then make ablation studies on Mini-K200. Furthermore, we compare with counterparts and state-of-the-arts on K-400, and finally show generalization of our TCPNet on remaining datasets.




Figure 2: Effect of feature dimension.

Method	Top-1 (%)	Top-5 (%)
GAP (Baseline)	71.9	90.6
Plain GCP	72.1	90.8
Plain GCP + MPN	74.5	91.8
+ f_{TCA}	75.1	92.2
+ f_{TSA}	75.6	92.5
+TC-COV $_{\kappa=5}$	76.1	92.7

Table 1: Impact of various key components.

Implementation Details. To construct our TCPNet, we integrate the proposed TCP with several video architectures, including TSN [60] with 2D CNNs (i.e., ResNet-50 [22] and ResNet-152 [22]), X3D-M [10] and 2D CNNs with temporal modeling (e.g., TEA [34]). ResNet-50 (R50) and ResNet-152 (R152) are pre-trained on ImageNet-1K and ImageNet-11K, respectively. As suggested in [32, 31], we discard the last downsampling operation when covariance representations are computed. Here we describe the settings of hyper-parameters on Mini-K200 and K-400. For training our TCPNet

¹The matrix square root is a special case of matrix power normalization with the power of 1/2, usually achieving the best performance.

Method	Dim.	Top-1 (%)	Top-5 (%)
GAP	2048	71.9	90.6
BCNN [38]	16384	72.7	91.4
CBP [15]	8256	72.5	91.5
TLE (BCNN-A) [7]	8256	72.2	91.0
TLE (CBP-A) [7]	8256	72.2	90.6
iSQRT [31]	8256	74.5	91.8
iSQRT-A	8256	72.0	90.5
iSQRT-B	8256	74.3	91.2
TCP (Ours)	8256	76.1	92.7

Table 2: Comparison of various global pooling methods on Mini-K200.

Method	Backbone	Input	Top-1 (%)
TBN [35]	2D R18	112	69.5
MARS [4]	3D RX101	112	72.8
V4D [71]	3D R50	224	80.7
BAT [3]	2D R50	224	70.6
RMS [29]	3D Slow50	224	78.6
CGNL [70]	3D R101	224	79.5
TCPNet (Ours)	2D R50	112	78.3
TCPNet (Ours)	2D R50	224	80.7

Table 3: Comparison of state-of-the-arts on Mini-K200.

with 2D CNNs, we adopt the same data augmentation strategy as [60], and number of segments is set to 8 or 16. A dropout with a rate of 0.5 is used for the last FC layer. TCPNet is optimized by mini-batch stochastic gradient descent (SGD) with a batch size of 96, a momentum of 0.9 and a weight decay of $1e-4$. The whole networks are trained within 50 epochs, where initial learning rate is 0.015 and decay by 0.1 every 20 epochs. For training our TCPNet with X3D-M, we process the images followed by [10], and 16 frames are sampled as inputs. The SGD with cosine training strategy is used to optimize the network parameters within 100 epochs, and the initial learning rate is set to 0.1. For training our TCPNet with TEA, we first train the backbone of TEA (i.e. Res2Net50 [14]) with covariance pooling on ImageNet-1K, and then fine-tune TCPNet with TEA using the same settings as [34]. During inference, the strategies of 1 crop \times 1 clip and 3 crops \times 10 clips are used for ablation study and comparison with state-of-the-arts. The training settings for the remaining datasets will be specified in the following parts. All programs are implemented by Pytorch and run on a PC equipped with four NVIDIA Titan RTX GPUs.

4.1 Ablation Study on Mini-K200

For efficiency, we make ablation studies on Mini-K200 with input size of 112×112 , and TSN with R50 is used as backbone model. We first assess the effects of feature dimension and key components on our TCP, and then compare with different global pooling methods and state-of-the-arts.

Effect of Feature Dimension As described in section 3.3, we introduce $1 \times 1 \times 1$ convolutions before our TCP to reduce dimension of convolution features, controlling the size of covariance representations. Here we assess the effect of feature dimension on recognition accuracy. Note that we do not use temporal attention module and temporal pooling for eliminating their effects and guaranteeing generality. Specifically, the feature dimension varies from 64 to 256, and the corresponding results are shown in Figure 2, where GAP is compared as a baseline. From Figure 2 we can see that global covariance pooling (GCP) consistently outperforms GAP. With increase of feature dimension, performance saturates at dimension of 128, but after that parameters significantly increase. GCP with dimension of 128 is better than GAP by 2.6%, and brings extra $\sim 2M$ parameters. For efficiency and performance, we set feature dimension to 128 throughout the following experiments, producing 8256-dimensional covariance representations after matrix triangulation.

Impact of Key Components Our TCP involves several key components, including temporal attention module, temporal covariance pooling TC-COV $_{\kappa}$ ($\kappa = 5$ and $\kappa = 9$ for 8 and 16 frame inputs) and matrix power normalization. To evaluate their effects, we give the results of our TCP under various configurations. As shown in Table 1, plain covariance pooling obtains little improvement over the original GAP and matrix power normalization (MPN) brings about 2.4% gains. Note that Plain GCP + MPN is a strong baseline established by our work. Although promising performance is achieved by this baseline, the temporal-based channel attention f_{TCA} achieves 0.6% gains over plain covariance pooling with matrix power normalization, while temporal-based spatial attention f_{TSA} further brings 0.5% gains. Furthermore, TC-COV $_{\kappa=5}$ achieves 0.5% improvement by performing temporal covariance pooling. Our TCP finally obtains the best results (76.1%), outperforming GAP and plain covariance pooling by 4.2% and 4.0% respectively. These results show the effectiveness of each key component in TCP.

(a) GAP vs. TCP (Ours) under various architectures.

Method	Pooling	Top-1	Top-5
TSN [60]+R50	GAP	70.6	89.2
	TCP	75.3 (\uparrow 4.7)	91.7 (\uparrow 2.5)
TSN [60]+R152	GAP	75.9	92.1
	TCP	78.3 (\uparrow 2.4)	93.7 (\uparrow 1.6)
X3D-M [10]	GAP	76.0	92.3
	TCP	77.0 (\uparrow 1.0)	92.7(\uparrow 0.4)
TEA [34]+R50	GAP	75.0	91.8
	TCP	76.8 (\uparrow 1.8)	92.9(\uparrow 1.1)

(b) Comparison with various 3D CNNs.

Method	Backbone	Fr \times Views	Top-1	Top-5
I3D [2]	3D Inc.	64 \times N/A	71.1	89.3
R(2+1)D [55]	3D R34	16 \times 10	72.0	90.0
Non-local [62]	3D R101	32 \times 30	76.0	92.1
ABM [72]	3D Inc.	150 \times 1	72.6	N/A
CorrNet [57]	3D R50	32 \times 10	77.2	N/A
SlowFast [11]	3D R50	64 \times 30	77.0	92.6
X3D [10]	X3D-M	16 \times 30	76.0	92.3
TCPNet (Ours)	X3D-M	16 \times 30	77.0	92.7

(c) Comparison with various 2D CNNs (8 frames).

Method	Backbone	Fr \times Views	Top-1	Top-5
TSN [36]	2D R50	8 \times 10	70.6	89.2
TPN [68]	2D R50	8 \times 10	73.5	N/A
TEA [34]	2D R50	8 \times 30	75.0	91.8
BAT [3]	2D R50	8 \times 30	75.8	N/A
GTA [21]	2D R50	8 \times 30	75.9	92.2
SmallBig [33]	2D R50	8 \times 30	76.3	92.5
TCPNet (Ours)	TEA R50	8 \times 30	76.8	92.9
TCPNet (Ours)	TSN R152	8 \times 30	78.3	93.7

(d) Comparison with various 2D CNNs (16 frames).

Method	Backbone	Fr \times Views	Top-1	Top-5
STM [28]	2D R50	16 \times 30	73.7	91.6
TSM [36]	2D R50	16 \times 10	74.7	N/A
TEA [34]	2D R50	16 \times 30	76.1	92.5
TEINet [40]	2D R50	16 \times 30	76.2	92.5
TAM [41]	2D R50	16 \times 12	76.9	92.9
bL-TAM [9]	2D R50	24 \times 9	73.5	91.2
TCPNet (Ours)	TEA R50	16 \times 30	77.2	93.1
TCPNet (Ours)	TSN R152	16 \times 30	79.3	94.0

Table 4: Comparison (in %) of state-of-the-arts on K-400 dataset under various configurations.

Comparison of Various Global Pooling Methods To further verify the effectiveness of our TCP, we compare with eight kinds of global pooling methods, including GAP, BCNN [38], CBP [15], iSQRT [31], TLE (BCNN-A & CBP-A) [7], iSQRT-A and iSQRT-B. Specifically, we compute GAP, BCNN, CBP and iSQRT by using all features $\{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_L\}$ in temporal orderless manner. For BCNN-A, CBP-A and iSQRT-A, we first perform average pooling along temporal dimension to obtain $\tilde{\mathbf{X}} = \sum_i^L \mathbf{X}_i$ upon which we use BCNN, CBP and iSQRT. For iSQRT-B, we first apply iSQRT to each frame \mathbf{X}_i , and then perform average pooling for all covariances. The results of all methods are presented in Table 2, based on which we can draw the following conclusions. First, all GCP methods outperform GAP, showing that GCP can produce more powerful video representations. It can also be seen that iSQRT performs significantly better than BCNN and CBP, clearly suggesting that the advantage of exploiting geometry of covariance spaces. Second, BCNN, CBP and iSQRT achieve better results than TLE and iSQRT-A/iSQRT-B, respectively; this suggests that the simple strategy to exploit temporal information [7] brings little gains. Finally, our TCP achieves the best results, which outperforms GAP and iSQRT by 4.2% and 1.6% in Top-1 accuracy, respectively.

Comparison with State-of-the-arts At the end of this part, we compare our TCPNet with several state-of-the-arts on Mini-K200. The compared results are given in Table 3. When 112×112 images are used as inputs, our TCPNet outperforms TBN [35] and MARS [4] by 8.8% and 5.5%, respectively. When size of input images increases to 224×224 , our TCPNet obtains 80.7% in Top-1 accuracy, performing best among all competing methods. Notably, for backbone of 2D R50, TCPNet improves BAT [3] over 10.1%. Finally, TCPNet with 2D backbone is superior to CGNL [70] and RMS [29] while obtaining the same result with V4D [71]; note that all three latter exploit 3D backbone models.

4.2 Comparison with State-of-the-arts on Kinetics-400

In this subsection, we compare our method with state-of-the-arts on K-400 under various configurations. As shown in Table 4 (a), our TCP consistently outperforms the original GAP under video architectures of TSN, X3D-M and TEA, which span all types of popular video architectures. Particularly, TCP achieves larger improvement if less temporal information is learned by backbone models. Table 4 (b) compares different 3D CNNs, where our TCPNet achieves

Method	Backbone	Top-1 (%)
TSN [60]	2D R50	19.7
TSM [36]	2D R50	45.6
TPN [68]	2D R50	49.0
TSM + NeXtVLAD [37]	Inc.-Res	41.9
TSM + PPAC [42]	Inc.-Res	43.7
GSM [52]	Inc.-V3	49.0
TSN + TCP (Ours)	2D R50	38.3 (\uparrow 18.6)
TSM + TCP (Ours)	2D R50	46.9 (\uparrow 1.3)
TPN + TCP (Ours)	2D R50	49.6 (\uparrow 0.6)

Table 5: Comparison on S-S.

Method	Backbone	Frames	mAP (%)
Non-local [62]	3D R101	128	39.5
SlowFast [66]	3D R50-NL	64	38.0
LFB [65]	3D R101-NL	32	42.5
STRG [63]	3D R101-NL	32	39.7
TimeCeption [25]	3D R101	128	41.1
SlowFast [11]	3D R101-NL	128	42.5
TCPNet (Ours)	3D R101-NL	32	42.6

Table 6: Comparison on Charades.

comparable performance with state-of-the-arts. Particularly, our TCPNet shows similar performance with SlowFast and CorrNet, but uses less input frames. As listed in Table 4 (c) and Table 4 (d), our TCPNet based on TEA with input of 8 frames outperforms other methods, when R50 is used. For input of 16 frames, TCPNet based on TEA is superior to all the compared methods by a clear margin. When R152 is employed, TCPNet obtains state-of-the-art results for inputs of both 8 and 16 frames.

4.3 Generalization to Other Benchmarks

We further verify the generalization ability of our TCP on S-S V1, Charades, UCF101 and HMDB51. For S-S V1, we integrate our TCP into three video architectures, i.e., TSN [60], TSM [36] and TPN [68]. Specifically, we train our TCPNet based on TSN and TSM following the settings in [36] with 8 frames and test with 1 clip \times 1 crop, while TCPNet based on TPN is trained using the same settings as [68]. As shown in Table 5, our TCP brings 18.6%, 1.3% and 0.6% gains for TSN, TSM and TPN, respectively. Based on TSM backbone, our TCP with 8 frame inputs is superior to NeXtVLAD and PPAC with 20 frame inputs by 5.0% and 3.2%, respectively. Besides, TCP outperforms GSM with backbone of Inception-V3 by 0.6%. For Charades dataset, we fine-tune TCPNet with backbone of 3D R101-NL [62] pre-trained on K-400 using learning rate of 0.3 within 50 epochs. The results of different methods in terms of mean Average Precision (mAP) are presented in Table 6, where our TCPNet obtains 42.6% in terms of mAP, outperforming all compared methods. On UCF101 and HMDB51, TCPNet is fine-tuned based on TSN and TEA with backbone of 2D R50 with initial learning rate of 1e-3 within 15 epochs. The results in top-1 accuracy are shown in Table 7, from it we can see that our TCPNet clearly outperforms TSN. Besides, TCPNet based on TSN is superior to other encoding methods (i.e., TLE, TBN, ActionVLAD and PPAC) by a large margin. Lastly, our TCPNet with TEA performs competitively on UCF101 and achieves state-of-the-art results on more challenging HMDB51. These results clearly show our TCP has a good generalization ability.

5 Conclusion

In this paper, we propose a temporal-attentive covariance pooling for generating powerful video representations, which can be seamlessly integrated with existing deep models and result in an effective video recognition architecture (namely TCPNet). Extensive experimental comparisons using various deep video architectures demonstrate the effectiveness of our TCPNet, and show our TCP is a compelling alternative to GAP in popular deep architectures for effective video recognition. By considering complex temporal dynamics and geometry of covariances, our TCP is also superior to its covariance pooling counterparts by a clear margin. In future, we will study more powerful global pooling methods (e.g., third-order pooling or mixture model) for more effective video signal modeling, and apply our proposed TCP to other deep architectures, e.g., vision Transformers [8].

Method	Backbone	Fr	UCF101	HMDB51
TSN [36]	2D R50	8	91.7	64.7
TLE [‡] [7]	2D Inc.	15	95.6	71.1
TBN [35]	2D R34	8	93.6	69.4
AVLAD [‡] [18]	VGG-16	25	92.7	66.9
PPAC [‡] [42]	2D R152	20	94.9	69.8
ECO [73]	Inc.+3D R18	8	91.7	65.6
TSM [36]	2D R50	8	95.9	73.5
TSM [†] [36]	2D R50	8	95.2	72.0
ECO [73]	Inc.+3D R18	16	92.8	68.5
TEA [34]	2D R50	16	96.9	73.3
STM [28]	2D R50	16	96.2	72.2
STC [6]	3D RX101	16	92.3	65.4
ART [59]	3D R18	16	94.3	70.9
I3D [2]	3D Inc.	64	95.4	74.5
ABM [72]	3D Inc.	64	95.1	72.7
TCPNet (Ours)	TSN R50	8	95.1	72.5
TCPNet (Ours)	TEA R50	8	96.4	76.7

Table 7: Results (% in Top-1) on UCF101 and HMDB51. †:Re-implementation. ‡: RGB+Flow. AVLAD: ActionVLAD.

Acknowledgements

Our work was sponsored by National Natural Science Foundation of China (Grant No.s 61971086, 61806140, 61471082, 61872118, 61925602), Natural Science Foundation of Tianjin City (Grant No. 20JCQNJC1530) and CCF- Baidu Open Fund (NO.2021PP15002000).

Appendix

In this section, we evaluate effect of kernel size κ in TC-COV $_{\kappa}$, design of temporal attention module and compare our TCP with plain GCP + MPN in more detail. In Section A1 and A2, the experiments are conducted on Mini-K200 using the same settings as in Section 4.1. The experiments in Section A3 are conducted on K-400 using the same settings as in Section 4.

A1 Effect of Kernel Size κ on TC-COV $_{\kappa}$

Our temporal covariance pooling TC-COV $_{\kappa}$ has a parameter, i.e., kernel size κ . Here we assess effect of kernel size κ on TC-COV $_{\kappa}$. For 8 frames input, kernel size κ varies from 1 to 7 and 16 frames input, κ varies from 1 to 15. As listed in Table A1 & A2, $\kappa = 5$ and $\kappa = 9$ respectively achieves the best results for 8 frames input and 16 frames input. They obtain 0.5% and 0.7% gains over those with $\kappa = 1$, respectively. The performance will decrease gradually when $\kappa > 5$ and $\kappa > 9$, interaction among too large range will bring side effect on performance, while increasing computation cost.

κ	1	3	5	7
Top-1 Acc (%)	75.6	76.0	76.1	76.0
Top-5 Acc (%)	92.5	92.7	92.7	92.6

Table A1: Results (8 frames) of TCPNet with various κ on Mini-K200.

κ	1	3	5	7	9	11	13	15
Top-1 Acc (%)	77.0	77.0	77.4	77.5	77.7	77.3	76.9	76.8
Top-5 Acc (%)	93.0	93.1	93.3	93.3	93.5	93.1	93.0	93.0

Table A2: Results (16 frames) of TCPNet with various κ on Mini-K200.

Channel-att	Spatial-att	TC-COV	Acc (%)	Backbone	GAP	Plain GCP+MPN	TCP
TCA	Non-local [62]	✓	75.6	TSN+R50	70.6	73.1	75.3
SE [24]	TSA	✓	75.6	TSN+R152	75.9	77.2	78.3
TCA	TSA	✓	76.1	X3D-M	76.0	76.5	77.0
Query	Key	Value	Acc (%)	TEA+R50	75.0	75.8	76.8
\mathbf{X}_{l-1}	\mathbf{X}_l	\mathbf{X}_l	75.5				
\mathbf{X}_{l-1}	\mathbf{X}_{l-2}	\mathbf{X}_l	76.1				

Table A3: Evaluation on temporal attention module design on Mini-K200.

Table A4: Comparison with GAP and plain GCP + MPN on K-400.

A2 Design of Temporal Attention Module

Here, we evaluate the effect of our temporal attention module. To this end, we first replace our TSA and TCA by using Non-local and SE blocks, respectively. According to the upper part of Table A3, we can draw the following conclusions: (1) Our TSA achieves 0.5% gains over Non-local in accuracy, while having smaller computational cost (i.e., 0.16 GFLOPs of our TSA vs. 0.7 GFLOPs of Non-local). (2) Our TCA outperforms SE by 0.5% in accuracy. For our TSA module, three consecutive frames are adopted to model spatial attention by exploiting temporal information. Here, we develop a variant of TSA to consider only two consecutive frames, which yields 75.5% in top-1 accuracy and is lower than TSA (76.1%) exploiting three consecutive frames. The results above verify the effectiveness of our temporal attention module.

A3 Comparison TCP with Plain GCP + matrix power normalization (MPN)

In this work, we develop a strong baseline, i.e., Plain GCP + matrix power normalization (MPN). Here, we compare our TCP with this strong baseline on K-400. The results are shown in Table A4, from which we can see that our TCP performs better than Plain GCP + MPN by 2.2%, 1.1%, 0.5% and 1.0% with deep video architectures of TSN + R50, TSN + R152, X3D-M and TEA + R50, respectively. Above results show our TCP can achieve non-trivial performance gains over this strong baseline (Plain GCP + MPN), demonstrating the effectiveness of our TCP in capturing complex temporal dynamics.

References

- [1] Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. Fast and simple calculus on tensors in the Log-Euclidean framework. In *MICCAI*, 2005.
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the Kinetics dataset. In *CVPR*, 2017.
- [3] Lu Chi, Zehuan Yuan, Yadong Mu, and Changhu Wang. Non-local neural networks with grouped bilinear attentional transforms. In *CVPR*, 2020.
- [4] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. MARS: Motion-augmented RGB stream for action recognition. In *CVPR*, 2019.
- [5] Lieven De Lathauwer. *Signal processing based on multilinear algebra*. Katholieke Universiteit Leuven, 1997.
- [6] Ali Diba, Mohsen Fayyaz, Vivek Sharma, M. Mahdi Arzani, Rahman Yousefzadeh, Juergen Gall, and Luc Van Gool. Spatio-temporal channel correlation networks for action classification. In *ECCV*, 2018.
- [7] Ali Diba, Vivek Sharma, and Luc Van Gool. Deep temporal linear encoding networks. In *CVPR*, 2017.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [9] Quanfu Fan, Chun-Fu Richard Chen, Hilde Kuehne, Marco Pistoia, and David Cox. More is less: Learning efficient video representations by big-little network and depthwise temporal aggregation. In *NeurIPS*, 2019.
- [10] Christoph Feichtenhofer. X3D: Expanding architectures for efficient video recognition. In *CVPR*, 2020.
- [11] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019.
- [12] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016.
- [13] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016.
- [14] Shanghua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip HS Torr. Res2Net: A new multi-scale backbone architecture. *IEEE TPAMI*, 43(6):652–662, 2021.
- [15] Yang Gao, Oscar Beijbom, Ning Zhang, and Trevor Darrell. Compact bilinear pooling. In *CVPR*, 2016.
- [16] Zilin Gao, Jiangtao Xie, Qilong Wang, and Peihua Li. Global second-order pooling convolutional networks. In *CVPR*, 2019.
- [17] Rohit Girdhar and Deva Ramanan. Attentional pooling for action recognition. In *NeurIPS*, 2018.
- [18] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. ActionVLAD: Learning spatio-temporal aggregation for action classification. In *CVPR*, 2017.
- [19] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic. The "something something" video database for learning and evaluating visual common sense. In *ICCV*, 2017.
- [20] Mehrtash Tafazzoli Harandi, Mathieu Salzmann, and Richard I. Hartley. Dimensionality reduction on SPD manifolds: The emergence of geometry-aware methods. *IEEE TPAMI*, 40(1):48–62, 2018.
- [21] Bo He, Xitong Yang, Zuxuanu Wu, Hao Chen, Ser-Nam Lim, and Abhinav Shrivastava. GTA: Global temporal attention for video action understanding. *arXiv preprint arXiv:2012.08510*, 2020.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [23] KuehneEmail Hilde, JhuanRainer Hueihan, hagen Stiefel, and Serre Thomas. HMDB: A large video database for human motion recognition. In *ICCV*, 2011.
- [24] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [25] N. Hussein, E. Gavves, and A. W. M. Smeulders. Timeception for complex action recognition. In *CVPR*, 2019.
- [26] Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu. Matrix backpropagation for deep networks with structured layers. In *ICCV*, 2015.
- [27] Hervé Jégou, Florent Perronnin, Matthijs Douze, Jorge Sánchez, Patrick Pérez, and Cordelia Schmid. Aggregating local image descriptors into compact codes. *IEEE TPAMI*, 34(9):1704–1716, 2012.

- [28] Boyuan Jiang, Mengmeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. STM: Spatiotemporal and motion encoding for action recognition. In *ICCV*, 2019.
- [29] Jinyung Kim, Seunghwan Cha, Dongyoon Wee, Soonmin Bae, and Junmo Kim. Regularization on spatio-temporally smoothed feature for action recognition. In *CVPR*, 2020.
- [30] Piotr Koniusz, Lei Wang, and Anoop Cherian. Tensor representations for action recognition. *IEEE TPAMI*, 2020.
- [31] Peihua Li, Jiangtao Xie, Qilong Wang, and Zilin Gao. Towards faster training of global covariance pooling networks by iterative matrix square root normalization. In *CVPR*, 2018.
- [32] Peihua Li, Jiangtao Xie, Qilong Wang, and Wangmeng Zuo. Is second-order information helpful for large-scale visual recognition? In *ICCV*, 2017.
- [33] Xianhang Li, Yali Wang, Zhipeng Zhou, and Yu Qiao. SmallBigNet: Integrating core and contextual views for video classification. In *CVPR*, 2020.
- [34] Yan Li, Bin Ji, Xintian Shi, Jianguo Zhang, Bin Kang, and Limin Wang. TEA: Temporal excitation and aggregation for action recognition. In *CVPR*, 2020.
- [35] Yanghao Li, Sijie Song, Yuqi Li, and Jiaying Liu. Temporal bilinear networks for video action recognition. In *AAAI*, 2018.
- [36] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal shift module for efficient video understanding. In *ICCV*, 2019.
- [37] Rongcheng Lin, Jing Xiao, and Jianping Fan. NeXtVLAD: An efficient neural network to aggregate frame-level features for large-scale video classification. In *ECCV Workshops*, 2018.
- [38] T. Lin, A. RoyChowdhury, and S. Maji. Bilinear CNN models for fine-grained visual recognition. In *ICCV*, 2015.
- [39] Georg Lindgren, Holger Rootzén, and Maria Sandsten. *Stationary stochastic processes for scientists and engineers*. CRC press, 2013.
- [40] Zhaoyang Liu, Donghao Luo, Yabiao Wang, Limin Wang, and Tong Lu. TEINet: Towards an efficient architecture for video recognition. In *AAAI*, 2020.
- [41] Zhaoyang Liu, Limin Wang, Wayne Wu, Chen Qian, and Tong Lu. TAM: Temporal adaptive module for video recognition. In *ICLR*, 2021.
- [42] Xiang Long, Gerard De Melo, Dongliang He, Fu Li, Zhizhen Chi, Shilei Wen, and Chuang Gan. Purely attention based local feature integration for video classification. *IEEE TPAMI*, 2020.
- [43] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A Riemannian framework for tensor computing. *IJCV*, 66(1):41–66, 2006.
- [44] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3D residual networks. In *ICCV*, 2017.
- [45] Jorge Sánchez, Florent Perronnin, Thomas Mensink, and Jakob Verbeek. Image classification with the Fisher Vector: Theory and practice. *IJCV*, 105(3):222–245, 2013.
- [46] Andres Sanin, Conrad Sanderson, Mehrtash Tafazzoli Harandi, and Brian C. Lovell. Spatio-temporal covariance descriptors for action and gesture recognition. In *WACV*, 2013.
- [47] Gunnar Sigurdsson, G Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2018.
- [48] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [49] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [50] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [51] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift networks for video action recognition. In *CVPR*, 2020.
- [52] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift networks for video action recognition. In *CVPR*, 2020.
- [53] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [54] Du Tran, Lubomir D. Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *ICCV*, 2015.
- [55] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018.

- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [57] Heng Wang, Du Tran, Lorenzo Torresani, and Matt Feiszli. Video modeling with correlation networks. In *CVPR*, 2020.
- [58] Lei Wang, Piotr Koniusz, and Du Q. Huynh. Hallucinating IDT descriptors and I3D optical flow features for action recognition with CNNs. In *ICCV*, 2019.
- [59] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification. In *CVPR*, 2018.
- [60] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE TPAMI*, 41(11):2740–2755, 2019.
- [61] Qilong Wang, Jiangtao Xie, Wangmeng Zuo, Lei Zhang, and Peihua Li. Deep CNNs meet global covariance pooling: Better representation and generalization. *IEEE TPAMI*, 2020.
- [62] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [63] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018.
- [64] Thomas Winterbottom, Sarah Xiao, Alistair McLean, and Noura Al Moubayed. Trying bilinear pooling in Video-QA. *arXiv preprint arXiv:2012.10285*, 2020.
- [65] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krähenbühl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019.
- [66] Chao-Yuan Wu, Ross Girshick, Kaiming He, Christoph Feichtenhofer, and Philipp Krahenbuhl. A multigrid method for efficiently training video models. In *CVPR*, 2020.
- [67] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018.
- [68] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *CVPR*, 2020.
- [69] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *ICCV*, 2017.
- [70] Kaiyu Yue, Sun Ming, Yuchen Yuan, Feng Zhou, Errui Ding, and Fuxin Xu. Compact generalized non-local network. In *NeurIPS*, 2018.
- [71] Shiwen Zhang, Sheng Guo, Weilin Huang, Matthew R Scott, and Limin Wang. V4D: 4D convolutional neural networks for video-level representation learning. In *ICLR*, 2020.
- [72] Xinqi Zhu, Chang Xu, Langwen Hui, Cewu Lu, and Dacheng Tao. Approximated bilinear modules for temporal modeling. In *ICCV*, 2019.
- [73] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. ECO: Efficient convolutional network for online video understanding. In *ECCV*, 2018.