

Modality Compensation Network: Cross-Modal Adaptation for Action Recognition

Sijie Song¹, *Student Member, IEEE*, Jiaying Liu¹, *Senior Member, IEEE*,
Yanghai Li¹, *Student Member, IEEE*, and Zongming Guo¹, *Member, IEEE*

Abstract—With the prevalence of RGB-D cameras, multi-modal video data have become more available for human action recognition. One main challenge for this task lies in how to effectively leverage their complementary information. In this work, we propose a Modality Compensation Network (MCN) to explore the relationships of different modalities, and boost the representations for human action recognition. We regard RGB/optical flow videos as *source* modalities, skeletons as *auxiliary* modality. Our goal is to extract more discriminative features from source modalities, with the help of auxiliary modality. Built on deep Convolutional Neural Networks (CNN) and Long Short Term Memory (LSTM) networks, our model bridges data from source and auxiliary modalities by a modality adaptation block to achieve adaptive representation learning, that the network learns to compensate for the loss of skeletons at test time and even at training time. We explore multiple adaptation schemes to narrow the distance between source and auxiliary modal distributions from different levels, according to the alignment of source and auxiliary data in training. In addition, skeletons are only required in the training phase. Our model is able to improve the recognition performance with source data when testing. Experimental results reveal that MCN outperforms state-of-the-art approaches on four widely-used action recognition benchmarks.

Index Terms—Modality compensation, multi-modal, action recognition.

I. INTRODUCTION

RECOGNITION of human actions is quite an important yet fundamental task in video analytics. Human action recognition technology is required for a wide range of applications, such as intelligent video surveillance, human-computer interaction and robotics. However, action recognition remains a challenging problem due to heterogeneous backgrounds, subtle inter-class differentiation and large intra-class variations. The key to the success of this task is enhancing the feature discriminations of various actions.

Manuscript received June 13, 2018; revised January 2, 2019 and August 25, 2019; accepted January 12, 2020. Date of publication January 23, 2020; date of current version February 4, 2020. This work was supported in part by the National Natural Science Foundation of China under Contract 61772043, in part by the Beijing Natural Science Foundation under Contract 4192025, in part by the Microsoft Research Asia (FY19-Research-Sponsorship-115), and in part by the Peking University Tencent Rhino Bird Innovation Fund. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Senem Velipasalar. (*Corresponding author: Jiaying Liu.*)

The authors are with the Wangxuan Institute of Computer Technology, Peking University, Beijing 100080, China (e-mail: ssj940920@pku.edu.cn; liujiaying@pku.edu.cn; lyttonhao@pku.edu.cn; guozongming@pku.edu.cn).

Digital Object Identifier 10.1109/TIP.2020.2967577

In the past decades, much work has been extensively studied on human action recognition in RGB videos [1]. Besides the methods relying on hand-crafted features [2], the two-stream architecture based on deep networks is very popular and successful in utilizing visual and temporal cues [3]. The spatial and temporal streams are trained independently and the final results are generated from the combination of each stream. Due to the strong performance of the two-stream architecture in action recognition, many works [4]–[6] following this line focus more on improving feature discriminations, and modelling the correlations between the spatial and temporal structures. Our starting point is the notice of the following difficulties in the task of action recognition based on the two-stream architecture:

- With limited training samples, it is hard for the network to extract invariant features and cover all kinds of possible action variations, such as the differences in viewpoints and scales.
- For the spatial stream, the extracted features can be disturbed by cluttered backgrounds, which leads to the mixture of human shapes and backgrounds.
- For the temporal stream, the background motions will introduce noises in the temporal features and disturb the learning process.

To address the above issues, we are inspired to seek other cues for more effective and robust feature learning. With the recent advent of low-cost depth cameras like Microsoft Kinect [7] and the maturity of pose estimation technologies [8], there is an increasing amount of visual data containing both color videos and 3D skeletons. These modalities provide multiple cues that are complimentary, which inspires us to explore their underlying common feature spaces. For example, the skeletons are invariant to different viewpoints and backgrounds. By exploring their common feature spaces, we expect the network to learn invariant and robust features from RGB/optical flow data.

As a high-level representation of human motion, biological observations have suggested that humans can recognize actions from skeleton data [9]. Skeleton-based representation is robust to the variations of illumination and scales, and 3D skeletons are even invariant to different viewpoints. Besides, it is able to exclude backgrounds, provide clear motion patterns and human posture descriptions. These advantages make skeleton data as an attractive option for action recognition [10]–[13]. Nevertheless, the lack of appearance information results in

the ambiguity for some actions whenever only observing skeletons.

To overcome the lack of appearance for skeleton data, many works integrate skeletons with other modalities to utilize complementary information. Most previous works focus on feature fusion, which is borrowed from the two-stream architecture. A recent work developed a chained multi-stream approach [14] built on Markov chain model to combine appearance, motion and pose features. Other works tend to explore the common high-level feature spaces from different modalities through cross-modal learning [15], [16]. The work in [16] is one of the early attempts. However, it only regularizes the feature learning of the RGB stream with skeletons, without taking optical flow into consideration. It is still underexplored that how to use skeletons to compensate for features from optical flow data. Besides, the modality adaptation under different scenarios where the modalities are aligned or misaligned should be explored.

In this work, we demonstrate that the skeleton-based representation is a valuable supplement to color videos in action recognition. We focus on adapting skeletons into modal spaces learnt from RGB/optical flow data. We define RGB/optical flow data as source data and skeletons as auxiliary data. Our key idea is to borrow the advantages of auxiliary skeletons and learn the compensated representations from source data. For the spatial stream, guided by skeletons, the network is encouraged to extract complementary features reflecting motion and detailed appearance contexts from source data. For the temporal stream, skeletons are used to supplement the shape reference, and facilitate the motion feature learning to exclude the motion noises in the backgrounds. Our work needs auxiliary data provided in training, but only requires source data in the testing phase. The main contributions of this paper can be summarized as follows:

- We propose a novel Modality Compensation Network for action recognition with adaptive representation learning, which aims to align the distributions of source and auxiliary data. Our model is able to extract compensated and more discriminative features from source data for action recognition.
- A modality adaptation block with residual feature learning is developed to bridge data from source and auxiliary modalities. We show that the residual structure is more effective in borrowing the advantages of auxiliary data.
- We explore different levels of modality adaptation schemes, including domain-, category- and sample-level, to cope with different scenarios according to the alignment of color and skeleton videos in the training data.
- We give comprehensive analysis on each component in our model to better understand the modality adaptation. Our approach is evaluated on the NTU RGB+D dataset, the MSR 3D Daily Activity dataset, the UCF-101 dataset and the JHMDB dataset, respectively. Experimental results show that our proposed model outperforms other state-of-the-art methods, thanks to more effective spatial and temporal feature learning.

The remainder of the paper is organized as follows. In Section II, we review the related works on action

recognition based on RGB videos or skeletons, and multi-modal feature learning. In Section III, we first give an overview of our baseline structure, and then introduce the proposed Modality Compensation Network in detail. Experiments and analysis are presented in Section IV. Finally, we conclude our work in Section V.

II. RELATED WORK

A. RGB-Based Action Recognition

Action recognition on RGB-based videos has attracted many research interests and been extensively studied in the past decades. For action recognition with RGB as inputs, a key branch is based on hand-crafted descriptors with decoupled classifiers [2], [17]–[19], other methods try to jointly learn features and classifiers. With the advent of deep learning, neural networks are recently employed for action recognition due to its powerful ability in learning robust feature representations [3]–[6], [20]–[27]. The two-stream architecture [3] is a pioneer work to employ deep convolutional network for action recognition in videos, and has become a backbone of many other approaches [4]–[6], [25], [27]–[29]. To address the aggregation of spatial and temporal features, [6] explored different score fusion schemes. The work in [5], [27] utilized residual connections to allow spatiotemporal interaction between to streams. Reference [22] used 3D convolutional networks (C3D) to learn discriminative spatio-temporal patterns jointly. Based on C3D, VLAD3 [30] was developed to model long-range dynamic information. Besides, the works in [23], [24] performed a hierarchical rank pooling to obtain video representations, which have high capacity of capturing informative frame-based feature representations. To further extend the temporal support, great efforts have been made with recurrent neural networks, learning how to integrate dynamics overtime [21], [31]–[33]. Towards good practice, [4] presents a Temporal Segment Network (TSN) and provides an effective way to model long-term temporal structure, which has brought the state-of-the-art to a new stage. However, the diversities in viewpoints and backgrounds make it challenging to extract discriminative features from RGB videos. In our work, we use skeleton data, which are invariant in backgrounds and scales, to encourage the network to learn more robust features from color videos.

B. Skeleton-Based Action Recognition

Skeleton-based action recognition can be regarded as a recognition task on time series for exploring spatio-temporal patterns. Many traditional works focus on hand-crafted features, which are generally based on the geometry relationships and high order encodings [34]–[38]. Benefiting from the merits of recurrent neural network for sequential data, some works adopt recurrent neural network to explore the spatial and temporal dynamics of skeletal data [10], [11], [13], [39]–[42]. With the help of convolutional neural networks, Ke *et al.* [43] proposed a new representation for 3D skeleton data, which transfers action recognition problem to the problem of image classification.

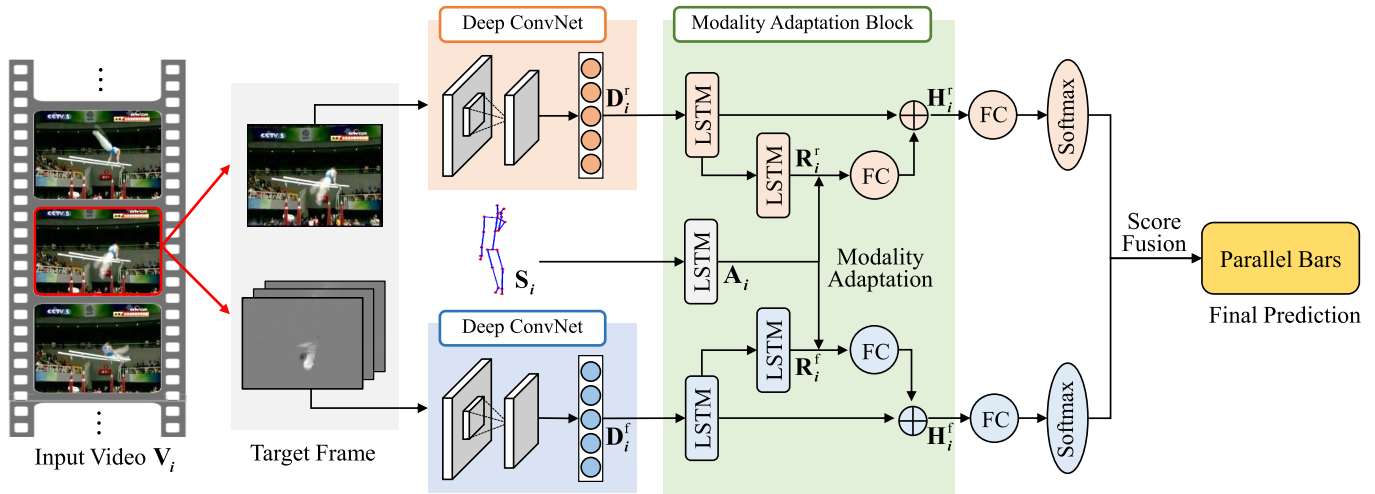


Fig. 1. The framework of Modality Compensation Network. Based on the two-stream architecture, a modality adaptation block is incorporated to compensate the feature learning of source data with the help of auxiliary data. (We use superscripts ‘r’ and ‘f’ to indicate RGB and flow streams, respectively, which are omitted in the text for simplicity.) Note that the LSTM with output A_i is not included in the testing phase.

Due to lack of appearance for skeleton data, many works integrate skeletons with data from other modalities to utilize the complementary information. A regularized LSTM is proposed [16] to constrain the RGB feature space with skeletons. But the relationship between skeletons and optical flow data is overlooked in [16]. Shi and Tim [15] proposed to achieve action recognition from depth sequences by learning an RNN with privileged information from skeletons. In addition, a chained multi-stream network [14] built on Markov chain model is developed to integrate appearance, motion and pose features. Liu *et al.* [44] jointly learned the regression and classification network with multi-modal data for action detection. Though the methods in [14], [15], [44], [45] illustrate that the introduction of multiple cues improve the performance for action analytics, they are limited by the strict data requirements. Reference [15] requires aligned skeletons and depth data in training. The models in [14] and [44] need aligned color videos and 3D skeletons (or 2D poses) both in training and testing. However, aligned multi-modal data are not always reliable. It is desirable to utilize unaligned data to facilitate feature learning. In our work, we explore the relationship not only between RGB videos and skeletons, but also between optical flow and skeleton data. Besides, our model is flexible in that it can be trained with aligned or unaligned multi-modal data. And only source modalities (*i.e.*, RGB and optical flow) are required at the test time.

C. Multi-Modal Feature Learning

Inspired by that different modalities provide complementary information, cross-modal feature learning has been widely studied these days. Much work tends to unify the distributions of different feature spaces, by regularizing the representation learning with data from other modalities [46], [47]. This can be regarded as a domain adaptation problem, which aims to minimize the distribution between source and target domain [48], [49]. Long *et al.* [50] introduced a deep adaptation network

to minimize the maximum mean discrepancy of the feature to learn transferable features in high layers of the network. However, the work in [51] claims that focusing only on the shared representation leads to the ignorance of individual characteristics. Following this line, Wang *et al.* [46] proposed a multi-modal feature learning framework for RGB-D object recognition to learn not only modal-specific patterns but also modal-shared features. The similar idea is also employed in the task of image segmentation [47], in which RGB and depth data sources are correlated with multi-kernel maximum mean discrepancy to discover common and specific features.

III. MODALITY COMPENSATION NETWORK

We propose a modality compensation network for action recognition. Our network is designed to implicitly compensate features from source modalities with the help of auxiliary data. We propose a modality adaptation block to improve the ability of adaptation learning with skeletons for action recognition. A hierarchical modality adaptation scheme is further explored for different data alignment cases to borrow information from the auxiliary data more effectively. Fig. 1 shows the overall architecture of our network. In the following, we first introduce our baseline model, then we discuss the modality adaptation block and our hierarchical modality adaptation scheme.

A. Baseline Model

Our backbone is built on the two-stream structure [3] as shown in Fig. 1: the spatial stream leverages the appearance information extracted from the RGB frames, while the temporal stream utilizes the motion contexts from the stacked horizontal and vertical optical flow data. The structure of each stream mainly consists of a CNN and an LSTM. The CNN is employed as an encoder to extract spatial features from each input frame, and the information over time is collected and integrated by the LSTM. Note that each stream is trained independently and the final results for two streams are given

by score fusion. Note, we use superscripts ‘r’ and ‘f’ in Fig. 1 to indicate RGB and optical flow streams, respectively. The following illustrations for our proposed model are in terms of a single stream and we omit the superscripts for simplicity.

1) *Convolutional Neural Network*: Compared with hand-crafted features, convolutional networks have shown superiority of interpreting and summarizing images. Given the i -th video sequence $\mathbf{V}_i = \{\mathbf{v}_{i,t} : t = 1, \dots, T\}$ within a batch, where $\mathbf{v}_{i,t}$ is the frame at time step t , we extract spatial features from the convolutional layers. Inspired by the pioneer work [4], in experiments we employ features after *global-pool* layer from the Inception with Batch Normalization (BN-Inception) network [52]. Then each frame is represented by a vector $\mathbf{d}_{i,t} \in \mathbb{R}^{1024}$. We then feed the sequence of CNN features $\mathbf{D}_i = \{\mathbf{d}_{i,t} : t = 1, \dots, T\}$ to the following recurrent layers.

2) *Long Short Term Memory Network*: To exploit the temporal dependencies in video sequences, LSTM [53] is employed to build our network. The self-connected architecture allows LSTM to maintain and process temporal information over time. At each time step, the network can choose to read, write or reset the memory cell governed by the input gate, forget gate and output gate.

3) *Loss Function*: To obtain the class predictions of \mathbf{V}_i , the hidden states $\mathbf{H}_i = [\mathbf{h}_{i,1}, \dots, \mathbf{h}_{i,T}]$ from LSTM layers is accumulated and then mapped to a 1-of- C encoding vector $\mathbf{g}_i = [g_{i,1}, \dots, g_{i,C}]$, representing the confidence scores of C classes of actions

$$\mathbf{g}_i = \mathbf{W}_g \left(\frac{1}{T} \sum_{t=1}^T \mathbf{h}_{i,t} \right) + \mathbf{b}_g, \quad (1)$$

where \mathbf{W}_g and \mathbf{b}_g are the parameters of the last FC layer before *softmax*. The predicted probability of the video \mathbf{V}_i being the c -th class is then normalized to $p(c|\mathbf{V}_i) = e^{g_{i,c}} / \sum_{j=1}^C e^{g_{i,j}}$ ($c = 1, \dots, C$). Our goal for each stream is to minimize the objective function as

$$\mathcal{L} = - \sum_{i=1}^n \sum_{c=1}^C l_{i,c} \log p(c|\mathbf{V}_i), \quad (2)$$

where $l_{i,c} = 1$ if the i -th video sequence belongs to the c -th class, and 0 otherwise.

For the final prediction, we obtain the probability for being the c -th class from each stream by score fusion, which can be formulated as

$$p^*(c|\mathbf{V}_i) = \sum_{m \in \{r,f\}} w^m p^m(c|\mathbf{V}_i), \quad (3)$$

where the superscript m indicates which stream the score is from. w^m denotes the fusion weight, and $\sum_{m \in \{r,f\}} w^m = 1$. In our experiments, we set $w_m = 0.5$ in two-stream fusion.

B. Adaptive Representation Learning

In the baseline model, the inter- and intra-class variations result in the difficulties in extracting robust features to adapt to different viewpoints and scales. Besides, the cluttered backgrounds can easily introduce noises in the process of feature

learning. In this work, we propose an adaptive representation learning method, for the purpose of compensating features and exploring more discriminative representations from the input frames. Thanks to the skeleton data, as a high level representation invariant to viewpoints and backgrounds, we are able to improve the feature learning of source data.

Modality Adaptation Block We design a modality adaptation block to achieve adaptive representation learning. For each stream, it consists of a main network and a residual subnetwork [54] using recurrent layers. The idea of residual LSTM networks has been explored in [55] and [56], both of which demonstrates the effectiveness of the structure. Our network is more similar to the staked LSTM [56]. Fig. 1 shows the detail structure of our network. On one hand, the network is able to keep the original information from source modalities. On the other hand, the compensated features are attained from the residual block through adaptive representation learning. The basic formulation of the residual subnetwork is defined as

$$\mathbf{z}^{l+1} = \mathcal{H}(\mathbf{z}^l, \{\theta^l\}) + \mathbf{z}^l, \quad (4)$$

where \mathbf{z}^l and \mathbf{z}^{l+1} are the input and output vectors of the layer considered, and $\mathcal{H}(\cdot)$ is a nonlinear residual mapping with parameters denoted as θ^l . More specifically, we use an LSTM layer and an FC layer as the residual mapping function in the residual subnetwork, as shown in Fig. 1.

The insight of our adaptive representation learning is to transform the source modal space to the auxiliary modal space, so that the advantages in the auxiliary modality can also be represented using data from source modalities, and then to achieve modality compensation. Note that the modality compensation is mainly achieved by the residual path, which pushes the feature representation of source modalities be close to the auxiliary skeleton features. For the features which are present in the source modalities but absent in the auxiliary modality, we use the skip connection in the residual block to preserve them. For a given skeleton sequence $\mathbf{S}_i = \{\mathbf{s}_{i,t} : t = 1, \dots, T\}$, we encode it into feature vectors $\mathbf{A}_i = \{\mathbf{a}_{i,t} : t = 1, \dots, T\}$ by a pre-trained LSTM network with parameters $\{\theta_a\}$, which can be formulated as $\mathbf{A}_i = f_a(\mathbf{S}_i, \{\theta_a\})$. The features from the source modalities are denoted as $\mathbf{R}_i = \{\mathbf{r}_{i,t} : t = 1, \dots, T\}$, which are extracted from the output of the LSTM layer in the residual subnetwork as shown in Fig. 1. We formulate it as $\mathbf{R}_i = f_r(\mathbf{V}_i, \{\theta_r\})$, where $\{\theta_r\}$ denotes the parameters. With fixed $\{\theta_a\}$, our goal is to learn features from source modalities similar to the pre-defined auxiliary features through optimizing the feature learning with $\{\theta_r\}$ in source data. The optimization can be achieved by minimizing the distance between the source modal space $\bar{\mathbf{R}}$ and the auxiliary modal space $\bar{\mathbf{A}}$ as follows

$$\text{dist}(\bar{\mathbf{R}}, \bar{\mathbf{A}}). \quad (5)$$

And it can be integrated with the loss function defined in Eq. (2)

$$\mathcal{L} = - \sum_{i=1}^n \sum_{c=1}^C l_{i,c} \log p(c|\mathbf{V}_i) + \lambda d, \quad (6)$$

TABLE I
REQUIREMENTS FOR MULTI-MODAL DATA OF EACH SCHEME

Level	Required Alignment	Granularity	Target Case
Domain	Not required	General	General & Specific
Category	Category-aligned	Mid-level	Specific
Sample	Counterparts	Specific	Specific

where $d = \text{dist}(\bar{\mathbf{R}}, \bar{\mathbf{A}})$, and we use λ to balance the contribution of the adaptive term. However, due to the huge gap between skeletons and color videos, it is hard to define the distance function $\text{dist}(\cdot)$ directly. Instead, we bridge the different modalities by adapting their contributions from different levels, namely, domain-level adaptation (d_D), category-level adaptation (d_C), and sample-level adaptation (d_S). Thus we have $d \in \{d_D, d_C, d_S\}$ which will be discussed in the next subsection.

C. Different Levels of Modality Adaptation

To calculate the distance of source and auxiliary modal spaces and achieve modality adaptation, we consider two cases here:

- General case: the modalities in the training data are not aligned individually.
- Specific case: each color video in the training data is provided with its corresponding skeleton sequence.

To deal with the different cases, we consider different levels of modality adaptation, including domain-, category- and sample-level adaptation. We summarize the requirements for multi-modal data of each adaptation scheme in Table I.

1) *Domain-Level Adaptation*: Domain-level adaptation is able to deal with the most general cases, regardless of the alignment of source and auxiliary data. Inspired by domain adaptation, we utilize Maximum Mean Discrepancy (MMD) [57] to align domain-level distributions and handle this modality adaptation problem. Given two sets of samples, $X = \{\mathbf{x}_i\}_{i=1}^{m_x}$ and $Y = \{\mathbf{y}_j\}_{j=1}^{m_y}$ generated from two distributions, respectively, the squared MMD calculates the norm of the difference between embeddings of the two different distributions as

$$\begin{aligned}
 \text{MMD}^2[X, Y] &= \|\mathbf{E}_x[\phi(\mathbf{x})] - \mathbf{E}_y[\phi(\mathbf{y})]\|^2 \\
 &= \frac{1}{m_x^2} \sum_{i=1}^{m_x} \sum_{i'=1}^{m_x} \phi(\mathbf{x}_i)^T \phi(\mathbf{x}_{i'}) + \frac{1}{m_y^2} \sum_{j=1}^{m_y} \sum_{j'=1}^{m_y} \phi(\mathbf{y}_j)^T \phi(\mathbf{y}_{j'}) \\
 &\quad - \frac{2}{m_x m_y} \sum_{i=1}^{m_x} \sum_{j=1}^{m_y} \phi(\mathbf{x}_i)^T \phi(\mathbf{y}_j), \quad (7)
 \end{aligned}$$

where ϕ is the explicit feature mapping function of MMD, m_x and m_y indicate the number of samples under the two different distributions, respectively.

For the assessment of the distance between source and auxiliary modal distributions, we consider the video-level representations to compute the MMD distance. Given the encoded features \mathbf{R}_i from video \mathbf{V}_i , and \mathbf{A}_i from skeleton sequence \mathbf{S}_i , respectively, we generate the video-level feature descriptors

by $\hat{\mathbf{r}}_i = \frac{1}{T} \sum_{t=1}^T \mathbf{r}_{i,t}$ for source data, and $\hat{\mathbf{a}}_i = \frac{1}{T} \sum_{t=1}^T \mathbf{a}_{i,t}$ for auxiliary data. Applying the associated kernel function $k(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle$ in Eq. (7), the MMD distance between the source and auxiliary data can be estimated as

$$\begin{aligned}
 d_D &= \frac{1}{n^2} \sum_{i=1}^n \sum_{i'=1}^n k(\hat{\mathbf{a}}_i, \hat{\mathbf{a}}_{i'}) + \frac{1}{n^2} \sum_{j=1}^n \sum_{j'=1}^n k(\hat{\mathbf{r}}_j, \hat{\mathbf{r}}_{j'}) \\
 &\quad - \frac{2}{n^2} \sum_{i=1}^n \sum_{j=1}^n k(\hat{\mathbf{a}}_i, \hat{\mathbf{r}}_j), \quad (8)
 \end{aligned}$$

where n is the batch size, and we employ the linear kernel $k(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y}$ in our work. In practice, for misaligned source and auxiliary data (*i.e.*, RGB and skeletons from different video sources in training), Eq. (8) is optimized within a mini-batch, to explore the underlying common spaces of two misaligned modalities.

2) *Category-Level Adaptation*: Taking the category information into account, the network is able to learn more specific features guided by the skeleton data from the same category, so it is easier to find the optimal. In our category-based MMD, we only consider the distributions of the source and auxiliary data sharing the same label. More specifically, with the video-level feature descriptors and a linear kernel, we use the following function to calculate category-based MMD within a mini batch with the size of n ,

$$\begin{aligned}
 d_C &= \frac{1}{C} \sum_{c=1}^C \left(\frac{1}{n_c^2} \sum_{i=1}^{n_c} \sum_{i'=1}^{n_c} k(\hat{\mathbf{a}}_{i,c}, \hat{\mathbf{a}}_{i',c}) \right. \\
 &\quad \left. + \frac{1}{m_c^2} \sum_{j=1}^{m_c} \sum_{j'=1}^{m_c} k(\hat{\mathbf{r}}_{j,c}, \hat{\mathbf{r}}_{j',c}) \right. \\
 &\quad \left. - \frac{2}{n_c m_c} \sum_{i=1}^{n_c} \sum_{j=1}^{m_c} k(\hat{\mathbf{a}}_{i,c}, \hat{\mathbf{r}}_{j,c}) \right), \quad (9)
 \end{aligned}$$

where there are C classes of actions in total, n_c and m_c denote the number of auxiliary and source modal samples of the c -th class within the batch (*i.e.*, $\sum_{c=1}^C n_c = \sum_{c=1}^C m_c = n$), $\hat{\mathbf{a}}_{i,c}$ and $\hat{\mathbf{r}}_{j,c}$ are the auxiliary and source modal features of the c -th class, respectively.

3) *Sample-Level Adaptation*: This scenario assumes that, each color video of the source modality is provided with its corresponding skeleton sequence. Intuitively, we can adopt the sample-level adaptation on the granularity of each frame or each video. However, despite that each input frame is corresponding to a skeleton, optical flow images could be flat whenever actions only involve subtle actions. Therefore, it is not feasible to narrow the frame-wise distance between skeleton descriptions and features from flat optical flow images. Meanwhile, no gain is observed when narrowing frame-wise distance between RGB images and skeletons. It is probably because the noises in skeleton data affect the feature learning from RGB images. Instead, we use the video-level descriptions mentioned in Eq. (8). Euclidean distance is employed to learn adaptive features as follows

$$d_S = \frac{1}{n} \sum_{i=1}^n \|\hat{\mathbf{a}}_i - \hat{\mathbf{r}}_i\|^2, \quad (10)$$

where $\hat{\mathbf{a}}_i$ and $\hat{\mathbf{r}}_i$ are the video-level descriptors of the i -th sample from the auxiliary and source modalities, respectively, and n is the batch size.

IV. EXPERIMENT RESULTS

For evaluation, we conduct our experiments on the following four datasets: the NTU RGB+D dataset [12], the MSR 3D Daily Activity Dataset [36], the UCF-101 dataset [58] and the JHMDB dataset [59]. To further explore the effectiveness of each component in our work, the ablation analysis is given. Note that our model needs data from auxiliary modality to enable better feature learning in training but only requires data from source modalities when testing.

A. Datasets and Settings

1) *NTU RGB+D Dataset (NTU)*: [12] is the largest action recognition dataset with individually aligned multi-modal data (*i.e.*, skeletons, RGB, depth and infrared). This dataset consists of 56880 video samples with more than 4 million frames. There are 60 action types performed by 40 subjects, including interactions with pairs and individual activities. Each skeleton has 25 joints. The cross subject (CS) and cross view (CV) settings are two protocols to evaluate the performance. Because the resolution of RGB frames is 1920×1080 , after resized to 224×224 , the actors only hold a small place in the scene. To avoid unnecessary degradation in performance, the original RGB images are cropped to increase the resolution of subjects. To crop RGB images, after mapping skeletons to RGB, the min and max of the joint coordinates on the RGB video ($x_{min}, x_{max}, y_{min}, y_{max}$) can be calculated. Then we crop the region of $[x_{min} - 250 : x_{max} + 250, y_{min} - 50 : y_{max} + 50]$, increasing the resolution of actors. We also perform data normalization [12] on skeletons to have position and view invariance. To accelerate the extraction of optical flow [4], the RGB and skeleton sequences are downsampled with a stride as 5 over the temporal axis.

2) *MSR 3D Daily Activity Dataset (MSR)*: [36] includes 16 action categories related to daily activities, such as *playing the guitar, using vacuum cleaner*. This dataset provides skeletons (20 joints for each skeleton), RGB and depth images. The total number of video sequences is 320 with image resolution of 640×480 . The first five actors are used for training and others for testing, which follows the cross-subject setting.

3) *UCF-101 Dataset (UCF-101)*: [58] consists of 101 action classes, over 13k fully-annotated action videos. Each video snippet lasts 3-10 seconds and contains 100-300 frames, with a fixed resolution of 320×240 . This dataset is challenging due to the large variations in action categories, viewpoints and backgrounds. We follow the evaluation in [4] and report the mean average accuracy over the three training/testing splits. The videos are downsampled and processed in sequences of 60 frames.

4) *JHMDB Dataset (JHMDB)*: [59] is composed of 928 RGB videos for 21 actions such as *brushing hair, running, shooting ball*. Each clip contains 15-40 frames of size 320×240 and there are 31838 frames in total. The dataset is

TABLE II
SETTINGS FOR EACH DATASET

	NTU		MSR	UCF-101		JHMDB	
Auxiliary data	NTU	PKU	MSR	NTU	Poses	NTU	Poses
N	1024	1024	100	1024	1024	100	100
batch size	256	256	16	192	192	16	16

divided into three train/test splits, and evaluation averages the results over the three splits.

5) *Auxiliary Data Sources*: The auxiliary data may come from internal or external sources. We consider different auxiliary data sources for each dataset. For the NTU dataset, we consider two settings here: (1) Using its own skeletons as the auxiliary data, since the NTU dataset is provided with aligned skeleton sequences and RGB videos. (2) Using external skeletons from PKU Multi-Modality Dataset (PKU) [60]. There are 52 action labels in the PKU dataset, and all of them are from the NTU dataset. Thus, the two datasets are category-aligned and we are able to further evaluate the effectiveness of domain-level and category-level adaptation schemes. For the MSR dataset, there are aligned skeleton sequences and RGB videos, and its own skeletons are used as auxiliary data. However, for the UCF-101 and JHMDB datasets, there are no counterpart skeletons available. We also consider two settings here to compensate the loss of skeletons: (1) Using external 3D skeletons from the NTU (CV) dataset as their auxiliary data. (2) Using 2D poses extracted from each video with an off-the-shelf pose extractor (*i.e.*, OpenPose [61]) as their internal auxiliary data. The summary for auxiliary data sources can be found in Table II.

6) *Implementation Details*: We list the parameter settings for each dataset in Table II. For the auxiliary data, we encode the skeleton representations using a 1-layer LSTM with N units, which is pretrained by the classification task. For the NTU and MSR datasets, we use skeletons from the corresponding training set to pre-train the auxiliary network. For example, for NTU dataset, when under the setting of cross view (CV), the auxiliary network is trained with skeletons from the training set of NTU-CV. While under the setting of cross subject (CS), we use skeletons from the training set of NTU-CS to pre-train the auxiliary network. For UCF-101 and JHMDB dataset using external sources as auxiliary data, we simply use skeletons from the training set of NTU-CV to achieve pretraining. And when using extracted poses as auxiliary data, we use the poses from corresponding training set to pre-train the auxiliary network. For the source data (RGB and optical flow), we pretrain the BN-Inception network in [4] on the four datasets, respectively. The features from the layer of *global pool* in BN-Inception are extracted into a 1024-dimensional vector and then fed into the recurrent layers. We use a 2-layer LSTM network with N units in each LSTM layer, including one in the residual subnetwork, as shown in Fig. 2(a). The number of units in the FC layer after the LSTM is also N . In practice, we fix the parameters for encoding skeleton representations and train the

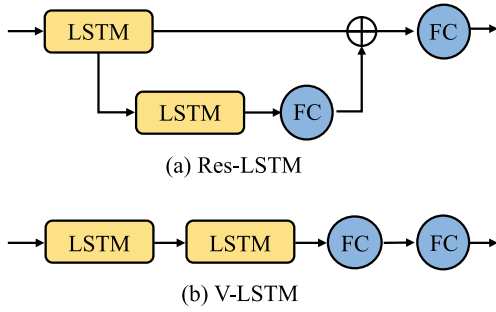


Fig. 2. The structures of Res-LSTM and V-LSTM, respectively. The input features are from CNN (omitted in the figure).

TABLE III

PERFORMANCE EVALUATION ON THE NTU AND MSR DATASETS IN ACCURACY (%). THE SUPERScript* DENOTES THE AUXILIARY DATA ARE FROM EXTERNAL SOURCES. FOR THE NTU DATASET, THE EXTERNAL SOURCE IS THE PKU DATASET

Methods	NTU-CS		NTU-CV		MSR	
	RGB	Flow	RGB	Flow	RGB	Flow
Res-LSTM	79.6	85.8	85.7	91.8	67.5	71.3
D-Res-LSTM*	80.9	86.6	85.1	91.9	–	–
C-Res-LSTM*	81.4	86.8	85.2	92.1	–	–
D-Res-LSTM	81.4	86.4	87.3	92.0	72.5	72.5
C-Res-LSTM	81.5	86.6	87.2	92.1	71.9	75.0
S-Res-LSTM	82.0	87.6	89.1	93.3	74.4	76.9
J-Res-LSTM	81.8	87.3	89.1	93.8	73.8	75.0

network for RGB/optical flow streams. Adam optimizer [62] is adopted to automatically modulate the learning rate with initial value as 0.001. Dropout [63] with a probability of 0.5 is used to mitigate overfitting. For the selection of the hyper parameter λ in Eq. (6), we approximate the weights per the order of magnitude for each adaptation scheme. Then we try different weights with a factor of 10 for the regularizations and choose the best weight per the performance on the validation set. Section IV-E further discusses the effect of λ .

B. Effectiveness of Modality Adaptation

To validate the effectiveness of modality adaptation, we compare the proposed adaptation schemes on the four datasets, respectively. We conduct experiments with the configurations targeting different cases in Table I as follows:

- **Res-LSTM**: Our baseline without using modality adaptation, as shown in Fig. 2(a).
- **D-Res-LSTM**: Res-LSTM with domain-level adaptation, targeting the general and specific cases.
- **C-Res-LSTM**: Res-LSTM with category-level adaptation, targeting the specific case.
- **S-Res-LSTM**: Res-LSTM with sample-level adaptation, targeting the specific case.
- **J-Res-LSTM**: Res-LSTM with domain-, category-, and sample-level adaptations jointly, targeting the specific case.

Table III shows the results for the NTU and MSR datasets on the modalities of RGB (RGB) and optical flow (Flow),

TABLE IV

PERFORMANCE EVALUATION WITH DIFFERENT AUXILIARY DATA SOURCES IN TRAINING ON THE UCF-101 (SPLIT1) AND JHMDB (SPLIT1) DATASETS IN ACCURACY (%). THE SUPERScript* DENOTES THE AUXILIARY DATA ARE FROM EXTERNAL SOURCES (*i.e.*, THE NTU (CV) DATASET)

Methods	UCF-101		JHMDB	
	RGB	Flow	RGB	Flow
Res-LSTM	81.6	83.8	50.9	61.5
D-Res-LSTM*	83.3	85.7	53.2	62.6
D-Res-LSTM	79.5	84.0	52.7	62.9
C-Res-LSTM	78.4	84.7	54.9	63.6
S-Res-LSTM	81.9	85.3	57.9	64.8
J-Res-LSTM	80.1	83.3	52.7	63.6

respectively. For the NTU dataset, when using external auxiliary skeletons from the PKU dataset, we are able to apply domain-level and category-level adaptation schemes since they are category-aligned. Results show that both the adaptation schemes (D-Res-LSTM* and C-Res-LSTM*) can improve the action recognition performance. Furthermore, the category-level adaptation outperforms domain-level adaptation. When using internal auxiliary data, it can be regarded as the specific case. We evaluate the domain-, category- and sample-level adaptation schemes, respectively. It is observed that the adaptive representation learning effectively boosts the results under all settings with internal auxiliary data. Since D-Res-LSTM and C-Res-LSTM adapt the feature learning in a coarse level, they achieve inferior results than sample-level adaptation. However, S-Res-LSTM narrows the distance between source and auxiliary modal distributions in a finer granularity. Compared with Res-LSTM, S-Res-LSTM brings about 2-3% improvement on the NTU dataset. We also notice that the jointly adaptation scheme J-Res-LSTM achieves comparable performance with S-Res-LSTM. It is mainly because S-Res-LSTM achieves the modality adaptation in the finest granularity. The distance between different categories and domains are also narrowed after sample-level adaptation. Thus, the training with jointly domain-, category-, and sample-level adaptation schemes does not bring more improvement. For the MSR dataset, we adopt different levels of adaptation schemes. Compared with baseline Res-LSTM, S-Res-LSTM improves the results on RGB and optical flow by 6.9% and 5.6%, respectively.

Table IV shows the results for the UCF-101 and JHMDB datasets with different auxiliary data sources. Note Res-LSTM is our baseline without using any auxiliary data. When using 3D skeletons from the NTU (CV) dataset, it can be regarded as the general case since the color videos are not aligned with skeletons. We use the domain-level adaptation on these two datasets (D-Res-LSTM*). Thanks to the introduction of adaptive representation learning guided by additional skeletons, the common spaces between two different modalities are explored. Thus, we are able to take advantage of the complementary information and improve the feature learning even if the two modalities are not aligned. The results of D-Res-LSTM* are improved about 1%-2% compared with the baseline Res-LSTM on both

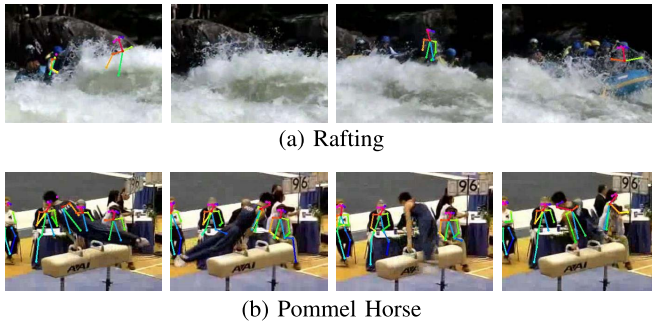


Fig. 3. Examples from the UCF-101 dataset: (a) rafting, (b) pommel horse. It is challenging to extract high-quality poses with [61] due to cluttered backgrounds and occlusions.

datasets. When using 2D poses extracted with [61], it can be regarded as the specific case and we are able to apply different levels of modal adaptation. The aligned poses provide much more specific guidance. It is observed that on the JHMDB dataset, the sample-level adaptation brings 7.0% and 3.3% improvement on RGB and optical flow, respectively. However, on the UCF dataset, the cluttered backgrounds and occlusions make it challenging to extract high-quality human poses, as shown in Fig. 3. It is less satisfactory to use extracted poses as auxiliary data. Besides, similar as Table III, the jointly adaptation scheme does not help boost the performance.

C. Effectiveness of Residual Subnetwork

The design of our residual subnetwork in Res-LSTM is motivated by two aspects: (1) The skip connection in the residual block aims to keep the original information from the source modalities. (2) The residual path is appended to incorporate compensated information, which is adapted from the auxiliary modality. To explore the effectiveness of the residual subnetwork and support our motivation, we compare Res-LSTM with another vanilla LSTM structure V-LSTM, as shown Fig. 2(b). There are N units for the LSTM layers and the following FC layer, so the number of parameters in V-LSTM is the same as Res-LSTM.

We conduct experiments on these two structures on the four datasets, as shown in Fig. 4. For V-LSTM, the modality adaptation is applied after the second LSTM layer. We adopt sample-level adaptation (S-V-LSTM) for the NTU and MSR datasets, and the domain-level scheme (D-V-LSTM) for the UCF-101 and JHMDB datasets. We obtain similar performance on V-LSTM and Res-LSTM when not using modality adaptation. However, once integrated with modality adaptation, better results are obtained with the architecture of Res-LSTM. The results illustrate that the residual subnetwork is more powerful and efficient in compensating auxiliary information for source data. Besides, our proposed modality adaptation is able to improve the performance with V-LSTM in most cases, which further confirms the effectiveness of the adaptive representation learning.

TABLE V
COMPARISONS ON NTU IN ACCURACY (%)

Methods	Modality	CS	CV
Trust Gate [13]	Ske.	69.2	77.7
STA-LSTM [11]	Ske.	73.4	81.2
VA-LSTM [64]	Ske.	79.4	87.6
LSTM	Poses	65.0	75.4
LSTM	Ske.	70.4	83.4
P-CNN [65]	RGB+Flow	53.8	61.7
TSN [4]	RGB+Flow	88.5	90.4
Chained MT [14]	Poses+RGB+Flow	80.8	–
Res-LSTM	RGB+Flow	88.8	94.1
Compensate-Train-Test Res-LSTM	Poses+RGB+Flow	88.7	94.4
Compensate-Train-Test Res-LSTM	Ske.+RGB+Flow	89.3	94.9
S-Res-LSTM	RGB+Flow	89.5	95.2
Compensate-Train-Test S-Res-LSTM	Poses+RGB+Flow	89.4	95.4
Compensate-Train-Test S-Res-LSTM	Ske.+RGB+Flow	90.0	96.3

TABLE VI
COMPARISONS ON MSR IN ACCURACY (%)

Methods	Modality	Acc.(%)
Action Ensemble [36]	Ske.	68.0
Moving Poselets [66]	Ske.	74.5
LSTM	Poses	63.8
LSTM	Ske.	72.5
P-CNN [65]	RGB+Flow	61.9
TSN [4]	RGB+Flow	88.1
Res-LSTM	RGB+Flow	84.4
Compensate-Train-Test Res-LSTM	Poses+RGB+Flow	85.0
Compensate-Train-Test Res-LSTM	Ske.+RGB+Flow	85.6
S-Res-LSTM	RGB+Flow	88.8
Compensate-Train-Test S-Res-LSTM	Poses+RGB+Flow	90.6
Compensate-Train-Test S-Res-LSTM	Ske.+RGB+Flow	91.9

TABLE VII
COMPARISONS ON UCF-101 IN ACCURACY (%). THE SUPERSCRIP^T* DENOTES AUXILIARY DATA ARE FROM EXTERNAL SOURCES (*i.e.*, THE NTU (CV) DATASET)

Methods	Modality	Acc.(%)
LSTM	Poses	36.0
LRCN [31]	RGB	82.9
C3D [22]	RGB	85.2
Two-Stream ConvNet [3]	RGB+Flow	88.0
Two-Stream + Fusion [6]	RGB+Flow	92.5
TSN [4]	RGB+Flow	94.0
Chained MT [14]	RGB+Flow	91.3
Res-LSTM	RGB+Flow	91.7
Compensate Train-Test Res-LSTM	Poses+RGB+Flow	91.8
D-Res-LSTM*	RGB+Flow	92.3
S-Res-LSTM	RGB+Flow	90.7

D. Comparisons to Other State-of-the-Arts

We compare our final results with other state-of-the-art methods, as shown in Table V, Table VI, Table VII, and Table VIII. For the specific case, we use S-Res-LSTM as our final model. And for the general case, we regard D-Res-LSTM* as our final model. In our work, the final prediction for a given video is generated by average score fusion of multiple streams [6].

To give a fair comparison, we mark the modalities contributed in the testing phase in Table V, Table VI, Table VII, and Table VIII. The involved modalities include 3D skeletons (ske.), 2D poses extracted with [61] (poses),

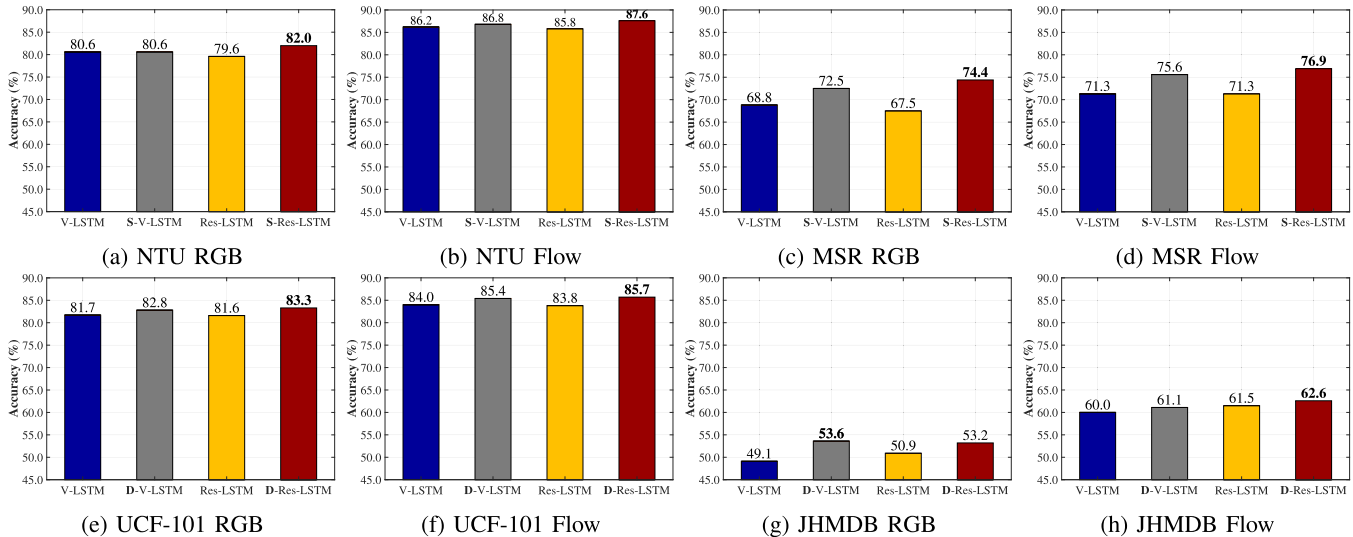


Fig. 4. Performance comparisons with different structures on different datasets. We apply sample-level adaptation on the NTU (CS) and MSR datasets, domain-level adaptation on the UCF-101 (split1) and JHMDB (split1) datasets, respectively. Our results show that modal adaptation can effectively improve the performance on Res-LSTM and V-LSTM. And the residual architecture is more efficient in compensating the feature learning and achieves better results.

TABLE VIII

COMPARISONS ON JHMDB IN ACCURACY (%). THE SUPERScript* DENOTES AUXILIARY DATA ARE FROM EXTERNAL SOURCES (*i.e.*, THE NTU (CV) DATASET)

Methods	Modality	Acc.(%)
LSTM	Poses	47.8
P-CNN [65]	RGB+Flow	61.1
Action Tubes [67]	RGB+Flow	62.5
TS R-CNN [68]	RGB+Flow	70.5
MR-TS R-CNN [68]	RGB+Flow	71.1
TSN [4]	RGB+Flow	73.4
Chained MT [14]	RGB+Flow	72.8
Res-LSTM	RGB+Flow	72.7
Compensate Train-Test Res-LSTM	Poses+RGB+Flow	72.9
D-Res-LSTM*	RGB+Flow	74.6
S-Res-LSTM	RGB+Flow	74.8

RGB videos (RGB) and optical flow data (Flow). We also show our recognition results only with skeleton/pose data using a 1-layer LSTM (LSTM). Our baseline results include Res-LSTM and Compensate-Train-Test Res-LSTM. For Res-LSTM, only source modalities are employed in the testing phase without any modal adaptation. For Compensate-Train-Test Res-LSTM, skeletons/poses are used in the testing phase and average fusion is employed to fuse RGB, optical flow and skeleton/pose streams.

In Table V and Table VI, our model S-Res-LSTM not only outperforms the baseline Res-LSTM, but also achieves higher performance than that even skeletons/poses are available in testing (Compensate-Train-Test Res-LSTM). It illustrates that our modal adaption scheme effectively compensates for the loss of skeletons at test time. In the meanwhile, we also find that fusing the results of LSTM and S-Res-LSTM, denoted as Compensate-Train-Test S-Res-LSTM, can further improve the performance. Furthermore, both S-Res-LSTM and Compensate-Train-Test S-Res-LSTM compare favorably to the existing methods on the NTU and MSR datasets. Note that we

implement P-CNN [65] and TSN [4] with the code provided by the authors. The result under the cross-view setting on the NTU dataset is not given in [14].

In Table VII and Table VIII, we report results over three splits on the UCF-101 and JHMDB datasets, respectively. On the UCF-101 dataset, as indicated in Table IV and Fig. 3, the inaccurate poses lead to poor action classification results, only 36.0% with LSTM. They also fail to improve the feature learning with S-Res-LSTM. However, D-Res-LSTM*, which is the two-stream fusion result with skeletons in NTU-CV as auxiliary data, outperforms Compensate Train-Test Res-LSTM and S-Res-LSTM. On the JHMDB dataset, we obtain similar results with D-Res-LSTM* and S-Res-LSTM, though they use different auxiliary data sources. And both of them effectively improve the baseline results. Compared with other state-of-the-arts, our model achieves comparable performance on the UCF-101 dataset and outperforms other methods on the JHMDB dataset. To give a fair comparison, we only compare our results using RGB and optical flow at the test time on both datasets. Note that the action categories in the UCF-101 dataset are very diverse, including *sky diving*, *rafting*, *surfing*, *etc.*, which are very different from those in the NTU dataset, such as *shaking hands*, *drinking*, *eating*, *etc.* We believe that training with skeletons in more relevant categories will bring further improvement on the UCF-101 dataset through the proposed modality adaption scheme.

E. Discussion

1) *Parameter Analysis*: The hyper parameter λ in Eq. (6) controls the degree of adaptive feature learning. To choose the scalar λ , we first observe the order of each item in the loss function, from which we roughly get the range of λ , then we vary different values to analyze its effect. The higher λ forces the network to narrow the distance between the

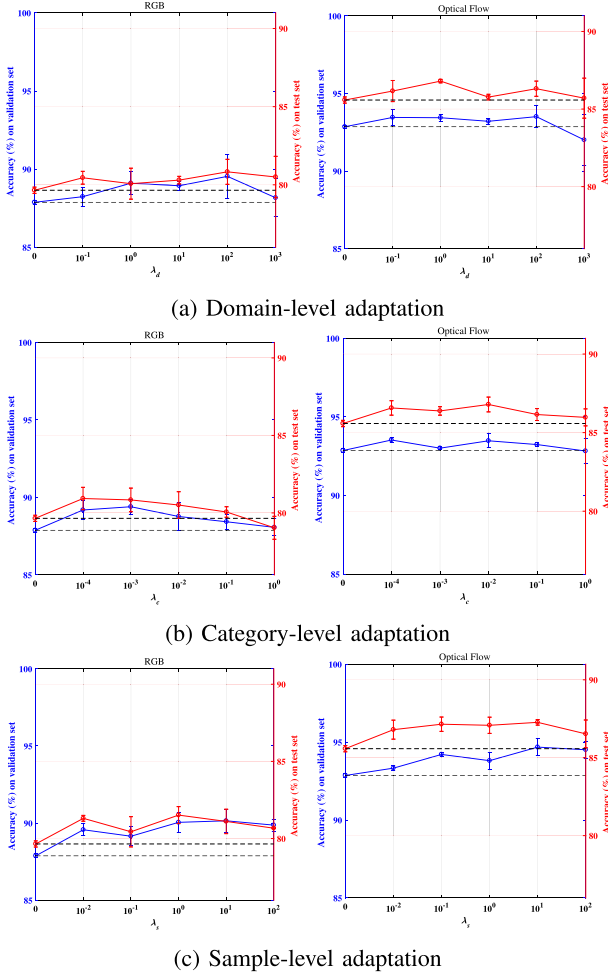


Fig. 5. Empirical study on the hyper parameter λ of different modal adaptation schemes on the NTU (CS) dataset.

source and auxiliary modalities as close as possible. However, due to the noises in the skeleton data, the features from the source modality could be interrupted and thus influence the performance. Moreover, with a lower λ , the network only borrows little complementary information from additional skeletons. Therefore, we study the effect of λ in Fig. 5 to seek the optimal. Note that $\lambda = 0$ indicates the Res-LSTM results when no modality adaptation is adopted, which is marked by horizontal black dash lines. We show the performance on validation and test sets with various λ in different levels of modality adaptation on the NTU (CS) dataset, respectively. In practice, the optimal λ is selected from the best results on the validation set. However, as shown in Fig. 5, different weights do not significantly influence the results. With λ in the proper range, different levels of adaptation bring improvement on the results consistently.

2) *Classification Analysis*: In Fig. 6, we list the histograms of the gain of S-Res-LSTM with respect to Res-LSTM. The gain values are calculated based on average precision for each class in the NTU dataset under the cross view setting. It is observed that for most classes, S-Res-LSTM outperforms the baseline. In the meantime, for RGB, we find the actions related with motions are improved much more, such as *clapping*

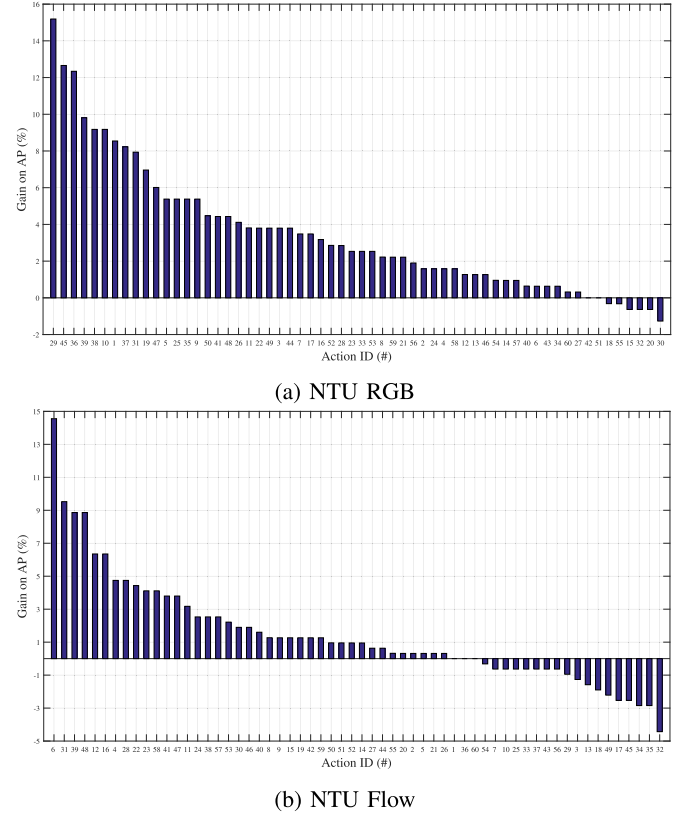


Fig. 6. Gain on average precision of S-Res-LSTM with respect to Res-LSTM on the NTU (CV) dataset. The index of the horizontal axis denotes the action ID provided in [12].

(#10), *putting the palms together* (#39), *saluting* (#38). With the auxiliary skeletons, our MCN is able to compensate source RGB features with motion information, increasing the discriminations of the actions. Meanwhile, the main contribution of skeleton data to optical flow is to provide extra posture information in actions where static frames are predominant, such as *pointing to something with finger* (#31), *nausea* (#48) and *writing* (#12). These actions with subtle optical flow are improved significantly with posture information compensated.

3) *Visualizations of Adapted Features*: To better understand the proposed modality adaptation method, we visualize the source and auxiliary feature vectors ($\hat{\mathbf{r}}_i$ and $\hat{\mathbf{a}}_i$ in Section III) that are learned from different modality adaptation levels. The visualizations are achieved with t-SNE techniques, as shown in Fig. 7. We randomly select five classes from the testing data in the NTU (CV) dataset. We denote different action classes with different colors. The features from source modalities are represented with triangles and those from skeletons are represented with squares. The source modality of the figures in the first row is RGB and the second row is optical flow. It is noticed that with a finer adaptation level, the features from source and auxiliary modalities become “closer” in the feature space. Compared with the distributions without modality adaptation, the source and auxiliary features after domain- or category-level adaptation stay tighter in general. However, when sample-level adaptation is adopted, the features in different modalities but sharing the same action class get

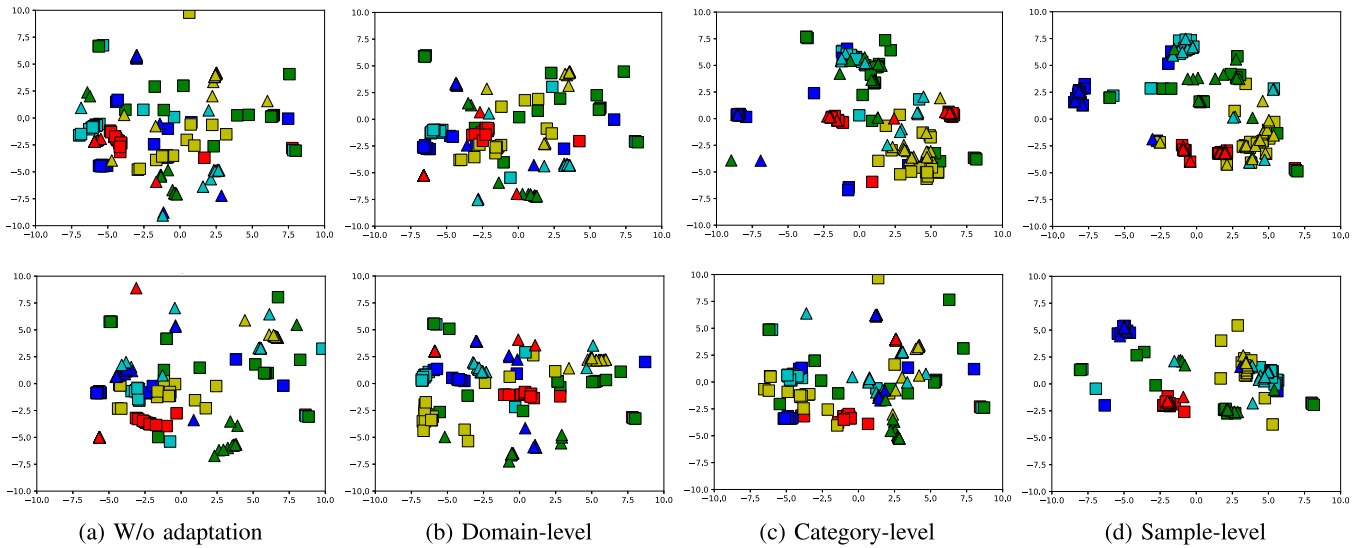


Fig. 7. The t-SNE results of source and auxiliary features with different modality adaptation levels (top: RGB & skeletons, bottom: optical flow & skeletons). We use the adapted features from the residual blocks in the Res-LSTM, D-Res-LSTM, C-Res-LSTM and S-Res-LSTM, respectively. Five classes are randomly selected from the testing data of the NTU (CV) dataset for visualization. Different colors denote different action classes. The features from source modalities are represented as triangles while those from skeletons are represented as squares.

together, which is able to further increase the discriminative performance.

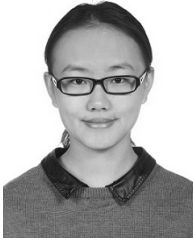
V. CONCLUSION

In this paper, we present a novel model for action recognition, Modality Compensation Network (MCN). Taking advantage of the auxiliary modality, we aim to compensate the feature learning in source modalities by adaptive representation learning. The modality adaptation block is developed to borrow the complementary information from the auxiliary modality, by narrowing the distance of source and auxiliary modal distributions. Adaptation schemes in different levels are employed according to the alignment of the training data. The ablation study illustrates the effectiveness of each component in our proposed network. Comprehensive analysis is given to better understand the modality adaptation scheme. Experiments demonstrate our model consistently improves the action recognition performance on four different datasets, and achieves remarkable performance compared with other state-of-the-arts.

REFERENCES

- [1] D. Weinland, R. Ronfard, and E. Boyerc, "A survey of vision-based methods for action representation, segmentation and recognition," *Comput. Vis. Image Understand.*, vol. 115, no. 2, pp. 224–241, 2011.
- [2] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu, "Action recognition by dense trajectories," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2011, pp. 3169–3176.
- [3] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 568–576.
- [4] L. Wang *et al.*, "Temporal segment networks: Towards good practices for deep action recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 20–36.
- [5] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Spatiotemporal multiplier networks for video action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 2097–2106.
- [6] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 1933–1941.
- [7] Z. Zhang, "Microsoft Kinect sensor and its effect," *IEEE MultiMedia*, vol. 19, no. 2, pp. 4–10, Feb. 2012.
- [8] J. Shotton *et al.*, "Real-time human pose recognition in parts from single depth images," *Commun. ACM*, vol. 56, no. 1, pp. 116–124, 2013.
- [9] G. Johansson, "Visual perception of biological motion and a model for it is analysis," *Perception Psychophys.*, vol. 14, no. 2, pp. 201–211, 1973.
- [10] W. Zhu *et al.*, "Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks," in *Proc. AAAI Conf. Artif. Intell.*, Aug. 2016, pp. 3697–3703.
- [11] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu, "An end-to-end spatio-temporal attention model for human action recognition from skeleton data," in *Proc. AAAI Conf. Artif. Intell.*, 2017, pp. 4263–4270.
- [12] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "NTU RGB+D: A large scale dataset for 3D human activity analysis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 1010–1019.
- [13] J. Liu, A. Shahroudy, D. Xu, and G. Wang, "Spatio-temporal LSTM with trust gates for 3D human action recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 816–833.
- [14] M. Zolfaghari, G. L. Oliveira, N. Sedaghat, and T. Brox, "Chained multi-stream networks exploiting pose, motion, and appearance for action classification and detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 2923–2932.
- [15] Z. Shi and T.-K. Kim, "Learning and refining of privileged information-based RNNs for action recognition from depth sequences," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 3461–3470.
- [16] B. Mahasseni and S. Todorovic, "Regularizing long short term memory with 3D human-skeleton sequences for action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 3054–3062.
- [17] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld, "Learning realistic human actions from movies," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [18] H. Wang, M. M. Ullah, A. Klaser, I. Laptev, and C. Schmid, "Evaluation of local spatio-temporal features for action recognition," in *Proc. Brit. Mach. Vis. Conf.*, 2009, pp. 1–124.
- [19] H. Wang and C. Schmid, "Action recognition with improved trajectories," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 3551–3558.
- [20] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1725–1732.

- [21] 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. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 4694–4702.
- [22] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3D convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 4489–4497.
- [23] B. Fernando, P. Anderson, M. Hutter, and S. Gould, "Discriminative hierarchical rank pooling for activity recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1924–1932.
- [24] H. Bilen, B. Fernando, E. Gavves, A. Vedaldi, and S. Gould, "Dynamic image networks for action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3034–3042.
- [25] Y. Wang, M. Long, J. Wang, and P. S. Yu, "Spatiotemporal pyramid network for video action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1529–1538.
- [26] Y. Zhang, L. Cheng, J. Wu, J. Cai, M. N. Do, and J. Lu, "Action recognition in still images with minimum annotation efforts," *IEEE Trans. Image Process.*, vol. 25, no. 11, pp. 5479–5490, Nov. 2016.
- [27] C. Feichtenhofer, A. Pinz, and R. Wildes, "Spatiotemporal residual networks for video action recognition," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3468–3476.
- [28] Y. Li, S. Song, Y. Li, and J. Liu, "Temporal bilinear networks for video action recognition," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 8674–8681.
- [29] H. Jiang, Y. Li, S. Song, and J. Liu, "Rethinking fusion baselines for multi-modal human action recognition," in *Proc. Pacific Rim Conf. Multimedia*, 2018, pp. 178–187.
- [30] Y. Li, W. Li, V. Mahadevan, and N. Vasconcelos, "VLAD3: Encoding dynamics of deep features for action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1951–1960.
- [31] J. Donahue *et al.*, "Long-term recurrent convolutional networks for visual recognition and description," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 2625–2634.
- [32] Z. Wu, Y.-G. Jiang, X. Wang, H. Ye, and X. Xue, "Multi-stream multi-class fusion of deep networks for video classification," in *Proc. ACM Multimedia Conf. (MM)*, 2016, pp. 791–800.
- [33] Z. Wu, X. Wang, Y.-G. Jiang, H. Ye, and X. Xue, "Modeling spatial-temporal clues in a hybrid deep learning framework for video classification," in *Proc. 23rd ACM Int. Conf. Multimedia (MM)*, 2015, pp. 461–470.
- [34] R. Vemulapalli, F. Arrate, and R. Chellappa, "Human action recognition by representing 3D skeletons as points in a lie group," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 588–595.
- [35] R. Vemulapalli and R. Chellappa, "Rolling rotations for recognizing human actions from 3D skeletal data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 4471–4479.
- [36] J. Wang, Z. Liu, Y. Wu, and J. Yuan, "Mining actionlet ensemble for action recognition with depth cameras," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 1290–1297.
- [37] Y. Goutsu, W. Takano, and Y. Nakamura, "Motion recognition employing multiple kernel learning of fisher vectors using local skeleton features," in *Proc. IEEE Int. Conf. Comput. Vis. Workshop (ICCVW)*, Dec. 2015, pp. 79–86.
- [38] F. Lv and R. Nevatia, "Recognition and segmentation of 3-D human action using HMM and multi-class adaboost," in *Proc. Eur. Conf. Comput. Vis.*, 2006, pp. 359–372.
- [39] Y. Du, W. Wang, and L. Wang, "Hierarchical recurrent neural network for skeleton based action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1110–1118.
- [40] Y. Du, Y. Fu, and L. Wang, "Representation learning of temporal dynamics for skeleton-based action recognition," *IEEE Trans. Image Process.*, vol. 25, no. 7, pp. 3010–3022, Jul. 2016.
- [41] H. Wang and L. Wang, "Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 499–508.
- [42] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu, "Spatio-temporal attention-based LSTM networks for 3D action recognition and detection," *IEEE Trans. Image Process.*, vol. 27, no. 7, pp. 3459–3471, Jul. 2018.
- [43] Q. Ke, M. Bennamoun, S. An, F. Sohel, and F. Boussaid, "A new representation of skeleton sequences for 3D action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4570–4579.
- [44] J. Liu, Y. Li, S. Song, J. Xing, C. Lan, and W. Zeng, "Multi-modality multi-task recurrent neural network for online action detection," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 29, no. 9, pp. 2667–2682, Sep. 2019.
- [45] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu, "Skeleton-indexed deep multi-modal feature learning for high performance human action recognition," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, Jul. 2018, pp. 1–6.
- [46] A. Wang, J. Cai, J. Lu, and T.-J. Cham, "MMSS: Multi-modal sharable and specific feature learning for RGB-D object recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1125–1133.
- [47] J. Wang, Z. Wang, D. Tao, S. See, and G. Wang, "Learning common and specific features for RGB-D semantic segmentation with deconvolutional networks," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 664–679.
- [48] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [49] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by back-propagation," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1180–1189.
- [50] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 97–105.
- [51] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, "Domain separation networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 343–351.
- [52] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," 2015, *arXiv:1502.03167*. [Online]. Available: <https://arxiv.org/abs/1502.03167>
- [53] A. Graves, *Supervised Sequence Labelling With Recurrent Neural Networks*, vol. 385. Berlin, Germany: Springer, 2012.
- [54] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [55] Y. Wang and F. Tian, "Recurrent residual learning for sequence classification," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2016, pp. 938–943.
- [56] S. A. Hasan *et al.*, "Neural paraphrase generation with stacked residual LSTM networks," in *Proc. Int. Conf. Comput. Linguistics, Tech. Papers*, 2016, pp. 2923–2934.
- [57] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *J. Mach. Learn. Res.*, vol. 13, pp. 723–773, Mar. 2012.
- [58] K. Soomro, A. R. Zamir, and M. Shah, "UCF101: A dataset of 101 human actions classes from videos in the wild," 2012, *arXiv:1212.0402*. [Online]. Available: <https://arxiv.org/abs/1212.0402>
- [59] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black, "Towards understanding action recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 3192–3199.
- [60] L. Chunhui, H. Yueyu, L. Yanghao, S. Sijie, and L. Jiaying, "PKU-MMD: A large scale benchmark for continuous multi-modal human action understanding," 2017, *arXiv:1703.07475*. [Online]. Available: <https://arxiv.org/abs/1703.07475>
- [61] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2D pose estimation using part affinity fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 7291–7299.
- [62] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [63] W. Zaremba, I. Sutskever, and O. Vinyals, "Recurrent neural network regularization," 2014, *arXiv:1409.2329*. [Online]. Available: <https://arxiv.org/abs/1409.2329>
- [64] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng, "View adaptive recurrent neural networks for high performance human action recognition from skeleton data," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 2117–2126.
- [65] G. Chéron, I. Laptev, and C. Schmid, "P-CNN: Pose-based CNN features for action recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2015, pp. 3218–3226.
- [66] L. Tao and R. Vidal, "Moving poselets: A discriminative and interpretable skeletal motion representation for action recognition," in *Proc. Int. Conf. Mach. Learn. Workshops*, 2015, pp. 61–69.
- [67] G. Gkioxari and J. Malik, "Finding action tubes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 759–768.
- [68] X. Peng and C. Schmid, "Multi-region two-stream R-CNN for action detection," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 744–759.



Sijie Song (Student Member, IEEE) received the B.S. degree in computer science from Peking University, Beijing, China, in 2016, where she is currently pursuing the Ph.D. degree with the Wangxuan Institute of Computer Technology. Her research interests include computer vision and image processing.



Jiaying Liu (Senior Member, IEEE) received the Ph.D. degree (Hons.) in computer science from Peking University, Beijing, China, 2010.

She was a Visiting Scholar with the University of Southern California, Los Angeles, from 2007 to 2008. She is currently an Associate Professor with the Wangxuan Institute of Computer Technology, Peking University. She has authored over 100 technical articles in refereed journals and proceedings, and holds 42 granted patents. Her current research interests include multimedia signal processing, compression, and computer vision. Dr. Liu is a Senior Member of CSIG and CCF. She was a Visiting Researcher with Microsoft Research Asia in 2015 supported by the Star Track Young Faculties Award. She has served as a member of Membership Services Committee in IEEE Signal Processing Society, a member of Multimedia Systems & Applications Technical Committee (MSA TC), Visual Signal Processing and Communications Technical Committee (VSPEC TC) in IEEE Circuits and Systems Society, a member of the Image, Video, and Multimedia (IVM) Technical Committee in APSIPA. She has served as an Associate Editor for the *IEEE TRANSACTION ON IMAGE PROCESSING*, and Elsevier *JVCI*, the Technical Program Chair of IEEE VCIP-2019/ACM ICMR-2021, the Publicity Chair of IEEE ICME-2020/ICIP-2019/VCIP-2018/MIPR 2020, and the Area Chair of ECCV-2020/ICCV-2019. She was the APSIPA Distinguished Lecturer from 2016 to 2017.



Yanghao Li (Student Member, IEEE) received the B.S. degree in computer science from Peking University, Beijing, China, in 2015, and the M.S. degree from the Wangxuan Institute of Computer Technology, Peking University, in 2018. His current research interests include action analysis and computer vision.



Zongming Guo (Member, IEEE) received the B.S. degree in mathematics, and the M.S. and Ph.D. degrees in computer science from Peking University, Beijing, China, in 1987, 1990, and 1994, respectively.

He is currently a Professor with the Wangxuan Institute of Computer Technology, Peking University. His current research interests include video coding, processing, and communication.

Dr. Guo is also the Executive Member of the China-Society of Motion Picture and Television Engineers. He was a recipient of the First Prize of the State Administration of Radio Film and Television Award in 2004, the First Prize of the Ministry of Education Science and Technology Progress Award in 2006, the Second Prize of the National Science and Technology Award in 2007, the Wang Xuan News Technology Award and the Chia Tai Teaching Award in 2008, the Government Allowance granted by the State Council in 2009, and the Distinguished Doctoral Dissertation Advisor Award of Peking University in 2012 and 2013.