RGB-D-based human motion recognition with deep learning: A survey[☆]Pichao Wang^a, Wanqing Li^a, Philip Ogunbona^a, Jun Wan^{*,b}, Sergio Escalera^c^a Advanced Multimedia Research Lab, University of Wollongong, NSW 2522, Australia^b Center for Biometrics and Security Research (CBSR) & National Laboratory of Pattern Recognition (NLPR), Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing 100190, China^c University of Barcelona and Computer Vision Center, Campus UAB, 08193 Bellaterra, Barcelona, Spain

ARTICLE INFO

Keywords:

Human motion recognition

RGB-D data

Deep learning

Survey

MSC:

68T10

68T30

65D19

ABSTRACT

Human motion recognition is one of the most important branches of human-centered research activities. In recent years, motion recognition based on RGB-D data has attracted much attention. Along with the development in artificial intelligence, deep learning techniques have gained remarkable success in computer vision. In particular, convolutional neural networks (CNN) have achieved great success for image-based tasks, and recurrent neural networks (RNN) are renowned for sequence-based problems. Specifically, deep learning methods based on the CNN and RNN architectures have been adopted for motion recognition using RGB-D data. In this paper, a detailed overview of recent advances in RGB-D-based motion recognition is presented. The reviewed methods are broadly categorized into four groups, depending on the modality adopted for recognition: RGB-based, depth-based, skeleton-based and RGB + D-based. As a survey focused on the application of deep learning to RGB-D-based motion recognition, we explicitly discuss the advantages and limitations of existing techniques. Particularly, we highlighted the methods of encoding spatial-temporal-structural information inherent in video sequence, and discuss potential directions for future research.

1. Introduction

Among the several human-centered research activities (e.g. human detection, tracking, pose estimation and motion recognition) in computer vision, human motion recognition is particularly important due to its potential application in video surveillance, human computer interfaces, ambient assisted living, human–robot interaction, intelligent driving, etc. A human motion recognition task can be summarised as the automatic identification of human behaviours from images or video sequences. The complexity and duration of the motion involved can be used as basis for broad categorization into four kinds namely gesture, action, interaction and group activity. A *gesture* can be defined as the basic movement or positioning of the hand, arm, body, or head that communicates an idea, emotion, etc. “Hand waving” and “nodding” are some typical examples of gestures. Usually, a gesture has relatively short duration. An *action* is considered as a type of motion performed by a single person during short time period and involves multiple body parts, in contrast with the few body parts that involved in gesture. An *activity* is composed by a sequence of actions. An *interaction* is a type of motion performed by two actors; one actor is human while the other may be human or an object. This implies that the interaction category

will include human–human or human–object interaction. “Hugging each other” and “playing guitar” are examples of these two kinds of interaction, respectively. *Group activity* is the most complex type of activity, and it may be a combination of gestures, actions and interactions. Necessarily, it involves more than two humans and from zero to multiple objects. Examples of group activities would include “two teams playing basketball” and “group meeting”.

Early research on human motion recognition was dominated by the analysis of still images or videos (Aggarwal and Cai, 1999; Guo and Lai, 2014; Poppe, 2010; Turaga et al., 2008; Wang et al., 2003; Zhu et al., 2016a). Most of these efforts used color and texture cues in 2D images for recognition. However, the task remains challenging due to problems posed by background clutter, partial occlusion, view-point, lighting changes, execution rate and biometric variation. This challenge remains even with current deep learning approaches (Asadi-Aghbolaghi et al., 2017; Herath et al., 2017).

With the recent development of cost-effective RGB-D sensors, such as Microsoft Kinect™ and Asus Xtion™, RGB-D-based motion recognition has attracted much attention. This is largely because the extra dimension (depth) is insensitive to illumination changes and includes rich 3D structural information of the scene. Additionally, 3D positions

[☆] The name of the Editor in chief Dr. Norimichi Ukita.

* Corresponding author.

E-mail addresses: pw212@uowmail.edu.au (P. Wang), jun.wan@nlpr.ia.ac.cn (J. Wan).

Table 1

Statistics of the public available benchmark datasets that are commonly used for evaluation with deep learning. Notation for the header: Seg: Segmented, Con: Continuous, D: Depth, S: Skeleton, Au: Audio, Ac: Accelerometer, IR: IR videos, #: number of, JI: Jaccard Index.

Dataset	year	Acquisition device	Seg/Con	Modality	#Class	#Subjects	#Samples	#Views	Metric
CMU Mocap	2001	Mocap	Seg	RGB,S	45	144	2235	1	Accuracy
HDM05	2007	Mocap	Seg	RGB,S	130	5	2337	1	Accuracy
MSR-Action3D	2010	Kinect v1	Seg	S,D	20	10	567	1	Accuracy
MSRC-12	2012	Kinect v1	Seg	S	12	30	594	1	Accuracy
MSR DailyActivity3D	2012	Kinect v1	Seg	RGB,D,S	16	10	320	1	Accuracy
UTKinect	2012	Kinect v1	Seg	RGB,D,S	10	10	200	1	Accuracy
G3D	2012	Kinect v1	Seg	RGB,D,S	5	5	200	1	Accuracy
SBU Kinect Interaction	2012	Kinect v1	Seg	RGB,D,S	7	8	300	1	Accuracy
Berkeley MHAD	2013	Mocap Kinect v1	Seg	RGB,D,S,Au,Ac	12	12	660	4	Accuracy
Multiview Action3D	2014	Kinect v1	Seg	RGB,D,S	10	10	1475	3	Accuracy
ChaLearn LAP IsoGD	2016	Kinect v1	Seg	RGB,D	249	21	47,933	1	Accuracy
NTU RGB + D	2016	Kinect v2	Seg	RGB,D,S,IR	60	40	56,880	80	Accuracy
ChaLearn2014	2014	Kinect v1	Con	RGB,D,S,Au	20	27	13,858	1	Accuracy JI etc.
ChaLearn LAP ConGD	2016	Kinect v1	Con	RGB,D	249	21	22,535	1	JI
PKU-MMD	2017	Kinect v2	Con	RGB,D,S,IR	51	66	1076	3	JI etc.

of body joints can be estimated from depth maps (Shotton et al., 2011). As a consequence, several methods based on RGB-D data have been proposed and the approach has proven to be a promising direction for human motion analysis (Table 1).

Several survey papers have summarized the research on human motion recognition using RGB-D data (Aggarwal and Xia, 2014; Chen et al., 2013; Cheng et al., 2016; Escalera et al., 2016; Han et al., 2017; Presti and La Cascia, 2016; Ye et al., 2013; Zhang et al., 2016b). Specifically, Chen et al. (2013) focused on depth sensors, pre-processing of depth data, depth-based action recognition methods and datasets. In their work, Ye et al. (2013) presented an overview of approaches using depth and skeleton modalities for tasks including activity recognition, head/hand pose estimation, facial feature detection and gesture recognition. The survey presented by Aggarwal and Xia (2014) summarized five categories of representations based on 3D silhouettes, skeletal joints/body part location, local spatial-temporal features, scene flow features and local occupancy features. The work of Cheng et al. (2016) focused on RGB-D-based hand gesture recognition datasets and summarized corresponding methods from three perspectives: static hand gesture recognition, hand trajectory gesture recognition and continuous hand gesture recognition. In another effort Escalera et al. (2016) reviewed the challenges and methods for gesture recognition using multimodal data. Some of the surveys have focused on available datasets for RGB-D research. For example, the work of Zhang et al. (2016b) described available benchmark RGB-D datasets for action/activity recognition and included 27 single-view datasets, 10 multi-view datasets and 7 multi-person datasets. Other works as Presti and La Cascia (2016) and Han et al. (2017) mainly reviewed skeleton-based representation and approaches for action recognition. A short survey on RGB-D action recognition using deep learning was recently presented in Asadi-Aghbolaghi et al. (2017), analysing RGB and depth cues in terms of 2DCNN, 3DCNN, and Deep temporal approaches.

All above surveys mainly focused on the analysis of handcrafted features. Here, we provide a comprehensive review of RGB-D-based human motion recognition using deep learning approaches. Even while focusing on deep learning approaches, the nature of the input data is still important. RGB-D data for human motion analysis comprises three modalities: RGB, depth and skeleton. The main characteristic of RGB data is its shape, color and texture which brings the benefits of extracting interesting points and optical flow. Compared to RGB videos, the depth modality is insensitive to illumination variations, invariant to color and texture changes, reliable for estimating body silhouette and skeleton, and provides rich 3D structural information of the scene. Differently from RGB and depth, skeleton data containing the positions of human joints, is a relatively high-level feature for motion recognition. The different properties of the three modalities have inspired the

various methods found in the literature. For example, optical flow-based methods with Convolutional Neural Networks (CNN) is very effective for RGB channel (Duan et al., 2016); depth rank pooling based-method with CNN is a good choice for depth modality (Wang et al., 2016c); sequence based method with Recurrent Neural Networks (RNN) (Liu et al., 2017a) and image-based method with CNN (Wang et al., 2016e) are effective for skeleton; and scene flow-based method using CNN are promising for RGB + D channels (Wang et al., 2017b). These methods are very effective for specific modalities, but not always the case for all the modalities. Given these observations, this survey identified four broad categories of methods based on the modality adopted for human motion recognition. The categories include RGB-based, depth-based, skeleton-based and RGB + D-based.

In each category, two sub-divisions are further identified, namely segmented human motion recognition and continuous/online motion recognition. For segmented motion recognition, the scenario of the problem can be simply described as classifying a well delineated sequence of video frames as one of a set of motion types. This is in contrast to continuous/online human motion recognition where there are no a priori given boundaries of motion execution. The online situation is compounded by the fact that the video sequence is not recorded and the algorithm must deal with frames as they are being captured, save for possibly a small data cache.

During the performance of a specified motion spatial information which refers to the spatial configuration of human body at an instant of time (e.g. relative positions of the human body parts) can be identified. Similarly, there is the temporal information which characterizes the spatial configuration of the body over time (i.e. the dynamics of the body). Lastly, the structural information encodes the coordination and synchronization of body parts over the period in which the action is being performed. It describes the relationship of the spatial configurations of human body across different time slots.

In reviewing the various methods, consideration has been given to the manner in which the spatial, temporal and structural information have been exploited. Hence, the survey discusses the advantages and limitations of the reviewed methods from the spatial-temporal-structural encoding viewpoint, and suggests potential directions for future research.

A key novelty of this survey is the focus on three architectures of neural networks used in the various deep learning methods reviewed namely CNN-based, RNN-based and other structured networks. Fig. 1 illustrates the taxonomy underpinning this survey.

This is one of the first surveys dedicated to RGB-D-based human motion recognition using deep learning. Apart from this claim, this survey distinguishes itself from other surveys through the following contributions:

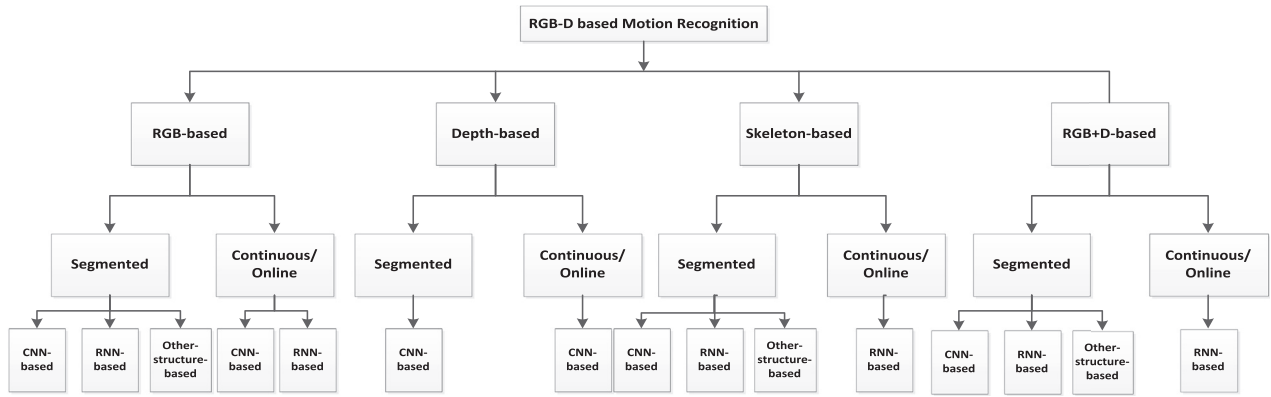


Fig. 1. Categorisation of the methods for RGB-D-based motion recognition using deep learning.

- Comprehensive coverage of the most recent and advanced deep learning-based methods developed in the last five years, thereby providing readers with a complete overview of recent research results and state-of-the-art methods.
- Insightful categorization and analysis of methods based on the different properties of the modalities; highlight of the pros and cons of the methods described in the reviewed papers from the viewpoint of spatial-temporal-structural encoding.
- Discussion of the challenges of RGB-D-based motion recognition; analysis of the limitations of available methods and discussion of potential research directions.

Additionally, several recently released or commonly used RGB-D-based benchmark datasets associated with deep learning are surveyed. The main application domain of interest in this survey paper is human motion recognition based on RGB-D data, including gesture recognition, action/activity recognition and interaction recognition. The lack of datasets focused on RGB-D-based group activity recognition has led to paucity of research on this topic and thus this survey does not cover this topic. Other RGB-D-based human-centered applications, such as human detection, tracking and pose estimation, are also not the focus of this paper. For surveys on RGB-D data acquisition readers are referred to Chen et al. (2013), Cheng et al. (2016), Han et al. (2017).

Subsequent sections of the this survey are organized as follows. Commonly used RGB-D-based benchmark datasets are described in Section 2. Sections 3 to 6 discuss methods of RGB-D-based motion recognition using deep learning from four perspectives: RGB-based motion recognition, depth-based motion recognition, skeleton-based motion recognition and RGB+D-based motion recognition. Challenges of RGB-D-based motion recognition and pointers to future directions are presented in Section 7. The survey provides concluding remarks in Section 8.

2. Benchmark datasets

Over the last decade, a number of RGB-D benchmark datasets have been collected and made publicly available for the research community. The sources of the datasets are mainly of three categories (Chen et al., 2013; Cheng et al., 2016; Han et al., 2017): Motion capture (Mocap) system, structured-light cameras (e.g. Kinect v1) and time-of-flight (ToF) cameras (e.g. Kinect v2). Hence the modalities of the datasets cover RGB, depth, skeleton and their combinations. With the advance of deep learning, deep methods have been developed for estimating skeletons directly from single images or video sequences, such as DeepPose (Toshev and Szegedy, 2014), Deepcrut (Insafutdinov et al., 2016) and Adversarial PoseNet (Chen et al., 2017). A comprehensive survey of these datasets have appeared in the literature (see e.g. Cheng et al., 2016 for hand gestures and (Zhang et al., 2016b) for action recognition). In the present survey only 15 large-scale datasets that have been

commonly adopted for evaluating deep learning-based methods are described. The reader is referred to Table 2 for a sample of works (publications) that have used these datasets. For the purpose of this survey, the datasets have been divided into two groups: segmented datasets and continuous/online datasets.

2.1. Segmented datasets

By segmented datasets we refer to those datasets where samples correspond to a whole begin-end action/gestures, with one segment for one action. They are mainly used for classification purposes. The following are several segmented datasets commonly used for the evaluation of methods based on deep learning.

2.1.1. CMU Mocap

CMU Graphics Lab Motion Capture Database (CMU Mocap) (CMU, 2001) (<http://mocap.cs.cmu.edu/>) is one of the earliest source of data that covers a wide range of human actions, including interactions between two subjects, human locomotion, interactions with uneven terrain, sports, and other human actions. This dataset consists of RGB and skeleton modalities.

2.1.2. HDM05

Motion Capture Database HDM05 (Müller et al., 2007) (<http://resources.mpi-inf.mpg.de/HDM05/>) was captured by an optical marker-based technology with the frequency of 120 Hz, which contains 2337 sequences for 130 actions performed by 5 non-professional actors, and 31 joints in each frame. Besides skeleton data, this dataset also provides RGB data.

2.1.3. MSR-Action3D

MSR-Action3D (Li et al., 2010) (<http://www.uow.edu.au/~wanqing/#MSRAction3DDatasets>) is the first public benchmark RGB-D action dataset collected using Kinect™ sensor by Microsoft Research, Redmond and University of Wollongong in 2010. The dataset contains 20 actions: *high arm wave, horizontal arm wave, hammer, hand catch, forward punch, high throw, draw x, draw tick, draw circle, hand clap, two hand wave, side-boxing, bend, forward kick, side kick, jogging, tennis serve, golf swing, pickup and throw*. Ten subjects performed these actions three times. All the videos were recorded from a fixed point of view and the subjects were facing the camera while performing the actions. The background of the dataset was removed by some post-processing. Specifically, if an action needs to be performed with one arm or one leg, the actors were required to perform it using right arm or leg.

2.1.4. MSRC-12

MSRC-12 dataset (Fothergill et al., 2012) (<http://research.microsoft.com/en-us/um/cambridge/projects/msrc12/>) was collected by Microsoft Research Cambridge and University of Cambridge in 2012. The

Table 2

Performance comparison among different methods on commonly used RGB-D datasets. Notation: D: Depth, S: Skeleton, Acc: Accuracy, JI: Jaccard Index, cs: cross-subject setting, cv: cross-view setting. Without specific notation, accuracy is used for metric.

Dataset	Reference	Modality	Method	Fusion method	Metric
CMU Mocap	(Zhu et al., 2016d)	Skeleton	Co-occurrence + LSTM	None	81.04%
	(Ke et al., 2017b)	Skeleton	Clips + CNN + MTLN	None	88.30%
HDM05	(Huang et al., 2017)	Skeleton	Deep Learning on Lie Group	None	75.78%
	(Du et al., 2015b)	Skeleton	HBRNN-L	None	96.92%
	(Zhu et al., 2016d)	Skeleton	Co-occurrence + LSTM	None	97.25%
MSR-Action3D	(Liu et al., 2016c)	Depth	3DCNN	None	84.07%
	(Veeriah et al., 2015)	Skeleton	dRNN	None	92.03%
	(Du et al., 2015b)	Skeleton	HBRNN-L	None	94.49%
	(Shi and Kim, 2017)	Depth + Skeleton	PRNN	Side information	94.90%
	(Wang et al., 2016b)	Depth	WHDMM + CNN	None	100.00%
	(Wang et al., 2017c)	Depth	S ² DDI	None	100.00%
MSRC-12	(Wang et al., 2016e)	Skeleton	JTM + CNN	None	93.12%
	(Hou et al., 2016)	Skeleton	SOS + CNN	None	94.27%
	(Liu et al., 2017b)	Skeleton	Enhanced Visualization + CNN	None	96.62%
MSRDaily Activity3D	(Wang et al., 2016b)	Depth	WHDMM + CNN	None	85.00%
	(Luo et al., 2017)	Depth	Unsupervised + ConvLSTM	None	86.90%
	(Shahroudy et al., 2017)	RGB + Depth	DSSCA-SSLM	Hierarchical fusion	97.50%
	(Wang et al., 2017c)	Depth	S ² DDI	None	97.50%
UTKinect	(Liu et al., 2016c)	Depth	3DCNN	None	82.00%
	(Wang et al., 2016b)	Depth	WHDMM + CNN	None	90.91%
	(Zhang et al., 2017c)	Skeleton	JL_d + RNN	None	95.96%
	(Lee et al., 2017)	Skeleton	Ensemble TS-LSTM	Ensemble	96.97%
	(Liu et al., 2016a)	Skeleton	ST-LSTM + Trust Gate	None	97.00%
G3D	(Huang et al., 2017)	Skeleton	Deep Learning on Lie Group	None	89.10%
	(Wang et al., 2016e)	Skeleton	JTM + CNN	None	94.24%
	(Hou et al., 2016)	Skeleton	SOS + CNN	None	95.45%
	(Wang et al., 2017c)	Depth	S ² DDI	None	96.06%
SBU Kinect interaction	(Shi and Kim, 2017)	Depth + Skeleton	PRNN	Side Information	89.20%
	(Zhu et al., 2016d)	Skeleton	Co-occurrence + LSTM	None	90.41%
	(Song et al., 2017)	Skeleton	STA-LSTM	None	91.51%
	(Liu et al., 2016a)	Skeleton	ST-LSTM + Trust Gate	None	93.30%
	(Ke et al., 2017a)	Skeleton	SkeletonNet(CNN)	None	93.47%
	(Ke et al., 2017b)	Skeleton	Clips + CNN + MTLN	None	93.57%
	(Zhang et al., 2017c)	Skeleton	JL_d + RNN	None	99.02%
Berkeley MHAD	(Ijjina and Krishna Mohan, 2016)	Skeleton	Stacked Autoencoder	None	98.03%
	(Du et al., 2015b)	Skeleton	HBRNN-L	None	100.00%
	(Zhang et al., 2017c)	Skeleton	JL_d + RNN	None	100.00%
	(Du et al., 2015a)	Skeleton	Skeleton Matrix + CNN	None	100.00%
	(Liu et al., 2016a)	Skeleton	ST-LSTM + Trust Gate	None	100.00%
Multiview Action3D	(Rahmani and Mian, 2016)	Depth	Fitting model + CNN	None	92.00%
	(Liu et al., 2017b)	Skeleton	Enhanced Visualization + CNN	None	92.61%
ChaLearn LAP IsoGD	(Wang et al., 2017b)	RGB + Depth	Sceneflow + CNN	Early Fusion	36.27%
	(Wang et al., 2016c)	Depth	DynamicImages + CNN	None	39.23%
	(Wang et al., 2018a)	Depth	DynamicMaps + CNN	None	43.72%
	(Wang et al., 2018b)	Depth + RGB	Cooperative CNN	None	44.80%
	(Zhu et al., 2016b)	RGB + Depth	Pyramidal C3D	Score Fusion	45.02%
	(Duan et al., 2016)	RGB + Depth	2SCVN-3DDSN	Score Fusion	49.17%
	(Li et al., 2016c)	RGB + Depth	C3D	Score Fusion	49.20%
	(Zhu et al., 2017a)	RGB + Depth	C3D + ConvLSTM	Score Fusion	51.02%
NTU RGB + D	(Luo et al., 2017)	RGB	Unsupervised + ConvLSTM	None	56.00%(cs)
	(Huang et al., 2017)	Skeleton	Deep Learning on Lie Group	None	61.37%(cs)
					66.95%(cv)
	(Shahroudy et al., 2016)	Skeleton	2Layer P-LSTM	None	62.93%(cs)
					70.27%(cv)
	(Luo et al., 2017)	Depth	Unsupervised + ConvLSTM	None	66.20%(cs)
	(Liu et al., 2016a)	Skeleton	ST-LSTM + Trust Gate	None	69.20%(cs)
					77.70%(cv)
	(Zhang et al., 2017c)	Skeleton	JL_d + RNN	None	70.26%(cs)
					82.39%(cv)
	(Wang et al., 2016e)	Skeleton	JTM + CNN	None	73.40%(cs)
					75.20%(cv)
	(Song et al., 2017)	Skeleton	STA-LSTM	None	73.40%(cs)
					81.20%(cv)
	(Kim and Reiter, 2017)	Skeleton	Res-TCN	None	74.30%(cs)
					83.10%(cv)
	(Lee et al., 2017)	Skeleton	Ensemble TS-LSTM	Ensemble	74.60%(cs)
					81.25%(cv)
	(Shahroudy et al., 2017)	RGB + Depth	DSSCA-SSLM	Hierarchical Fusion	74.86%(cs)
	(Ke et al., 2017a)	Skeleton	SkeletonNet(CNN)	None	75.94%(cs)
					81.16%(cv)
	(Li et al., 2017a)	Skeleton	JDM + CNN	None	76.20%(cs)
					82.30%(cv)
	(Ke et al., 2017b)	Skeleton	Clips + CNN + MTLN	None	79.57%(cs)

(continued on next page)

Table 2 (continued)

Dataset	Reference	Modality	Method	Fusion method	Metric
ChaLearn2014	(Liu et al., 2017b)	Skeleton	Enhanced Visualization + CNN	None	84.83%(cv) 80.03%(cs) 87.21%(cv)
	(Zolfaghari et al., 2017)	RGB	Chained Multi-stream	Markov chain	80.80%(cs)
	(Wang et al., 2018b)	RGB + Depth	Cooperative CNN	None	86.42%(cs) 89.08%(cv)
	(Wang et al., 2018a)	Depth	DynamicMaps + CNN	None	87.08%(cs) 84.22%(cv)
	(Du et al., 2015a)	Skeleton	Skeleton Matrix + CNN	None	91.16%
	(Pigou et al., 2016)	RGB	Temp Conv + RNN,LSTM	None	94.49%
	(Molchanov et al., 2016)	RGB + Depth	CTC + RNN	Score Fusion	0.906(JI) 98.20%
	(Wu et al., 2016)	RGB + Depth + Skeleton	DDNN	Late Fusion	0.980(JI) 0.809(JI)
	(Wang et al., 2016d)	Depth	IDMM + CNN	None	0.240(JI)
	(Chai et al., 2016)	RGB + Depth	RNN	Score Fusion	0.266(JI)
ChaLearn LAP ConGD	(Camgoz et al., 2016)	RGB	C3D	None	0.343(JI)
	(Wang et al., 2018a)	Depth	DynamicMaps + CNN	None	0.391(JI)
	(Song et al., 2017)	S	SlidingWindow + STA-LSTM	None	0.427(JI)(cs) 0.468(JI)(cv)
	(Li et al., 2016a)	S	JCR-RNN	None	0.479(JI)(cs) 0.728(JI)(cv)

authors provided three familiar and easy to prepare instruction modalities and their combinations to the participants. The modalities are (1) descriptive text breaking down the performance kinematics, (2) an ordered series of static images of a person performing the gesture with arrows annotating as appropriate, and (3) video (dynamic images) of a person performing the gesture. There are 30 participants in total and for each gesture, the data were collected as: Text (10 people), Images (10 people), Video (10 people), Video with text (10 people), Images with text (10 people). The dataset was captured using one Kinect™ sensor and only the skeleton data are made available.

2.1.5. MSRDailyActivity3D

MSRDailyActivity3D Dataset (Wang et al., 2012) (<http://www.uow.edu.au/~wanqing/#MSRAAction3DDatasets>) was collected by Microsoft and the Northwestern University in 2012 and focused on daily activities. The motivation was to cover human daily activities in the living room. The actions were performed by 10 actors while sitting on the sofa or standing close to the sofa. The camera was fixed in front of the sofa. In addition to depth data, skeleton data are also recorded, but the joint positions extracted by the tracker are very noisy due to the actors being either sitting on or standing close to the sofa.

2.1.6. UTKinect

UTKinect dataset (Xia et al., 2012) (<http://cvrc.ece.utexas.edu/KinectDatasets/HOJ3D.html>) was collected by the University of Texas at Austin in 2012. Ten types of human actions were performed twice by 10 subjects. The subjects performed the actions from a variety of views. One challenge of the dataset is due to the actions being performed with high actor-dependent variability. Furthermore, human-object occlusions and body parts being out of the field of view have further increased the difficulty of the dataset. Ground truth in terms of action labels and segmentation of sequences are provided.

2.1.7. G3D

Gaming 3D dataset (G3D) (Bloom et al., 2012) (<http://dipersec.king.ac.uk/G3D/>) captured by Kingston University in 2012 focuses on real-time action recognition in gaming scenario. It contains 10 subjects performing 20 gaming actions. Each subject performed these actions thrice. Two kinds of labels were provided as ground truth: the onset and offset of each action and the peak frame of each action.

2.1.8. SBU Kinect Interaction Dataset

SBU Kinect Interaction Dataset (Yun et al., 2012) (http://www3.cs.stonybrook.edu/~kyun/research/kinect_interaction/index.html) was collected by Stony Brook University in 2012. It contains eight types of interactions. All videos were recorded with the same indoor background. Seven participants were involved in performing the activities which have interactions between two actors. The dataset is segmented into 21 sets and each set contains one or two sequences of each action category. Two kinds of ground truth information are provided: action labels of each segmented video and identification of “active” actor and “inactive” actor.

2.1.9. Berkeley MHAD

Berkeley Multimodal Human Action Database (Berkeley MHAD) (Ofli et al., 2013) (http://tele-immersion.citris-uc.org/berkeley_mhad#dl), collected by University of California at Berkeley and Johns Hopkins University in 2013, was captured in five different modalities to expand the fields of application. The modalities are derived from: optical mocap system, four multi-view stereo vision cameras, two Microsoft Kinect v1 cameras, six wireless accelerometers and four microphones. Twelve subjects performed 11 actions, five times each. Three categories of actions are included: (1) actions with movement in full body parts, e.g., *jumping in place*, *jumping jacks*, *throwing*, etc., (2) actions with high dynamics in upper extremities, e.g., *waving hands*, *clapping hands*, etc. and (3) actions with high dynamics in lower extremities, e.g., *sit down*, *stand up*. The actions were executed with style and speed variations.

2.1.10. Northwestern-UCLA Multiview Action 3D

Northwestern-UCLA Multiview Action 3D (Wang et al., 2014) (http://users.eecs.northwestern.edu/~jwa368/my_data.html) was collected by Northwestern University and University of California at Los Angeles in 2014. This dataset contains data taken from a variety of viewpoints. The actions were performed by 10 actors and captured by three simultaneous Kinect™v1 cameras.

2.1.11. ChaLearn LAP IsoGD

ChaLearn LAP IsoGD Dataset (Wan et al., 2016b) (<http://www.cbsr.ia.ac.cn/users/jwan/database/isogd.html>) is a large RGB-D dataset for segmented gesture recognition, and it was collected by Kinect v1 camera. It includes 47,933 RGB-D depth sequences, each RGB-D video representing one gesture instance. There are 249 gestures performed by 21 different individuals. The dataset is divided into training, validation

and test sets. All three sets consist of samples of different subjects to ensure that the gestures of one subject in the validation and test sets will not appear in the training set.

2.1.12. NTU RGB+D

NTU RGB+D Dataset (Shahroudy et al., 2016) (<https://github.com/shahroudy/NTURGB-D>) is currently the largest action recognition dataset in terms of the number of samples per action. The RGB-D data is captured by Kinect v2 cameras. The dataset has more than 56 thousand sequences and 4 million frames, containing 60 actions performed by 40 subjects aging between 10 and 35. It consists of front view, two side views and left, right 45° views.

2.2. Continuous/online datasets

Continuous/online datasets refer to those datasets where each video sequence may contain one or more actions/gestures, and the segmented position between different motion classes are unknown. These datasets are mainly used for action detection, localization and online prediction. There are few datasets for this type.

2.2.1. ChaLearn2014 Multimodal Gesture Recognition

ChaLearn2014 Multimodal Gesture Recognition (Escalera et al., 2014) (<http://gesture.chalearn.org/2014-looking-at-people-challenge>) is multi-modal dataset collected by Kinect v1 sensor, including RGB, depth, skeleton and audio modalities. In all sequences, a single user is recorded in front of the camera, performing natural communicative Italian gestures. The starting and ending frames for each gesture are annotated along with the gesture class label. It contains nearly 14K manually labeled (beginning and ending frame) gesture performances in continuous video sequences, with a vocabulary of 20 Italian gesture categories. There are 1, 720, 800 labeled frames across 13, 858 video fragments of about 1–2 min sampled at 20 Hz. The gestures are performed by 27 different individuals under diverse conditions; these include varying clothes, positions, backgrounds and lighting.

2.2.2. ChaLearn LAP ConGD

The ChaLearn LAP ConGD Dataset (Wan et al., 2016b) (<http://www.cbsr.ia.ac.cn/users/jwan/database/congd.html>) is a large RGB-D dataset for continuous gesture recognition. It was collected by Kinect v1 sensor and includes 47,933 RGB-D gesture instances in 22,535 RGB-D gesture videos. Each RGB-D video may contain one or more gestures. There are 249 gestures performed by 21 different individuals. The dataset is divided into training, validation and test sets. All three sets consist of samples of different subjects to ensure that the gestures of one subject in the validation and test sets will not appear in the training set.

2.2.3. PKU-MMD

PKU-MMD (Chunhui et al., 2017) (<http://www.icst.pku.edu.cn/struct/Projects/PKUMMD.html>) is a large scale dataset for continuous multi-modality 3D human action understanding and covers a wide range of complex human activities with well annotated information. It was captured via the Kinect v2 sensor. PKU-MMD contains 1076 long video sequences in 51 action categories, performed by 66 subjects in three camera views. It contains almost 20,000 action instances and 5.4 million frames in total. It provides multi-modality data sources, including RGB, depth, Infrared Radiation and Skeleton.

Table 1 shows the statistics of publicly available benchmark datasets that are commonly used for evaluation of deep learning-based algorithms. It can be seen that the surveyed datasets cover a wide range of different types of actions including gestures, simple actions, daily activities, human–object interactions, human–human interactions. It also covers both segmented and continuous/online datasets, with different acquisition devices, modalities, and views. Sample images from different datasets are shown in Fig. 2. In this section, we introduce the deep learning concepts and architectures that are relevant or have been

applied to RGB-D-based motion recognition. Readers who are interested in more background and techniques are referred to the book by Goodfellow et al. (2016).

3. RGB-based motion recognition with deep learning

RGB is one important channel of RGB-D data with characteristics including shape, color and texture that bear rich features. These properties also make it effective to directly use networks, such as 2D CNNs (He et al., 2016; Krizhevsky et al., 2012; Simonyan and Zisserman, 2014b), to extract frame-level features. Even though most of the surveyed methods for this section are not adapted to RGB-D-based datasets, we argue that the following methods could be directly adapted to RGB modality of RGB-D datasets. We define three categories namely, CNN-based, RNN-based and other-architecture-based approaches for segmented motion recognition; the first two categories are for continuous/online motion recognition.

3.1. Segmented motion recognition

3.1.1. CNN-based approach

For this group of methods, currently there are mainly four approaches to encode spatial-temporal-structural information. The first approach applies CNN to extract features from individual frames and later, fuse the temporal information. For example Karpathy et al. (2014) investigated four temporal fusion methods, and proposed the concept of slow fusion where higher layers get access to progressively more global information in both spatial and temporal dimensions (see Fig. 4). Furthermore, several temporal pooling methods have been explored and the suggestion is that max pooling in the temporal domain is preferable (Ng et al., 2015).

The second approach extends convolutional operation into temporal domain. In one such implementation, Ji et al. (2013) proposed 3D-convolutional networks using 3D kernels (filters extended along the time axis) to extract features from both spatial and temporal dimensions. This work empirically showed that the 3D-convolutional networks outperform their 2D frame-based counterparts. With modern deep architectures, such as VGG (Simonyan and Zisserman, 2014b), and large-scale supervised training datasets, such as Sports-1M (Karpathy et al., 2014), Tran et al. (2015) extended the work presented in Ji et al. (2013) by including 3D pooling layers, and proposed a generic descriptor named C3D by averaging the outputs of the first fully connected layer of the networks (see Fig. 3). However, both of these works break the video sequence into short clips and aggregate video-level information by late score fusion. This is likely to be suboptimal when considering some long action sequence, such as walking or swimming that lasts several seconds and spans tens or hundreds of video frames. To handle this problem, Varol et al. (2018) investigated the learning of long-term video representations and proposed the Long-term Temporal Convolutions (LTC) at the expense of decreasing spatial resolution to keep the complexity of networks tractable. Despite the fact that this is straightforward and mainstream, extending spatial kernels to 3D spatio-temporal derivative inevitably increases the number of parameters of the network. To relieve the drawbacks of 3D kernels, Sun et al. (2015) factorized a 3D filter into a combination of 2D and 1D filters.

The third approach is to encode the video into dynamic images that contain the spatio-temporal information and then apply CNN for image-based recognition. Bilen et al. (2016) proposed to adopt rank pooling (Fernando et al., 2016b) to encode the video into one dynamic set of images and used pre-trained models over ImageNet (Krizhevsky et al., 2012) for fine-tuning (see Fig. 6). The end-to-end learning methods with rank pooling has also been proposed in Bilen et al. (2016) and Fernando and Gould (2016). Hierarchical rank pooling (Fernando et al., 2016a) is proposed to learn higher order and non-linear representations compared to the original work. Generalized rank pooling

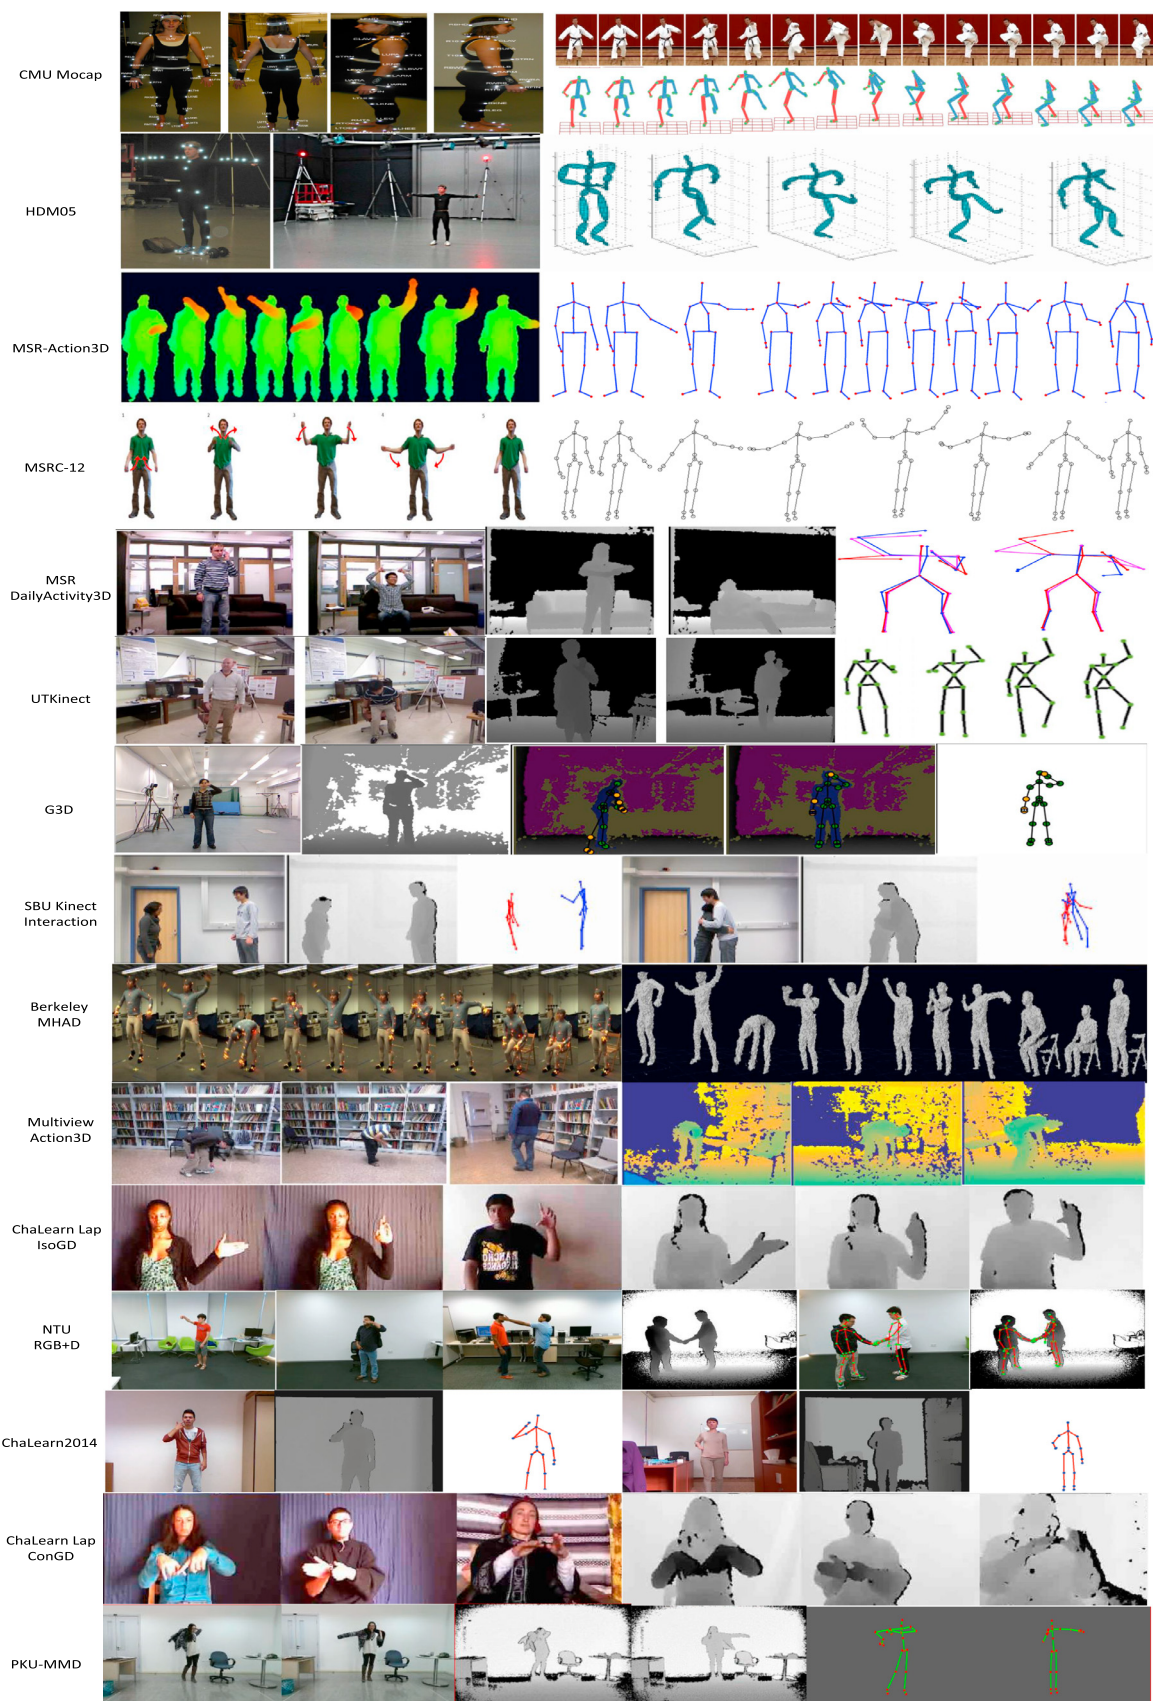


Fig. 2. Sample images from different datasets.




Fig. 3. C3D net has 8 convolutions, 5 max-poolings, and 2 fully connected layers, followed by a softmax output layer. All 3D convolution kernels are $3 \times 3 \times 3$ with stride 1 in both spatial and temporal dimensions. Number of filters are indicated in each box. The 3D pooling layers are as indicated from *pool1* to *pool5*. All pooling kernels are $2 \times 2 \times 2$, except for *pool1* which is $1 \times 2 \times 2$. Each fully connected layer has 4096 output units. Figure from Tran et al. (2015).




Fig. 4. Different approaches for fusing information over temporal dimension through the network. Red, green and blue boxes indicate convolutional, normalization and pooling layers respectively. In the Slow Fusion model, the depicted columns share parameters. Figure from Karpathy et al. (2014). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(Cherian et al., 2017) is introduced to improve the original method via a quadratic ranking function which jointly provides a low-rank approximation to the input data and preserves their temporal order in a subspace.

Besides the above works that aim to adopt one network to exploit both spatio-temporal information contained in the video, the fourth approach separates the two factors and adopt multiple stream networks. Simonyan and Zisserman (2014a) proposed one spatial stream network fed with raw video frames, and one temporal stream network accepting optical flow fields as input. The two streams are fused together using the softmax scores (see Fig. 5 for the two-stream architecture). Wang et al. (2015a) extended the two-stream networks by integrating improved trajectories (Wang et al., 2013), where trajectory-constrained sampling and pooling are used to encode deep features learned from deep CNN architecture, into effective descriptors. To incorporate long-range temporal structure using the two-stream networks, Wang et al. (2016a) devised a temporal segment network (TSN) that uses a sparse sampling scheme to extract short snippets over a long video sequence. With the removal of redundancy from consecutive frames and a segmental structure, aggregated information is obtained from the sampled snippets. To reduce the expensive calculation of optical flow, Zhang et al. (2016a) accelerated this two stream structure by replacing optical flow with motion vector which can be obtained directly from compressed videos without extra calculation. Wang et al. (2016g) leveraged semantic cues in video by using a two-stream semantic region-based CNNs (SR-CNNs) to incorporate human/object detection results into the framework. In their work, Chéron et al. (2015) exploit spatial structure of the human pose and extract a pose-based convolutional neural network (P-CNN) feature

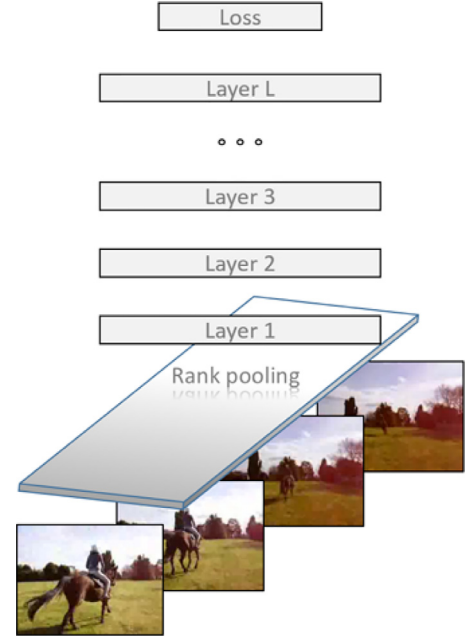


Fig. 6. Rank pooling encodes the RGB video into one dynamic image and CNN is adopted for feature extraction and classification. Figure from Bilen et al. (2016).

from both RGB frames and optical flow for fine-grained action recognition. The work presented in Wang et al. (2016f) formulated the problem of action recognition from a new perspective and model an action as a transformation which changes the state of the environment before the action to the state after the action. They designed a Siamese network which models the action as a transformation on a high-level feature space based on the two-stream model. Based on the two-stream framework, Zhu et al. (2016c) proposed a key volume mining deep framework for action recognition, where they identified key volumes and conducted classification simultaneously. Inspired by the success of Residual Networks (ResNets) (He et al., 2016), Feichtenhofer et al. (2016a) injected residual connections between the two streams to allow spatial-temporal interaction between them. Instead of using optical flow for temporal stream, Lea et al. (2016) adopted Motion History Image (MHI) (Bobick and Davis, 2001) as the motion clue. The MHI was combined with RGB frames in a spatio-temporal CNN for fine-grained action recognition. However, all the

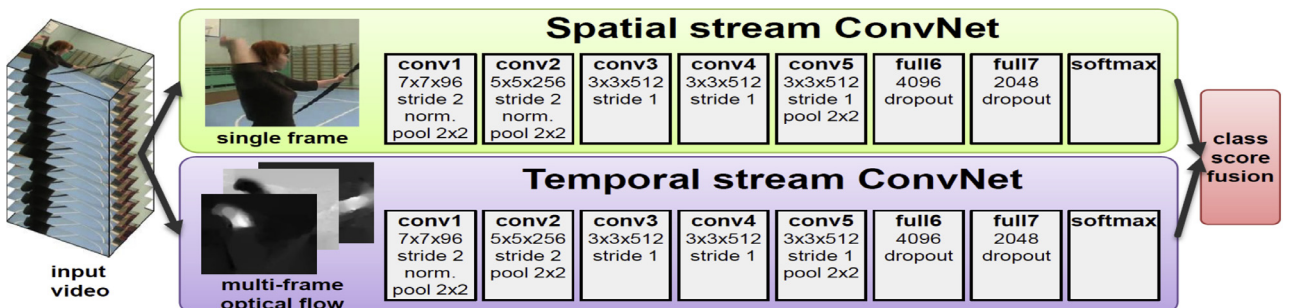


Fig. 5. Two-stream architecture for RGB-based motion recognition. Figure from Simonyan and Zisserman (2014a).

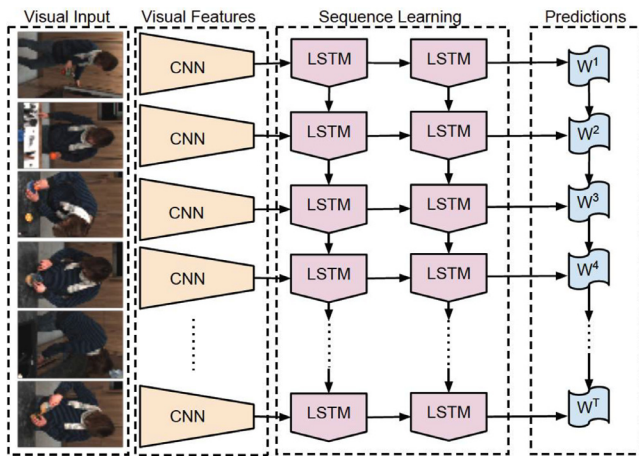


Fig. 7. LRCN processes the variable-length visual input with a CNN, whose outputs are fed into a stack of recurrent sequence models. The output is a variable-length prediction. Figure from Donahue et al. (2015).

methods reviewed above incorporated the two streams from separate training regimes; any registration of the two streams was neglected. In order to address this gap and propose a new architecture for spatial-temporal fusion of the two streams Feichtenhofer et al. (2016b) investigated three aspects of fusion for the two streams: (i) how to fuse the two networks with consideration for spatial registration, (ii) where to fuse the two networks and, (iii) how to fuse the networks temporally. One of their conclusions was that the results suggest the importance of learning correspondences between highly abstract ConvNet features both spatially and temporally.

3.1.2. RNN-based approach

For RNN-based approach, Baccouche et al. (2011) tackled the problem of action recognition through a cascade of 3D CNN and LSTM, in which the two networks were trained separately. Differently from the separate training, Donahue et al. (2015) proposed one Long-term Recurrent Convolutional Network (LRCN) to exploit end-to-end training of the two networks (see illustration in Fig. 7). To take full advantage of both CNN and RNN, Ng et al. (2015) aggregated CNN features with both temporal pooling and LSTM for temporal exploitation, and fused the output scores from the feature pooling and LSTM network to conduct final action recognition. Pigou et al. (2016) proposed an end-to-end trainable neural network architecture incorporating temporal convolutions and bidirectional LSTM for gesture recognition. This provided opportunity to mine temporal information that is much discriminative for gesture recognition. In their work, Sharma et al. (2016) proposed a soft attention model for action recognition based on LSTM (see Fig. 8). The attention model learns the parts in the frames that are relevant for the task at hand and attaches higher importance to them. Previous attention-based methods have only utilized video-level category as supervision to train LSTM. This strategy may lack a detailed and dynamical guidance and consequently restrict their capacity for modelling complex motions in videos. Du et al. (2017) address this problem by proposing a recurrent pose-attention network (RPAN) for action recognition in videos, which can adaptively learn a highly discriminative pose-related feature for every-step action prediction of LSTM. To take advantage of both Fisher Vector (Sánchez et al., 2013) and RNN, Lev et al. (2016) introduced a Recurrent Neural Network Fisher Vector (RNN-FV) where the GMM probabilistic model in the fisher vector is replaced by a RNN and thus avoids the need for the assumptions of data distribution in the GMM. Even though RNN is remarkably capable of modeling temporal dependences, it lacks an intuitive high-level spatial-temporal structure. The spatio-temporal structural information has been mined by Jain et al. (2016) through a combination of the powers of spatio-temporal graphs and RNN for

action recognition. Recently, Sun et al. (2017) proposed a Lattice-LSTM (L2STM) network, which extends LSTM by learning independent hidden state transitions of memory cells for individual spatial locations. This method effectively enhances the ability to model dynamics across time and addresses the non-stationary issue of long-term motion dynamics without significantly increasing the model complexity. Differently from previous methods that using only feedforward connections, Shi et al. (2017) proposed a biologically-inspired deep network, called ShuttleNet1. Unlike traditional RNNs, all processors inside ShuttleNet are connected in a loop to mimic the human brain's feedforward and feedback connections. In this manner, the processors are shared across multiple pathways in the loop connection. Attention mechanism is then employed to select the best information flow pathway.

3.1.3. Other-architecture-based approach

Besides the commonly used CNN- and RNN-based methods for motion recognition from RGB modality, there are several other structures that have been adopted for this task. Jhuang et al. (2007) used a feedforward hierarchical template matching architecture for action recognition with pre-defined spatio-temporal filters in the first layer. In his thesis, Chen (2010) adopted the convolutional RBM (CRBM) as the basic processing unit and proposed the so-called space-time Deep Belief Network (ST-DBN) that alternates the aggregation of spatial and temporal information so that higher layers capture longer range statistical dependencies in both space and time. Taylor et al. (2010) extended the Gated RBM (GRBM) (Memisevic and Hinton, 2007) to convolutional GRBM (convGRBM) that shares weights at all locations in an image and inference is performed through convolution. Le et al. (2011) presented an extension of the independent subspace analysis algorithm (Theis, 2007) to learn invariant spatio-temporal features from unlabeled video data. They scale up the original ISA to larger input data by employing two important ideas from convolutional neural networks: convolution and stacking. This convolutional stacking idea enables the algorithm to learn a hierarchical representation of the data suitable for recognition. Yan et al. (2014) proposed Dynencoder, a three layer auto-encoder, to capture video dynamics. Dynencoder is shown to be successful in synthesizing dynamic textures, and one can think of a Dynencoder as a compact way of representing the spatio-temporal information of a video. Similarly, Srivastava et al. (2015) introduced a LSTM autoencoder model as shown in Fig. 10. The LSTM autoencoder model consists of two RNNs, namely, the encoder LSTM and the decoder LSTM. The encoder LSTM accepts a sequence as input and learns the corresponding compact representation. The states of the encoder LSTM contain the appearance and dynamics of the sequence. The decoder LSTM receives the learned representation to reconstruct the input sequence. Inspired by the Generative Adversarial Networks (GAN) (Goodfellow et al., 2014), Mathieu et al. (2016) adopted the adversarial mechanism to train a multi-scale convolutional network to generate future frames given an input sequence. To deal with the inherently blurry predictions obtained from the standard Mean Squared Error (MSE) loss function, they proposed three different and complementary feature learning strategies: a multi-scale architecture, an adversarial training method, and an image gradient difference loss function.

3.2. Continuous/online motion recognition

Most of the action recognition methods reviewed above rely heavily on segmented videos for model training. However, it is very expensive and time-consuming to acquire a large-scale trimmed video dataset. The availability of untrimmed video datasets (e.g. Chunhui et al., 2017; De Geest et al., 2016; Heilbron et al., 2015; Wan et al., 2016b; Yeung et al., 2017) have encouraged research and challenges/contests in motion recognition in this domain.

3.2.1. CNN-based approach

Inspired by the success of region-proposal-based object detection

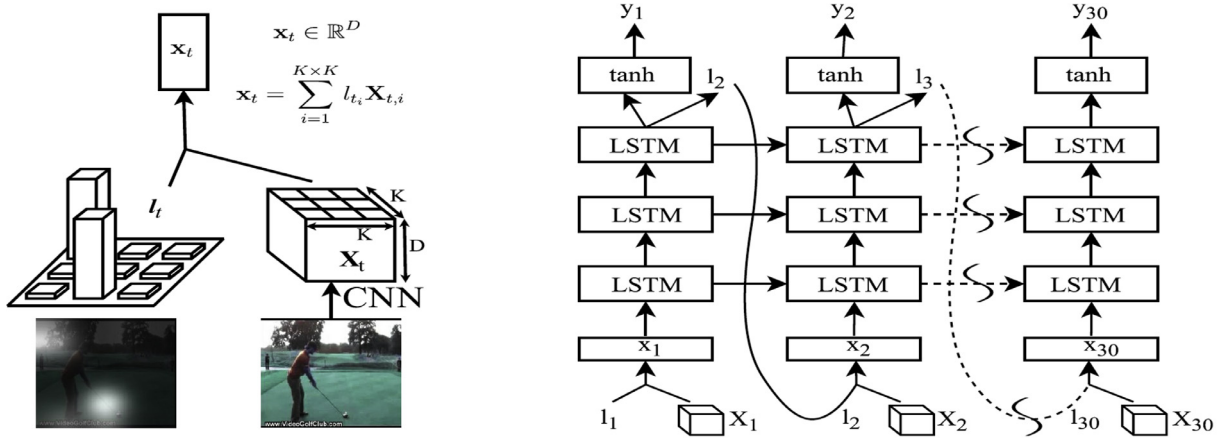


Fig. 8. The CNN takes the video frames as its input and produces a feature tube. The model computes the current input \mathbf{x}_t as an average of the feature slices weighted according to the location softmax l_t . At each time-step t , the recurrent network takes a feature slice \mathbf{x}_t as input. It then propagates \mathbf{x}_t through three layers of LSTMs and predicts the next location probabilities l_{t+1} and the class label y_t . Figure from Sharma et al. (2016).

using R-CNN (Girshick, 2015; Girshick et al., 2014; Ren et al., 2015), several proposal-based action recognition methods from untrimmed video are proposed. These methods first generate a reduced number of candidate temporal windows, and then an action classifier discriminates each proposal independently into one of the actions of interest. For instance, based on the two-stream concept (Simonyan and Zisserman, 2014a), Gkioxari and Malik (2015) classified frame-based region proposals of interest using static and motion cues. The regions are then linked across frames based on the predictions and their spatial overlap; thus producing action tubes respectively for each action and video. Weinzaepfel et al. (2015) also started from the frame-level proposals, selected the highest scoring ones, tracked them throughout the video, and adopted a multi-scale sliding window approach over tracks to detect the temporal content of an action. Shou et al. (2016) proposed a multi-stage segment-based 3D CNNs to generate candidate segments, that are used to recognize actions and localize temporal boundaries. Peng and Schmid (2016) generated rich proposals from both RGB and optical flow data by using region proposal networks for frame-level action detection, and stacked optical flows to enhance the discriminative power of motion R-CNN. Wang et al. (2017a) proposed an UntrimmedNet to generate clip proposals that may contain action instances for untrimmed action recognition. Based on these clip-level representations, the classification module aims to predict the scores for each clip proposal and the selection module tries to select or rank those clip proposals. Similarly in the same direction, Zhao et al. (2017) adopted explicit structural modeling in the temporal dimension. In their model, each complete activity instance is considered as a composition of three major stages, namely *starting*, *course*, and *ending*, and they introduced structured temporal pyramid pooling to produce a global representation of the entire proposal. Differently from previous methods, Zhu et al. (2017b) proposed a framework that integrates the complementary spatial and temporal information into an end-to-end trainable system for video action proposal, and a novel and efficient path trimming method is proposed to handle untrimmed video by examining actionness and background score pattern without using extra detectors. To generalize R-CNN from 2D to 3D, Hou et al. (2017a) proposed an end-to-end 3D CNN-based approach for action detection in videos. A Tube Proposal Network was introduced to leverage skip pooling in temporal domain to preserve temporal information for action localization in 3D volumes, and Tube-of-Interest pooling layer was proposed to effectively alleviate the problem with variable spatial and temporal sizes of tube proposals. Saha et al. (2017) proposed a deep net framework capable of regressing and classifying 3D region proposals spanning two successive video frames. The core of the framework is an evolution of classical region proposal networks (RPNs) to 3D RPNs.

Similarly, Kalogeiton et al. (2017) extended the Single Shot MultiBox Detector (SSD) (Liu et al., 2016b) framework from 2D to 3D by proposing an Action Tubelet detector. In order to quickly and accurately generate temporal action proposals, Gao et al. (2017) proposed a Temporal Unit Regression Network (TURN) model, that jointly predicts action proposals and refines the temporal boundaries by temporal coordinate regression using CNN. Similarly, Singh et al. (2017) designed an efficient online algorithm to incrementally construct and label ‘action tubes’ from the SSD frame level detections, making it the first real-time (up to 40fps) system able to perform online S/T action localisation on the untrimmed videos. Besides the proposal-based methods discussed above, Lea et al. (2017) introduced a new class of temporal models, called Temporal Convolutional Networks (TCNs), that use a hierarchy of temporal convolutions to perform fine-grained action segmentation or detection.

3.2.2. RNN-based approach

Apart from the proposal-based methods that use CNN, there are several proposal-based methods using RNN for temporal modeling. Escorcia et al. (2016) introduced the Deep Action Proposals (DAPs) that generate temporal action proposals from long untrimmed videos for action detection and classification (see Fig. 9). They adopted C3D network (Tran et al., 2015) for visual encoder and LSTM for sequence encoder. However, all of these methods generated proposals by a sliding window approach, dividing the video into short overlapping temporal window, which is computationally expensive. To reduce the number of proposals, Buch et al. (2017) proposed a single-stream temporal action proposal generation method that does not the need to divide input into short overlapping clips or temporal windows for batch processing.

Besides the proposal-based methods, there are several methods that are proposal-free. Yeung et al. (2016) proposed an end-to-end training model that is formulated as a recurrent neural network-based agent. This agent learns a policy for sequentially forming and refining hypotheses about action instances based on the intuition that the process of detecting actions is naturally one of observation and refinement. They adopted two networks namely, observation network and recurrent network, for this purpose. Singh et al. (2016) presented a multi-stream bi-directional recurrent neural network for fine-grained action detection. They adopted a tracking algorithm to locate a bounding box around the person and trained two streams on motion and appearance cropped to the tracked bounding box. The video sequence was split into fixed long chunks for the input of two-stream networks, and bi-directional LSTM was used to model long-term temporal dynamics within and between actions. Ma et al. (2016) introduced a novel ranking loss

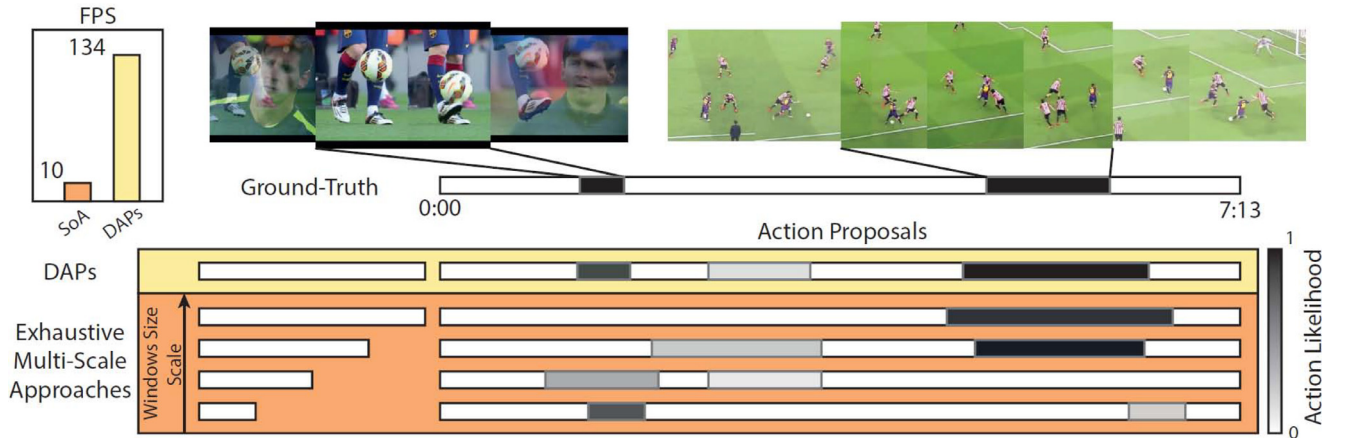


Fig. 9. The Deep Action Proposals algorithm can localize segments of varied duration around actions occurring along a video without exhaustively exploring multiple temporal scales. Figure from Escorcia et al. (2016).

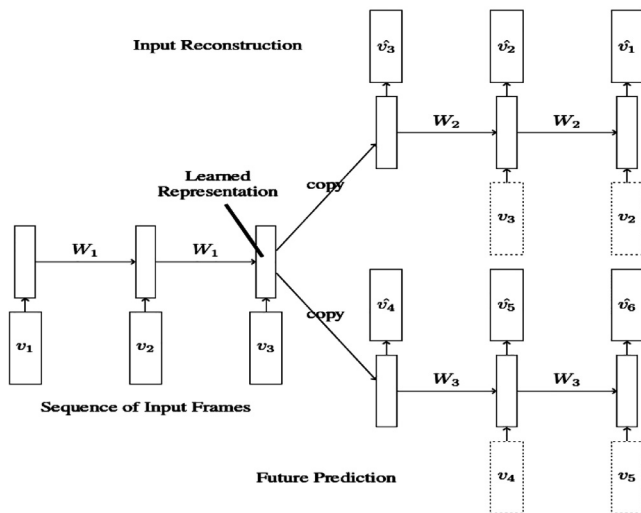


Fig. 10. The LSTM autoencoder model and LSTM future predictor model. Figure from Srivastava et al. (2015).

within the RNN objective so that the trained model better captures progression of activities. The ranking loss constrains the detection score of the correct category to be monotonically non-decreasing as the activity progresses. The same time, the detection score margin between the correct activity category and all other categories is monotonically non-decreasing. Huang et al. (2016) proposed a weakly-supervised

framework for action labeling in video (see Fig. 11), where only the order of occurring actions is required during training. They proposed an Extended Connectionist Temporal Classification (ECTC) framework to efficiently evaluate all possible alignments between the input and label sequences via dynamic programming and explicitly enforce their consistency with frame-to-frame visual similarities. Taking inspiration from classical linear dynamic systems theory for modeling time series, Dave et al. (2017) derived a series of recurrent neural networks that sequentially make top-down predictions about the future and then correct those predictions with bottom-up observations. Their predictive-corrective architecture allows the incorporation of insights from time-series analysis: adaptively focus computation on “surprising” frames where predictions require large corrections; simplify learning in that only “residual-like” corrective terms need to be learned over time and naturally decorrelate an input stream in a hierarchical fashion, producing a more reliable signal for learning at each layer of a network.

4. Depth-based motion recognition with deep learning

Compared with RGB videos, the depth modality is insensitive to illumination variations, invariant to color and texture changes, reliable for estimating body silhouette and skeleton, and provides rich 3D structural information of the scene. However, there are only few published results on depth based action recognition using deep learning methods. Two reasons can be adduced for this situation. First, the absence of color and texture in depth maps weakens the discriminative representation power of CNN models (Liu et al., 2016c). Second, existing depth data is relatively small-scale. The conventional pipelines

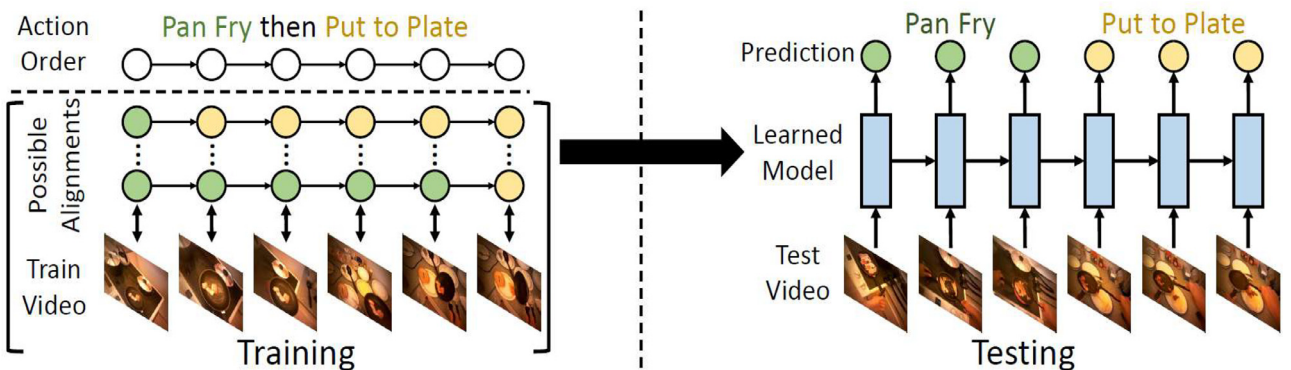


Fig. 11. The problem of weakly supervised action labelling is tackled where only the order of the occurring actions is given during training. The temporal model is trained by maximizing the probability of all possible frame-to-label alignments. At testing time, no annotation is given. The learned model encodes the temporal structure of videos and could predict the actions without further information. Figure from Huang et al. (2016).

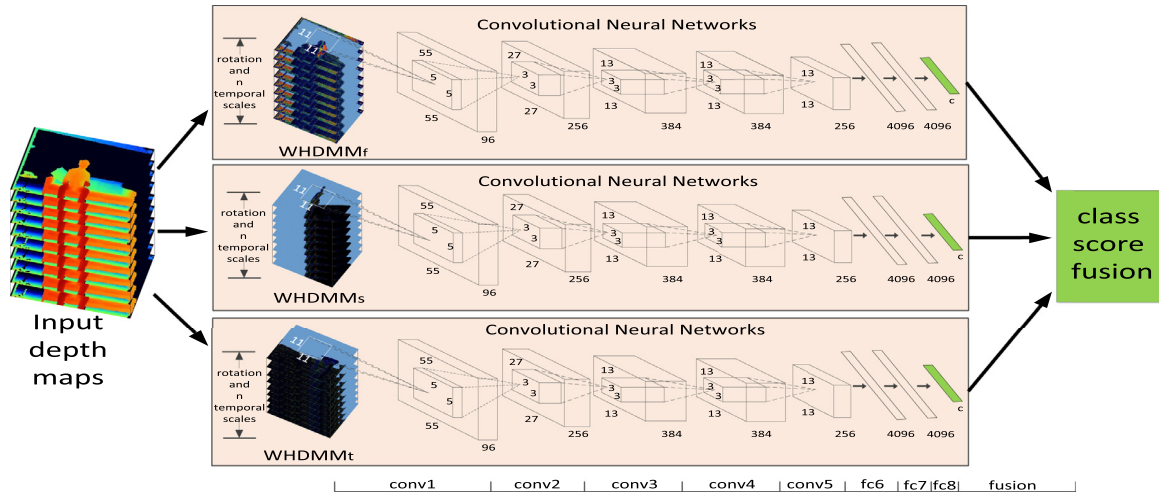


Fig. 12. Depth map sequences are encoded into texture color images by using the concepts of Depth Motion Maps (DMM) (Yang et al., 2012) and pseudo-coloring, and at the same time enlarged the training data by scene rotation on the 3D point cloud. Three channel of CNN are adopted for feature extraction and classification. Figure from Wang et al. (2016b).




Fig. 13. The JTM framework for skeleton-based motion recognition with CNN. Figure from Wang et al. (2016e).

are purely data-driven and learn representation directly from the pixels. Such model is likely to be at risk of overfitting when the network is optimized on limited training data. Currently, there are only CNN-based methods for depth-based motion recognition.

4.1. Segmented motion recognition

4.1.1. CNN-based approach

Wang et al. (2015b, 2016b) proposed a method called Weighted Hierarchical Depth Motion Maps (WHDMM) + 3ConvNet, as shown in Fig. 12 for human action recognition from depth maps on small training datasets. Three strategies were developed to leverage the capability of ConvNets in mining discriminative features for recognition. Firstly, different viewpoints are mimicked by rotating the 3D points of the captured depth maps. This not only augments the data, but also makes the trained ConvNets view-tolerant. Secondly, WHDMMs at several temporal scales were constructed to encode the spatio-temporal motion patterns of actions into 2D spatial structures. The 2D spatial structures are further enhanced for recognition by converting the WHDMMs into pseudo-color images. Lastly, the three ConvNets were initialized with the models obtained from ImageNet and fine-tuned independently on the color-coded WHDMMs constructed in three orthogonal planes. Inspired by the promising results achieved by rank pooling method (Bilen et al., 2016) on RGB data, Wang et al. (2016c) encoded the depth map sequences into three kinds of dynamic images with rank pooling: Dynamic Depth Images (DDI), Dynamic Depth Normal Images (DDNI) and Dynamic Depth Motion Normal Images (DDMNI). These three representations capture the posture and motion information from three

different levels for gesture recognition. Specifically, DDI exploits the dynamics of postures over time and DDNI and DDMNI exploit the 3D structural information captured by depth maps. Wang et al. (2018a) replaced the bidirectional rank pooling in the method of Wang et al. (2016c) with *hierarchical* and *bidirectional* rank pooling to capture both high order and non-linear dynamics effectively for both gesture and action recognition. Recently, Wang et al. (2017c) proposed to represent a depth map sequence into three pairs of structured dynamic images at body, part and joint levels respectively through bi-directional rank pooling. Different from previous works that applied one ConvNet for each part/joint separately, one pair of structured dynamic images is constructed from depth maps at each granularity level and serves as the input of a ConvNet. The structured dynamic image not only preserves the spatial-temporal information but also enhances the structure information across both body parts/joints and at different temporal scales. In addition, it requires low computational cost and memory to construct. This new representation, referred to as Spatially Structured Dynamic Depth Images (S²DDI), aggregates from global to fine-grained motion and structure information in a depth sequence, and enables us to fine-tune the existing ConvNet models trained on image data for classification of depth sequences, without a need for training the models afresh. Similarly, Hou et al. (2017b) extended S²DDI to Spatially and Temporally Structured Dynamic Depth Images (STSDDI), where a hierarchical bidirectional rank pooling method was adopted to exploit the spatio-temporal-structural information contained in the depth sequence and it is applied to interactions of two subjects. Differently from the above texture image encoding method, Rahmani and Mian (2016) proposed a cross-view action recognition based on depth sequence. Their method comprises two steps: (i) learning a general view-invariant human pose model from synthetic depth images and, (ii) modeling the temporal action variations. To enlarge the training data for CNN, they generated the training data synthetically by fitting realistic synthetic 3D human models to real mocap data and then rendering each pose from a large number of viewpoints. For spatio-temporal representation, they used group sparse Fourier Temporal Pyramid which encodes the action-specific discriminative output features of the proposed human pose model.

4.2. Continuous/online motion recognition

4.2.1. CNN-based approach

For continuous gesture recognition, Wang et al. (2016d) first segmented the continuous depth sequence into segmented sequences using

quantity of movement (QOM) (Jiang et al., 2015), and then adopted improved DMM (IDMM) to encode the dynamics of depth sequences into texture images for large-scale continuous gesture recognition. To improve the encoding quality of depth sequences, Wang et al. (2018a) proposed three simple, compact yet effective representations of depth sequences, referred to respectively as Dynamic Depth Images (DDI), Dynamic Depth Normal Images (DDNI) and Dynamic Depth Motion Normal Images (DDMNI), for continuous action recognition. These dynamic images are constructed from a segmented sequence of depth maps using hierarchical bidirectional rank pooling to effectively capture the spatio-temporal information. Specifically, DDI exploits the dynamics of postures over time while DDNI and DDMNI extract the 3D structural information captured by depth maps. The image-based representations enable us to fine-tune the existing ConvNet models trained on image data without training a large number of parameters from scratch.

5. Skeleton-based motion recognition with deep learning

Differently from RGB and depth, skeleton data contains the positions of human joints, which can be considered relatively high-level features for motion recognition. There are two common ways to estimate skeletons, one is to use MOCAP systems and the other is to estimate skeletons directly from depth maps or RGB images/video. Skeletons from MOCAP systems are often robust to scale and illumination changes and can be invariant to viewpoints as well as human body rotation and motion speed; Skeletons estimated from depth maps or RGB images/video are prone to errors caused by a number of factors including viewpoints and occlusion since both factors can lead to significant different appearance of same actions. Currently, there are mainly three approaches to skeleton-based motion recognition using deep learning: (i) RNN-based, (ii) CNN-based and other-architecture-based approaches for segmented motion recognition and, (iii) RNN-based approaches for continuous/online motion recognition.

5.1. Segmented motion recognition

5.1.1. CNN-based approach

The main step in this approach is to convert the skeleton sequences into images where the spatio-temporal information is reflected in the image properties including color and texture. Du et al. (2015a) represented a skeleton sequence as a matrix by concatenating the joint coordinates at each instant and arranging the vector representations in a chronological order. The matrix is then quantified into an image and normalized to handle the variable-length problem. The final image is fed into a CNN model for feature extraction and recognition. Wang et al. (2016e) proposed to encode spatio-temporal information contained in the skeleton sequence into multiple texture images, namely, Joint Trajectory Maps (JTM), as shown in Fig. 13 by mapping the trajectories into HSV (hue, saturation, value) space. Pre-trained models over Imagenet is adopted for fine-tuning over the JTMs to extract features and recognize actions. Similarly, Hou et al. (2016) drew the skeleton joints with a specific pen to three orthogonal canvases, and encodes the dynamic information in the skeleton sequences with color encoding. Li et al. (2017a) proposed to encode the pair-wise distances of skeleton joints of single or multiple subjects into texture images, namely, Joint Distance Maps (JDM), as the input of CNN for action recognition. Compared with the works reported by Wang et al. (2016e) and Hou et al. (2016), JDM is less sensitive to view variations. Liu et al. (2017b) introduced an enhanced skeleton visualization method to represent a skeleton sequence as a series of visual and motion enhanced color images. They proposed a sequence-based view invariant transform to deal with the view variation problem, and multi-stream CNN fusion method is adopted to conduct recognition. Ke et al. (2017a) designed vector-based features for each body part of human skeleton sequences, which are translation, scale and rotation

invariant, and transformed the features into images to feed into CNN for learning high level and discriminative representation. In another effort, Ke et al. (2017b) represented the sequence as a clip with several gray images for each channel of the 3D coordinates, which reflects multiple spatial structural information of the joints. The images are fed to a deep CNN to learn high-level features, and the CNN features of all the three clips at the same time-step are concatenated in a feature vector. Each feature vector represents the temporal information of the entire skeleton sequence and one particular spatial relationship of the joints. A Multi-Task Learning Network (MTLN) is adopted to jointly process the feature vectors of all time-steps in parallel for action recognition. Kim and Reiter (2017) approached the problem differently and proposed to use the Temporal Convolutional Neural Networks (TCN) (Lea et al., 2017) for skeleton based action recognition. They re-designed the original TCN into Res-TCN by factoring out the deeper layers into additive residual terms that yields both interpretable hidden representations and model parameters.

5.1.2. RNN-based approach

In this class of approaches, skeleton features are input to an RNN in order to exploit the temporal evolution. For instance, in a series of works Du et al. (2016, 2015b) divided the whole skeleton sequence into five parts according to the human physical structure, and separately fed them into five bidirectional RNNs/LSTMs. As the number of layers increases, the representations extracted by the subnets are hierarchically fused to build a higher-level representation. The process is illustrated in Fig. 14. This method explicitly encodes the spatio-temporal-structural information into high level representation. Veeriah et al. (2015) proposed a differential gating scheme for the LSTM neural network, which emphasizes the change in information gain caused by the salient motions between the successive frames. This work is one of the first aimed at demonstrating the potential of learning complex time-series representations via high-order derivatives of states. Zhu et al. (2016d) designed two types of regularizations to learn effective features and motion dynamics. In the fully connected layers, they introduced regularization to drive the model to learn co-occurrence features of the joints at different layers. Furthermore, they derived a new dropout and apply it to the LSTM neurons in the last LSTM layer, which helps the network to learn complex motion dynamics. Instead of keeping a long-term memory of the entire body's motion in the cell, Shahroudy et al. (2016) proposed a part-aware LSTM human action learning model (P-LSTM) wherein memory is split across part-based cells. It is argued that keeping the context of each body part independent and representing the output of the P-LSTM unit as a combination of independent body part context information is more efficient. Previous RNN-based 3D-action recognition methods have adopted RNN to model the long-term contextual information in the temporal domain for motion-based dynamics representation. However,




Fig. 14. The LSTM autoencoder model and LSTM future predictor model. Figure from Du et al. (2015b).

there is also strong dependency between joints in the spatial domain. In addition the spatial configuration of joints in video frames can be highly discriminative for 3D-action recognition task. To exploit this dependency, Liu et al. (2016a) proposed a spatio-temporal LSTM (ST-LSTM) network which extends the traditional LSTM-based learning to both temporal and spatial domains. Rather than concatenate the joint-based input features, ST-LSTM explicitly models the dependencies between the joints and applies recurrent analysis over spatial and temporal domains concurrently. Besides, they introduced a trust gate mechanism to make LSTM robust to noisy input data. Song et al. (2017) proposed a spatio-temporal attention model with LSTM to automatically mine the discriminative joints and learn the respective and different attentions of each frame along the temporal axis. Similarly, Liu et al. (2017a) proposed a Global Context-Aware Attention LSTM (GCA-LSTM) to selectively focus on the informative joints in the action sequence with the assistance of global context information. Differently from previous works that adopted the coordinates of joints as input, Zhang et al. (2017c) investigated a set of simple geometric features of skeleton using 3-layer LSTM framework, and showed that using joint-line distances as input requires less data for training. Based on the notion that LSTM networks with various time-step sizes can model various attributes well, Lee et al. (2017) proposed an ensemble Temporal Sliding LSTM (TS-LSTM) networks for skeleton-based action recognition. The proposed network is composed of multiple parts containing short-term, medium-term and long-term TS-LSTM networks, respectively. Li et al. (2017b) proposed an adaptive and hierarchical framework for fine-grained, large-scale skeleton-based action recognition. This work was motivated by the need to distinguish fine-grained action classes that are intractable using a single network, and adaptivity to new action classes by model augmentation. In the framework, multiple RNNs are effectively incorporated in a tree-like hierarchy to mitigate the discriminative challenge and thus using a divide-and-conquer strategy. To deal with large view variations in captured human actions, Zhang et al. (2017b) proposed a self-regulated view adaptation scheme which re-positions the observation viewpoints dynamically, and integrated the proposed view adaptation scheme into an end-to-end LSTM network which automatically determines the “best” observation viewpoints during recognition.

5.1.3. Other-architecture-based approach

Besides the RNN- and CNN-based approaches, there are several other deep learning-based methods. Salakhutdinov et al. (2013) proposed a new compositional learning architecture that integrates deep learning models with structured hierarchical Bayesian models. Specifically, this method learns a hierarchical Dirichlet process (HDP) (Teh et al., 2004) prior over the activities of the top-level features in a deep Boltzmann machine (DBM). This compound HDP-DBM model learns novel concepts from very few training examples by learning: (i) low-level generic features, (ii) high-level features that capture correlations among low-level features and, (iii) a category hierarchy for sharing priors over the high-level features that are typical of different kinds of concepts. Wu and Shao (2014) adopted deep belief networks (DBN) to model the distribution of skeleton joint locations and extract high-level features to represent humans at each frame in 3D space. Ijina and Krishna Mohan (2016) adopted stacked auto encoder to learn the underlying features of input skeleton data. Huang et al. (2017) incorporated the Lie group structure into a deep learning architecture to learn more appropriate Lie group features for skeleton based action recognition (see Fig. 16).

5.2. Continuous/online motion recognition

5.2.1. RNN-based approach

Differently from previous methods that recognize motion from segmented skeleton sequences, Li et al. (2016a) proposed a multi-task end-to-end Joint Classification-Regression Recurrent Neural Network to

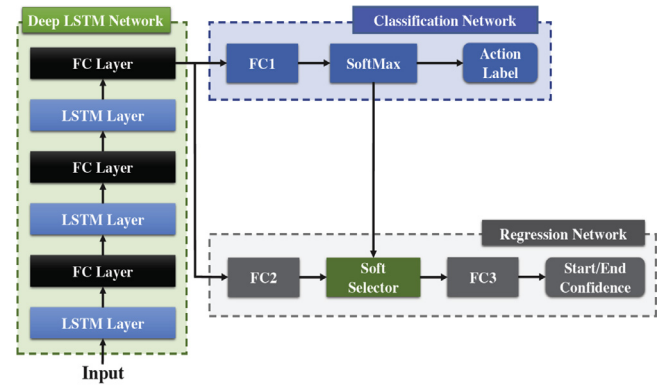


Fig. 15. The joint classification-regression RNN framework for online action detection and forecasting. Figure from Li et al. (2016a).

explore the action type and temporal localization information. They adopted LSTM to capture the complex long-range temporal dynamics, which avoids the typical sliding window design and thus ensures high computational efficiency. Furthermore, the subtask of regression optimization provides the ability to forecast the action prior to its occurrence. The framework is shown in Fig. 15.

6. RGB + D-based motion recognition with deep learning

As discussed in previous sections, RGB, depth and skeleton modalities have their own specific properties, and how to combine the strengths of these modalities with deep learning approach is important. To address this problem, several methods have been proposed. In general, these methods can be categorized as (i) CNN-based, (ii) RNN-based and other-architecture-based approaches for segmented motion recognition and, (iii) RNN-based continuous/online motion recognition.

6.1. Segmented motion recognition

6.1.1. CNN-based approach

Zhu et al. (2016b) fused RGB and depth in a pyramidal 3D convolutional network based on C3D (Tran et al., 2015) for gesture recognition. They designed pyramid input and pyramid fusion for each modality and late score fusion was adopted for final recognition. Duan et al. (2016) proposed a convolutional two-stream consensus voting network (2SCVN) which explicitly models both the short-term and long-term structure of the RGB sequences. To alleviate distractions from background, a 3D depth-saliency ConvNet stream (3DDSN) was aggregated in parallel to identify subtle motion characteristics. Later score fusion was adopted for final recognition. The methods described so far considered RGB and depth as separate channels and fused them later. Wang et al. (2017b) took a different approach and adopted scene flow to extract features that fused the RGB and depth from the onset. The new representation based on CNN and named Scene Flow to Action Map (SFAM) was used for motion recognition, as shown in Fig. 18. Different from previous methods, Wang et al. (2018b) proposed to co-operatively train a single convolutional neural network (named c-ConvNet) on both RGB and depth features, and deeply aggregate the two kinds of features for action recognition. While the conventional ConvNet learns the deep separable features for homogeneous modality-based classification with only one softmax loss function, the c-ConvNet enhances the discriminative power of the deeply learned features and weakens the undesired modality discrepancy by jointly optimizing a ranking loss and a softmax loss for both homogeneous and heterogeneous modalities. Rahmani and Bennamoun (2017) proposed an end-to-end learning model for action recognition from depth and skeleton data. The proposed model learned to fuse features from depth and

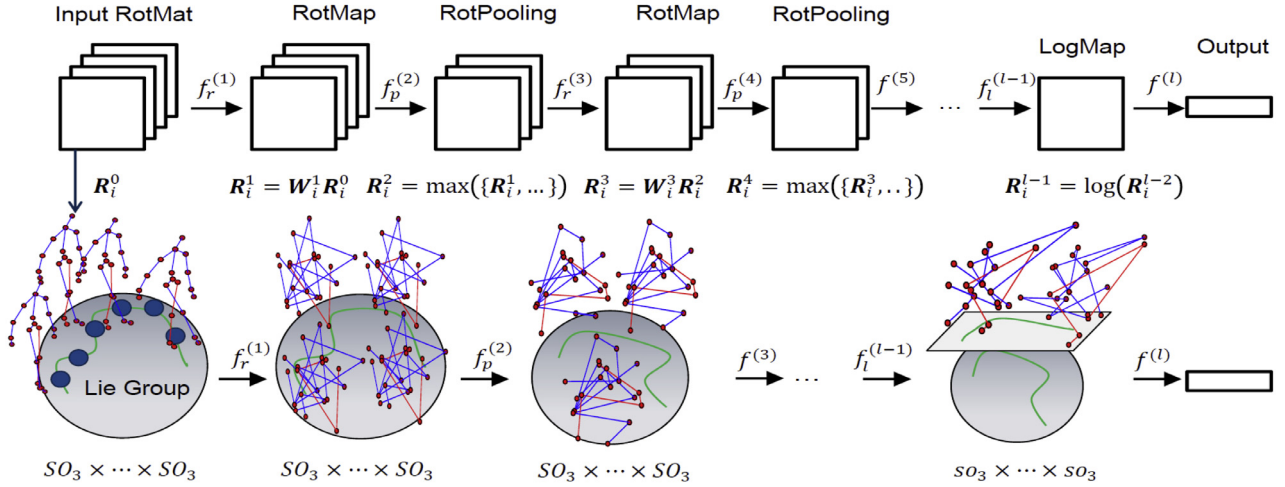


Fig. 16. Conceptual illustration of LieNet architecture. In the network structure, the data space of each RotMap/RotPooling layer corresponds to a Lie group, while the weight spaces of the RotMap layers are Lie groups as well. Figure from Huang et al. (2017).

skeletal data, capture the interactions between body-parts and/or interactions with environmental objects, and model the temporal structure of human actions in an end-to-end learning framework. The proposed method was made robust to viewpoint changes, by introducing a deep CNN which transfers visual appearance of human body-parts acquired from different unknown views to a view-invariant space.

6.1.2. RNN-based approach

For RGB and depth fusion, Pigou et al. (2016) directly considered the depth as the fourth channel and CNN was adopted to extract frame-based appearance features. Temporal convolutions and RNN were combined to capture the temporal information. Li et al. (2016b) adopted C3D (Tran et al., 2015) to extract features separately from RGB and depth modalities, and used the concatenated for SVM classifier. Zhu et al. (2017a) presented a gesture recognition method using C3D (Tran et al., 2015) and convolutional LSTM (convLSTM) (Xingjian et al., 2015) based on depth and RGB modalities (see Fig. 17). The major drawback of traditional LSTM in handling spatio-temporal data is its usage of full connections in input-to-state and state-to-state transitions

in which no spatial information is encoded. The ConvLSTM determines the future state of a certain cell in the grid by the inputs and past states of its local neighbors. Average score fusion was adopted to fuse the two separate channel networks for the two modalities. Luo et al. (2017) proposed to use a RNN-based encoder-decoder framework to learn a video representation by predicting a sequence of basic motions described as atomic 3D flows. The learned representation is then extracted from the generated model to recognize activities.

Shi and Kim (2017) fused depth and skeleton in a so-called privileged information (PI)-based RNN (PRNN) that exploits additional knowledge of skeleton sequences to obtain a better estimate of network parameters from depth map sequences. A bridging matrix is defined to connect softmax classification loss and regression loss by discovering latent PI in the refinement step. The whole process is illustrated in Fig. 19.

For RGB and skeleton fusion, Mahasseni and Todorovic (2016) presented a regularization of LSTM learning where the output of another encoder LSTM (eLSTM) grounded on 3D human-skeleton training data is used as the regularization. This regularization rests on the hypothesis that since videos and skeleton sequences are about human motions their respective feature representations should be similar. The skeleton sequences, being view-independent and devoid of background clutter, are expected to facilitate capturing important motion patterns of human-body joints in 3D space.

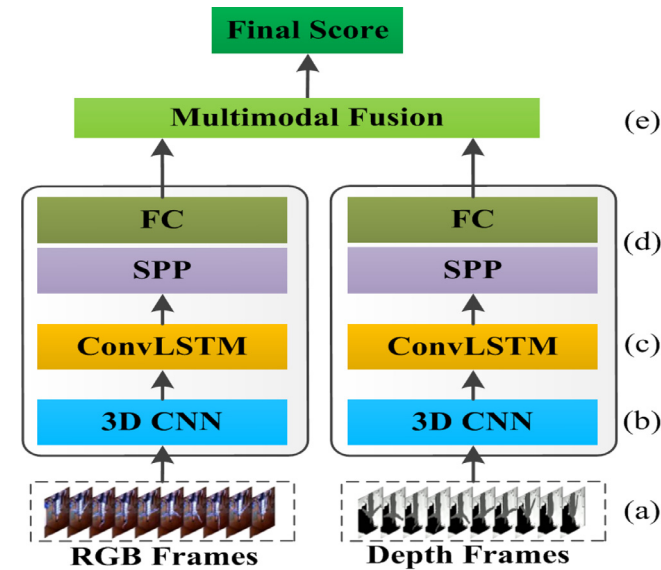


Fig. 17. The deep architecture is composed of five components: (a) input pre-processing; (b) 3D CNN (C3D); (c) ConvLSTM; (d) spatial pyramid pooling and fully connected layers; (e) multimodal score fusion. Figure from Zhu et al. (2017a).

6.1.3. Other-architecture-based approach

Shahroudy et al. (2017) extracted hand-crafted features which are neither independent nor fully correlated from RGB and depth, and embedded the input feature into a space of factorized common and modality-specific components. The combination of shared and specific components in input features can be very complex and highly non-




Fig. 18. The framework of using scene flow for motion recognition. Scene flow vectors are first transformed into Scene Flow Maps (SFM), and then using Channel Transform Kernels to transform SFM into an analogous RGB space to take advantage of pre-train models over ImageNet. Figure from Wang et al. (2017b).

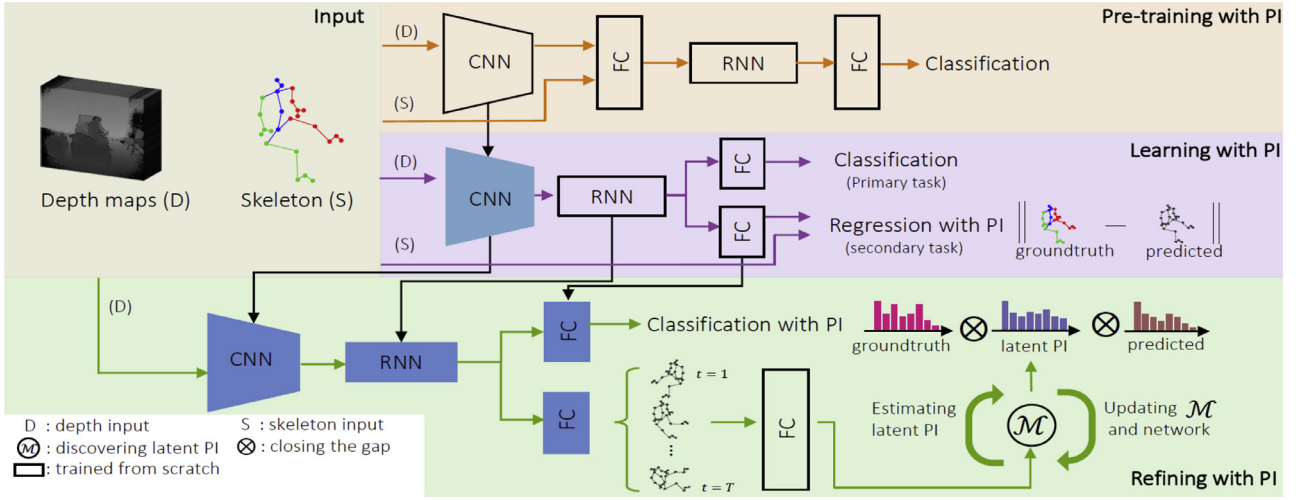


Fig. 19. The framework of PI-based RNNs. It consists of three steps: (1) the pre-training step taking both depth maps and skeleton as input. An embedded encoder is trained in a standard CNN-RNN pipeline. (2) The trained encoder is used to initialize the learning step. A multi-task loss is applied to exploit the PI in the regression term as a secondary task. (3) Finally, refinement step aims to discover the latent PI by defining a bridging matrix, in order to maximize the effectiveness of the PI. The latent PI is utilized to close the gap between different information. The latent PI, bridging matrix and the network are optimized iteratively in an EM procedure. Figure from Shi and Kim (2017).

linear. In order to disentangle them, they stacked layers of non-linear auto encoder-based component factorization to form a deep shared-specific analysis network.

In a RGB, depth and skeleton fusion method, Wu et al. (2016) adopted Gaussian–Bernoulli Deep Belief Network (DBN) to extract high-level skeletal joint features and the learned representation is used to estimate the emission probability needed to infer gesture sequences. A 3D Convolutional Neural Network (3DCNN) was used to extract features from 2D multiple channel inputs such as depth and RGB images stacked along the 1D temporal domain. In addition, intermediate and late fusion strategies were investigated in combination with the temporal modeling. The result of both mechanisms indicates that multiple-channel fusion can outperform individual modules.

6.2. Continuous/online motion recognition

6.2.1. RNN-based approach

Chai et al. (2016) proposed to fuse RGB and depth in a two-stream RNN (2S-RNN) for gesture recognition. They designed a fusion layer for depth and RGB before the LSTM layer. Molchanov et al. (2016) presented an algorithm for joint segmentation and classification of dynamic hand gestures from continuous video streams. They proposed a network that employs a recurrent C3D with connectionist temporal classification (CTC) (Graves et al., 2006). They trained a separate network for each modality and averaged their scores for final recognition. Beside RGB and depth modalities, they also adopted stereo-IR modality in their work.

7. Discussion

We presented a comprehensive overview of RGB-D based motion recognition using deep learning. We defined a taxonomy covering two groups: segmented and continuous/online motion recognition, with four categories in each group based on the adopted modalities. From the viewpoint of encoding spatio-temporal-structural information contained in the video sequences, CNN, RNN and other networks adopted for motion recognition are discussed in each category. In subsequent sections, the relative performance of the different methods on several commonly used RGB-D datasets are analysed, and from the comparisons we highlight some challenges. The discussion on performance and challenges then provides a basis for outlining potential future research

directions.

7.1. Performance analysis of the current methods

In this section, we compare the accuracy of different methods using several commonly used datasets, including CMU Mocap, HDM05, MSR-Action3D, MSRC-12, MSRDailyActivity3D, UTKinect, G3D, SBU Kinect Interaction, Berkeley MHAD, Northwestern-UCLA Multiview Action3D, ChaLearn LAP IsoGD, NTU RGB + D, ChaLearn2014, ChaLearn LAP ConGD, and PKU-MMD. These datasets cover motion capture sensor system, structured light cameras (Kinect v1) and ToF cameras (Kinect v2). The last three datasets are continuous datasets while the others are segmented datasets. The performance is evaluated using *accuracy* for segmented motion recognition, and *Jaccard Index* is added as another criteria for continuous motion recognition. The accuracy is calculated as the proportion of accurately labelled samples. The Jaccard index measures the average relative overlap between true and predicted sequences of frames for a given gesture/action. For a sequence s , let $G_{s,i}$ and $P_{s,i}$ be binary indicator vectors for which 1-values correspond to frames in which the i th gesture/action label is being performed. The Jaccard Index for the i th class is defined for the sequence s as:

$$J_{s,i} = \frac{G_{s,i} \cap P_{s,i}}{G_{s,i} \cup P_{s,i}}, \quad (1)$$

where $G_{s,i}$ is the ground truth of the i th gesture/action label in sequence s , and $P_{s,i}$ is the prediction for the i th label in sequence s . When $G_{s,i}$ and $P_{s,i}$ are empty, $J_{s,i}$ is defined to be 0. Then for the sequence s with l_s true labels, the Jaccard Index J_s is calculated as:

$$J_s = \frac{1}{l_s} \sum_{i=1}^{l_s} J_{s,i}. \quad (2)$$

For all test sequences $S = s_1, \dots, s_n$ with n gestures/actions, the mean Jaccard Index \bar{J}_S is used as the evaluation criteria and calculated as:

$$\bar{J}_S = \frac{1}{n} \sum_{j=1}^n J_{s_j}. \quad (3)$$

The detailed comparison of different methods is presented in following Table 2. From the Table we can see that there is no single approach that is able to produce the best performance over all datasets. Generally speaking, methods using multi-modal information can have

better performance than their single modality counterpart due to the complementary properties of the three different modalities. On some datasets, such as NTU RGB+D dataset, current results suggest that CNN-based methods tend to be better than RNN-based methods. This is probably due to fact that CNN-based methods includes human empirical knowledge in the coding process, and could take advantage of pre-trained models over large image set, such as ImageNet. The combination of CNN and RNN seems to be a good choice for motion recognition, for instance, the C3D+ConvLSTM (Zhu et al., 2017a) method achieved promising results on ChaLearn LAP IsoGD dataset. For continuous motion recognition, RNN-based methods tend to achieve good results.

7.2. Challenges

The advent of low-cost RGB-D sensors that have access to extra depth and skeleton data, has motivated the significant development of human motion recognition. Promising results have been achieved with deep learning approaches (Liu et al., 2016a; Wang et al., 2016b; Zhang et al., 2017c), on several constrained simple datasets, such as MSR-Action3D, Berkeley MHAD and SBU Kinect Interaction. Despite this success, results are far from satisfactory on some large complex datasets, such as ChaLearn LAP IsoGD and NTU RGB+D datasets and especially the continuous/online datasets. In fact, it is still very difficult to build a practical intelligent recognition system. Such goal poses several challenges:

Encoding temporal information. As discussed, there are several methods to encode temporal information. We can use CNN to extract frame-based features and then conduct temporal fusion (Karpathy et al., 2014), or adopt 3D filter and 3D pooling layers to learn motion features (Tran et al., 2015), or use optical/scene flow to extract motion information (Simonyan and Zisserman, 2014a; Wang et al., 2017b), or encode the video into images (Bilen et al., 2016; Wang et al., 2016b; 2016e), or use RNN/LSTM to model the temporal dependencies (Donahue et al., 2015; Du et al., 2015b; Liu et al., 2017a). However, all these approaches have their drawbacks. Temporal fusion method tends to neglect the temporal order; 3D filters and 3D pooling filters have a very rigid temporal structure and they only accept a predefined number of frames as input which is always short; optical/scene flow methods are computationally expensive; sequence to images methods inevitably loses temporal information during encoding; the weight sharing mechanism of RNN/LSTM methods make the sequence matching imprecise, but rather approximated, so an appropriate distance function must be used to predict the match probability. In fact, there is still no perfect method for temporal encoding, and how to model temporal information is a big challenge.

Small training data. Most of available deep learning methods rely on large labeled training data (Karpathy et al., 2014; Tran et al., 2015). However, in practical scenarios, obtaining large labeled training data is costly and laborious, even impossible, especially in medical-related applications. It has been shown that fine-tuning motion-based networks with spatial data (ImageNet) is more effective than training from scratch (Bilen et al., 2016; Simonyan and Zisserman, 2014a; Wang et al., 2017b; 2016e). Strategies for data augmentation are also commonly used (Wang et al., 2016b). Likewise, training mechanisms to avoid overfitting and control learning rate have also been studied (Srivastava et al., 2014). However, it is still a challenge to effectively train deep networks from small training data.

Viewpoint variation and occlusion. When skeletons are estimated from RGB images/video or depth maps, viewpoint variation may cause significantly different appearance of same actions, and occlusion would “crash” the skeleton data. Occlusion includes inter-occlusion caused by other subjects or objects, and self-occlusion created by the object/subject itself. Most of available datasets require subjects to perform actions in a visible and restricted view to avoid occlusion, and this results in limited view data collection and less occlusion. However,

occlusion is inevitable in practical scenarios, especially for interactions. This makes it challenging to isolate individuals in overlapping area and extract features of a unique person; leading to the ineffectiveness of many of available approaches (Du et al., 2015b; Li et al., 2017a; Shahroudy et al., 2016). Possible solutions to handle viewpoint variation and occlusion include the use of multi-sensor systems (Chunhui et al., 2017; Ofli et al., 2013; Shahroudy et al., 2016; Wang et al., 2014). The multi-camera systems is able to generate multi-view data, but the drawback is the requirement of synchronization and feature/recognition fusion among different views. This usually increases processing complexity and computation cost. Several methods have been proposed to handle the viewpoint variation and occlusion. Wang et al. (2015b) proposed to rotate the depth data in 3D point clouds through different angles to deal with viewpoint invariance; spherical coordinate system corresponding to body center was developed to achieve view-independent motion recognition (Huang et al., 2017). However, these methods become less effective when occlusion occurs. How to effectively handle occlusion using deep learning methods is a new challenge.

Execution rate variation and repetition. The execution rate may vary due to the different performing styles and states of individuals. The varying rate results in different frames for the same motion. Repetition also bring about this issue. The global encoding methods (Hou et al., 2016; Ke et al., 2017a; Liu et al., 2017b) would become less effective due to the repetition. The commonly used methods to handle this problem is up/down sampling (Li et al., 2017a; Zhang et al., 2017c; Zhu et al., 2016d). However, sampling methods would inevitably bring redundant or loss of useful information. Effective handling of this problem remains a challenge.

Cross-datasets. Many research works have been carried out to recognize human actions from RGB-D video clips. To learn an effective action classifier, most of the previous approaches rely on enough training labels. When being required to recognize the action in a different dataset, these approaches have to re-train the model using new labels. However, labeling video sequences is a very tedious and time-consuming task, especially when detailed spatial locations and time durations are required. Even though some works have studied this topic (Cao et al., 2010; Sultani and Saleemi, 2014; Zhang et al., 2017a), they are all based on hand-crafted features, and the results are far from satisfactory due to the large distribution variances between different datasets, including different scenarios, different modalities, different views, different persons, and even different actions. How to deal with cross-datasets RGB-D motion recognition is a big challenge.

Online motion recognition. Most of available methods rely on segmented data, and their capability for online recognition is quite limited. Even though continuous motion recognition is one improved version where the videos are untrimmed, it still assumes that all the videos are available before processing. Thus, proposal-based methods (Shou et al., 2016; Wang et al., 2017a) can be adopted for offline processing. Differently from continuous motion recognition, online motion recognition aims to receive continuous streams of unprocessed visual data and recognize actions from an unsegmented stream of data in a continuous manner. So far two main approaches can be identified for online recognition, sliding window-based and RNN-based. Sliding window-based methods (Chunhui et al., 2017) are simple extension of segmented-based action recognition methods. They often consider the temporal coherence within the window for prediction and the window-based predictions are further fused to achieve online recognition. However, the performance of these methods are sensitive to the window size which depends on actions and is hard to set. Either too large or too small a window size could lead to significant drop in recognition. For RNN-based methods (Li et al., 2016a; Molchanov et al., 2016), even though promising results have been achieved, it is still far from satisfactory in terms of performance. How to design effective practical online recognition system is a big challenge.

Action prediction. We are faced with numerous situations in which

we must predict what actions other people are about to do in the near future. Predicting future actions before they are actually executed is a critical ingredient for enabling us to effectively interact with other humans on a daily basis (Hoai and De la Torre, 2014; Lan et al., 2014; Ryoo, 2011; Sadeh Aliakbarian et al., 2017; Vu et al., 2014). There are mainly two challenges for this task: first, we need to capture the subtle details inherent in human movements that may imply a future action; second, predictions usually should be carried out as quickly as possible in the social world, when limited prior observations are available. Predicting the action of a person before it is actually executed has a wide range of applications in autonomous robots, surveillance and health care. How to develop effective algorithms for action prediction is really challenging.

7.3. Future research directions

The discussion on the challenges faced by available methods allows us to outline several future research directions for the development of deep learning methods for motion recognition. While the list is not exhaustive, they point at research activities that may advance the field.

Hybrid networks. Most of previous methods adopted one type of neural networks for motion recognition. As discussed, there is no perfect solution for temporal encoding using single networks. Even though available works such as C3D+ConvLSTM (Zhu et al., 2017a) used two types of networks, the cascaded connection makes them dependent on each other during training. How to cooperatively train different kinds of networks would be a good research direction; for example, using the output of CNN to regularize RNN training in parallel.

Simultaneous exploitation of spatial-temporal-structural information. A video sequence has three important inherent properties that should be considered for motion analysis: spatial information, temporal information and structural information. Several previous methods tend to exploit the spatio-temporal information for motion recognition, however, structural information contained in the video is rarely explicitly mined. Concurrent mining of these three kinds of information with deep learning would be an interesting topic in the future (Jain et al., 2016).

Fusion of multiple modalities. While significant progress has been achieved by singly using RGB, skeleton or depth modality, effective deep networks for fusion of multi-modal data would be a promising direction. For example, methods such as SFAM (Wang et al., 2017b) and PRNN (Shi and Kim, 2017) have pioneered the research in this direction. The work SFAM (Wang et al., 2017b) proposed to extract scene flow for motion analysis. The strategy of fusing the RGB and depth modalities at the outset allowed the capture of rich 3D motion information. In PRNN (Shi and Kim, 2017) the concept of privileged information (side information) was introduced for deep networks training and showed some promise. Zolfaghari et al. (2017) proposed the integration of different modalities via a Markov chain, which leads to a sequential refinement of action labels. So far, most methods considered the three modalities as separate channels and fused them at a later or scoring stage using different fusion methods without cooperatively exploiting their complementary properties. Cooperative training using different modalities would be a promising research area.

Large-scale datasets. With the development of data-hungry deep learning approach, there is demand for large scale RGB-D datasets. Even though there are several large datasets, such as NTU RGB+D Dataset (Shahroudy et al., 2016) and ChaLearn LAP IsoGD Dataset (Wan et al., 2016b), they are focused on specific tasks. Various large-scale RGB-D datasets are needed to facilitate research in this field. For instance, large-scale fine-grained RGB-D motion recognition datasets and large-scale occlusion-based RGB-D motion recognition datasets are urgently needed.

Zero/one-shot learning. As discussed, it is not always easy to collect large scale labeled data. Learning from a few examples remains a key challenge in machine learning. Despite recent advances in important domains such as vision and language, the standard supervised deep

learning paradigm does not offer a satisfactory solution for learning new concepts rapidly from little data. How to adopt deep learning methods for zero/one shot RGB-D-based motion recognition would be an interesting research direction. Zero/one-shot learning is about being able to recognize gesture/action classes that are never seen or only one training sample per class before. This type of recognition should carry embedded information universal to all other gestures/actions. In the past few years, there are some works on zero/one-shot learning. For example, Wan et al. (2016a) proposed the novel spatial-temporal features for one-shot learning gesture recognition and have got promising performances on ChaLearn Gesture Dataset CGD (Guyon et al., 2014). For zero-shot learning, Madapana and Wachs (2017) proposed a new paradigm based on adaptive learning which it is possible to determine the amount of transfer learning carried out by the algorithm and how much knowledge is acquired for a new gesture observation. However, the mentioned works used traditional methods (such as bag of visual words model Wan et al., 2013). Mettes and Snoek (2017) proposed a spatial-aware object embedding for zero-shot action localization and classification. The spatial-aware embedding generate action tubes by incorporating word embeddings, box locations for actors and objects, as well as their spatial relations. However, how to effectively adopt deep learning methods for zero/one shot RGB-D based motion recognition would be still an interesting research direction especially when using only very few training samples.

Outdoor practical scenarios. Although lots of RGB-D datasets have been collected during the last few years, there is a big gap between the collected datasets and wild environment due to constrained environment setting and insufficient categories and samples. For example, most available datasets do not involve much occlusion cases probably due to the collapse of skeleton dataset in case of occlusion. However, in practical scenarios, occlusion is inevitable. How to recover or find cues from multi-modal data for such recognition tasks would be an interesting research direction. Besides, with the development of depth sensors, further distances could be captured, and recognition in outdoor practical scenarios will gain the attention of researchers.

Unsupervised learning/self-learning. Collecting labeled datasets are time-consuming and costly, hence learning from unsupervised video data is required. Mobile robots mounted with RGB-D cameras need to continuously learn from the environment and without human intervention. How to automatically learn from the unlabeled data stream to improve the learning capability of deep networks would be a fruitful and useful research direction. Generative Adversarial Net (GAN) (Ho and Ermon, 2016) has got much processes recently in image generation task, such as face generation, text-to-image task. Besides, it also can be used for recognition task. For example, Tran et al. (2017) proposed a Disentangled Representation learning Generative Adversarial Networks (DR-GAN) for pose-invariant face recognition. Therefore, we believe the GAN-based techniques also can be used for action/gesture recognition, which is a great exciting direction for research. Vondrick et al. (2016b) proposed a generative adversarial network for video with spatial-temporal convolutional architecture that untangles the scene's foreground from backgrounds. This is an initial work to capitalize on large amounts of unlabeled video in order to learn a model of scene dynamic for both video recognition tasks (e.g. action classification) and video generation tasks (e.g. future prediction). Increasing research will be reported in the coming years on GAN-based methods for video-based recognition.

Online motion recognition and prediction. Online motion recognition and prediction is required in practical applications, and arguably this is the final goal of motion recognition systems. Differently from segmented recognition, online motion recognition requires the analysis of human behavior in a continuous manner, and prediction aims to recognize or anticipate actions that would happen. How to design effective online recognition and prediction systems with deep learning methods has attracted some attention. For example, Vondrick et al. (2016a) introduced a framework that capitalizes on

temporal structure in unlabeled video to learn to anticipate human actions and objects based on CNN, and it is likely to emerge as an active research area.

8. Conclusion

This paper presents a comprehensive survey of RGB-D based motion recognition using deep learning. We provide a brief overview of existing commonly used datasets and pointed at surveys that focused mainly on datasets. The available methods are grouped into four categories according to the modality: RGB-based, depth-based, skeleton-based and RGB+D-based. The three modalities have their own specific features and lead to different choices of deep learning methods to take advantages of their properties. Spatial, temporal and structural information inherent in a video sequence is defined, and from the viewpoint of spatio-temporal-structural encoding, we analyse the pros and cons of available methods. Based on the insights drawn from the survey, several potential research directions are described, indicating the numerous opportunities in this field despite the advances achieved to date.

Acknowledgment

Jun Wan is partially supported by the National Natural Science Foundation of China [61502491]. Sergio Escalera is partially supported by Spanish project [TIN2016-74946-P] (MINECO/FEDER, UE) and CERCA Programme/Generalitat de Catalunya.

References

- Aggarwal, J., Cai, Q., 1999. Human motion analysis: a review. *Proceedings of the Computer Vision and Image Understanding*.
- Aggarwal, J.K., Xia, L., 2014. Human activity recognition from 3D data: a review. *Pattern Recognit. Lett.* 48, 70–80.
- Asadi-Aghbolaghi, M., Clapes, A., Bellantonio, M., Escalante, H.J., Ponce-López, V., Baró, X., Guyon, I., Kasaei, S., Escalera, S., 2017. A survey on deep learning based approaches for action and gesture recognition in image sequences. *Proceedings of the 2017 Twelfth IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*. IEEE, pp. 476–483.
- Baccouche, M., Mamalet, F., Wolf, C., Garcia, C., Baskurt, A., 2011. Sequential deep learning for human action recognition. *Proceedings of the International Workshop on Human Behavior Understanding*. Springer, pp. 29–39.
- Bilen, H., Fernando, B., Gavves, E., Vedaldi, A., Gould, S., 2016. Dynamic image networks for action recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Bloom, J., Makris, D., Argyriou, V., 2012. G3D: a gaming action dataset and real time action recognition evaluation framework. *Proceedings of the CVPRW*.
- Bobick, A.F., Davis, J.W., 2001. The recognition of human movement using temporal templates. *IEEE Trans. Pattern Anal. Mach. Intell.* 23 (3), 257–267.
- Buch, S., Escorcia, V., Shen, C., Ghanem, B., Niebles, J.C., 2017. SST: single-stream temporal action proposals. *Proceedings of the Computer Vision and Pattern Recognition*.
- CMU, 2001. Graphics Lab Motion Capture Database, <http://mocap.cs.cmu.edu/>.
- Camgoz, N.C., Hadfield, S., Koller, O., Bowden, R., 2016. Using convolutional 3D neural networks for user-independent continuous gesture recognition. *Proceedings of the 2016 twenty-third International Conference on Pattern Recognition (ICPR)*. pp. 49–54.
- Cao, L., Liu, Z., Huang, T.S., 2010. Cross-dataset action detection. *Proceedings of the Computer vision and pattern recognition (CVPR)*. IEEE, pp. 1998–2005.
- Chai, X., Liu, Z., Yin, F., Liu, Z., Chen, X., 2016. Two streams recurrent neural networks for large-scale continuous gesture recognition. *Proceedings of the International Conference on Pattern Recognition Workshops*.
- Chen, B., 2010. Deep learning of invariant spatio-temporal features from video. University of British Columbia Ph.D. thesis.
- Chen, L., Wei, H., Ferryman, J., 2013. A survey of human motion analysis using depth imagery. *Pattern Recognit. Lett.* 34 (15), 1995–2006.
- Chen, Y., Shen, C., Wei, X.-S., Liu, L., Yang, J., 2017. Adversarial posenet: A structure-aware convolutional network for human pose estimation. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Cheng, H., Yang, L., Liu, Z., 2016. Survey on 3D hand gesture recognition. *IEEE Trans. Circuits Syst. Video Technol.* 26 (9), 1659–1673.
- Cherian, A., Fernando, B., Harandi, M., Gould, S., 2017. Generalized rank pooling for activity recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Chéron, G., Laptev, I., Schmid, C., 2015. P-CNN: pose-based CNN features for action recognition. *Proceedings of the International Conference on Computer Vision*. pp. 3218–3226.
- Chunhui, L., Yueyu, H., Yanghao, L., Sijie, S., Jiaying, L., 2017. PKU-MMD: a large scale benchmark for continuous multi-modal human action understanding. *Proceedings of the Workshop on Visual Analysis in Smart and Connected Communities (VSCC'17)*. ACM, pp. 1–8.
- Dave, A., Russakovsky, O., Ramanan, D., 2017. Predictive-corrective networks for action detection. *Proceedings of the Computer Vision and Pattern Recognition*.
- De Geest, R., Gavves, E., Ghodrati, A., Li, Z., Snoek, C., Tuytelaars, T., 2016. Online action detection. *Proceedings of the European Conference on Computer Vision*. pp. 269–284.
- Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T., 2015. Long-term recurrent convolutional networks for visual recognition and description. *Computer Vision and Pattern Recognition*.
- Du, W., Wang, Y., Qiao, Y., 2017. Rpan: An end-to-end recurrent pose-attention network for action recognition in videos. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3725–3734.
- Du, Y., Fu, Y., Wang, L., 2015. Skeleton based action recognition with convolutional neural network. *Proceedings of the 2015 Third IAPR Asian Conference on Pattern Recognition (ACPR)*. IEEE, pp. 579–583.
- Du, Y., Fu, Y., Wang, L., 2016. Representation learning of temporal dynamics for skeleton-based action recognition. *IEEE Trans. Image Process.* 25 (7), 3010–3022.
- Du, Y., Wang, W., Wang, L., 2015b. Hierarchical recurrent neural network for skeleton based action recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Duan, J., Zhou, S., Wan, J., Guo, X., Li, S.Z., 2016. A unified framework for multi-modal isolated gesture recognition. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 9 (4), 39:1–39:17.
- Escalera, S., Athitsos, V., Guyon, I., 2016. Challenges in multimodal gesture recognition. *J. Mach. Learn. Res.* 17 (72), 1–54.
- Escalera, S., Baró, X., Gonzalez, J., Bautista, M.A., Madadi, M., Reyes, M., Ponce-López, V., Escalante, H.J., Shotton, J., Guyon, I., 2014. Chalearn looking at people challenge 2014: Dataset and results. *Proceedings of the Workshop at the European Conference on Computer Vision*. Springer, pp. 459–473.
- Escorcia, V., Heilbron, F.C., Niebles, J.C., Ghanem, B., 2016. Daps: Deep action proposals for action understanding. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 768–784.
- Feichtenhofer, C., Pinz, A., Wildes, R., 2016a. Spatiotemporal residual networks for video action recognition. *Proceedings of the Advances in Neural Information Processing Systems*. pp. 3468–3476.
- Feichtenhofer, C., Pinz, A., Zisserman, A., 2016b. Convolutional two-stream network fusion for video action recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1933–1941.
- Fernando, B., Anderson, P., Hutter, M., Gould, S., 2016a. Discriminative hierarchical rank pooling for activity recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Fernando, B., Gavves, S., Mogrovejo, O., Antonio, J., Ghodrati, A., Tuytelaars, T., 2016b. Rank pooling for action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 39, 773–787.
- Fernando, B., Gould, S., 2016. Learning end-to-end video classification with rank-pooling. *Proceedings of the International Conference on Machine Learning*.
- Fothergill, S., Mentis, H.M., Nowozin, S., Kohli, P., 2012. Instructing people for training gestural interactive systems. *Proceedings of the ACM Conference on Computer-Human Interaction (ACM HCI)*.
- Gao, J., Yang, Z., Sun, C., Chen, K., Nevatia, R., 2017. Turn tap: temporal unit regression network for temporal action proposals. *Proceedings of the International Conference on Computer Vision*.
- Girshick, R., 2015. Fast R-CNN. *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1440–1448.
- Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. *Proceedings of the Computer Vision and Pattern Recognition*.
- Gkioxari, G., Malik, J., 2015. Finding action tubes. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 759–768.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press. <http://www.deeplearningbook.org>
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. *Proceedings of the Advances in Neural Information Processing Systems*. pp. 2672–2680.
- Graves, A., Fernández, S., Gomez, F., Schmidhuber, J., 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. *Proceedings of the Twenty-third International Conference on Machine Learning*. ACM, pp. 369–376.
- Guo, G., Lai, A., 2014. A survey on still image based human action recognition. *Pattern Recognit.* 47 (10), 3343–3361.
- Guyon, I., Athitsos, V., Jangyodsuk, P., Escalante, H.J., 2014. The Chalearn gesture dataset (CGD 2011). *Mach. Vis. Appl.* 25 (8), 1929–1951.
- Han, F., Reilly, B., Hoff, W., Zhang, H., 2017. Space-time representation of people based on 3D skeletal data: a review. *Proceedings of the Computer Vision Image Understanding*.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C., 2015. Activitynet: a large-scale video benchmark for human activity understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 961–970.
- Herath, S., Harandi, M., Porikli, F., 2017. Going deeper into action recognition: a survey. *Image Vis. Comput.* 60, 4–21.
- Ho, J., Ermon, S., 2016. Generative adversarial imitation learning. *Proceedings of the*

- Advances in Neural Information Processing Systems. pp. 4565–4573.
- Hoai, M., De la Torre, F., 2014. Max-margin early event detectors. *Int. J. Comput. Vis.* 107 (2), 191–202.
- Hou, R., Chen, C., Shah, M., 2017a. Tube convolutional neural network (t-CNN) for action detection in videos. *Proceedings of the International Conference on Computer Vision*.
- Hou, Y., Li, Z., Wang, P., Li, W., 2016. Skeleton optical spectra based action recognition using convolutional neural networks. *Proceedings of the IEEE Transactions on Circuits and Systems for Video Technology*. pp. 1–5.
- Hou, Y., Wang, S., Wang, P., Gao, Z., Li, W., 2017b. Spatially and temporally structured global to local aggregation of dynamic depth information for action recognition. *IEEE Access*.
- Huang, D.-A., Fei-Fei, L., Niebles, J.C., 2016. Connectionist temporal modeling for weakly supervised action labeling. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 137–153.
- Huang, Z., Wan, C., Probst, T., Van Gool, L., 2017. Deep learning on lie groups for skeleton-based action recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Ijjina, E.P., Krishna Mohan, C., 2016. Classification of human actions using pose-based features and stacked auto encoder. *Pattern Recognit. Lett.* 83, 268–277.
- Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., Schiele, B., 2016. DeepCUT: a deeper, stronger, and faster multi-person pose estimation model. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 34–50.
- Jain, A., Zamir, A.R., Savarese, S., Saxena, A., 2016. Structural-RNN: Deep learning on spatio-temporal graphs. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5308–5317.
- Jhuang, H., Serre, T., Wolf, L., Poggio, T., 2007. A biologically inspired system for action recognition. *Proceedings of the IEEE Eleventh International Conference on Computer Vision*. pp. 1–8.
- Ji, S., Xu, W., Yang, M., Yu, K., 2013. 3D Convolutional neural networks for human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (1), 221–231.
- Jiang, F., Zhang, S., Wu, S., Gao, Y., Zhao, D., 2015. Multi-layered gesture recognition with Kinect. *J. Mach. Learn. Res.* 16 (2), 227–254.
- Kalogeiton, V., Weinzaepfel, P., Ferrari, V., Schmid, C., 2017. Action tubelet detector for spatio-temporal action localization. *Proceedings of the International Conference on Computer Vision*.
- Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L., 2014. Large-scale video classification with convolutional neural networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1725–1732.
- Ke, Q., An, S., Bennamoun, M., Soheli, F., Boussaid, F., 2017a. Skeletonnet: mining deep part features for 3D action recognition. *IEEE Signal Process. Lett.*
- Ke, Q., Bennamoun, M., An, S., Soheli, F., Boussaid, F., 2017b. A new representation of skeleton sequences for 3D action recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Kim, T.S., Reiter, A., 2017. Interpretable 3D human action analysis with temporal convolutional networks. *Proceedings of the Computer Vision and Pattern Recognition*.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet classification with deep convolutional neural networks. *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*. pp. 1106–1114.
- Lan, T., Chen, T.-C., Savarese, S., 2014. A hierarchical representation for future action prediction. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 689–704.
- Le, Q.V., Zou, W.Y., Yeung, S.Y., Ng, A.Y., 2011. Learning hierarchical invariant spatio-temporal features for action recognition with independent subspace analysis. *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, pp. 3361–3368.
- Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D., 2017. Temporal convolutional networks for action segmentation and detection. *Proceedings of the Proceedings of the Computer Vision and Pattern Recognition*.
- Lea, C., Reiter, A., Vidal, R., Hager, G.D., 2016. Segmental spatiotemporal CNNs for fine-grained action segmentation. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 36–52.
- Lee, I., Kim, D., Kang, S., Lee, S., 2017. Ensemble deep learning for skeleton-based action recognition using temporal sliding LSTM networks. *Proceedings of the International Conference on Computer Vision*. pp. 1012–1020.
- Lev, G., Sadeh, G., Klein, B., Wolf, L., 2016. RNN fisher vectors for action recognition and image annotation. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 833–850.
- Li, C., Hou, Y., Wang, P., Li, W., 2017a. Joint distance maps based action recognition with convolutional neural networks. *IEEE Signal Process. Lett.* 24 (5), 624–628.
- Li, W., Wen, L., Chang, M.-C., Nam Lim, S., Lyu, S., 2017b. Adaptive RNN tree for large-scale human action recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1444–1452.
- Li, W., Zhang, Z., Liu, Z., 2010. Action recognition based on a bag of 3D points. *Proceedings of the Conference on Computer Vision and Pattern Recognition Workshop*.
- Li, Y., Lan, C., Xing, J., Zeng, W., Yuan, C., Liu, J., 2016a. Online human action detection using joint classification-regression recurrent neural networks. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 203–220.
- Li, Y., Miao, Q., Tian, K., Fan, Y., Xu, X., Li, R., Song, J., 2016b. Large-scale gesture recognition with a fusion of RGB-D data based on the C3D model. *Proceedings of the 2016 Twenty-third International Conference on Pattern Recognition (ICPR)*. IEEE, pp. 25–30.
- Li, Y., Miao, Q., Tian, K., Fan, Y., Xu, X., Li, R., Song, J., 2016c. Large-scale gesture recognition with a fusion of RGB-D data based on the C3D model. *Proceedings of International Conference on Pattern Recognition Workshop*.
- Liu, J., Shahroudy, A., Xu, D., Wang, G., 2016a. Spatio-temporal LSTM with trust gates for 3D human action recognition. *Proceedings of the European Conference on Computer Vision*.
- Liu, J., Wang, G., Hu, P., Duan, L.-Y., Kot, A.C., 2017a. Global context-aware attention LSTM networks for 3D action recognition. *Proceedings of the Computer Vision and Pattern Recognition*.
- Liu, M., Liu, H., Chen, C., 2017b. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recognit.* 68, 346–362.
- Liu, W., Angelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016b. SSA: Single shot multibox detector. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 21–37.
- Liu, Z., Zhang, C., Tian, Y., 2016. 3D-Based deep convolutional neural network for action recognition with depth sequences. *Image Vis. Comput.* 55, 93–100.
- Luo, Z., Peng, B., Huang, D.-A., Alahi, A., Fei-Fei, L., 2017. Unsupervised learning of long-term motion dynamics for videos. *Proceedings of the Computer Vision and Pattern Recognition*.
- Ma, S., Sigal, L., Sclaroff, S., 2016. Learning activity progression in LSTMs for activity detection and early detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1942–1950.
- Madapana, N., Wachs, J.P., 2017. A semantical & analytical approach for zero shot gesture learning. *Proceedings of the FG Workshop*.
- Mahasseni, B., Todorovic, S., 2016. Regularizing long short term memory with 3D human-skeleton sequences for action recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3054–3062.
- Mathieu, M., Couprie, C., LeCun, Y., 2016. Deep multi-scale video prediction beyond mean square error. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Memisevic, R., Hinton, G., 2007. Unsupervised learning of image transformations. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 1–8.
- Mettes, P., Snoek, C.G., 2017. Spatial-aware object embeddings for zero-shot localization and classification of actions. *Proceedings of the International Conference on Computer Vision*.
- Molchanov, P., Yang, X., Gupta, S., Kim, K., Tyree, S., Kautz, J., 2016. Online detection and classification of dynamic hand gestures with recurrent 3D convolutional neural network. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4207–4215.
- Müller, M., Röder, T., Clausen, M., Eberhardt, B., Krüger, B., Weber, A., 2007. Documentation Mocap Database HDM05. Technical Report, CG-2007-2. Universität Bonn.
- Ng, J., Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G., 2015. Beyond short snippets: deep networks for video classification. *Proceedings of the Computer Vision and Pattern Recognition*.
- Ofl, F., Chaudhry, R., Kurillo, G., Vidal, R., Bajcsy, R., 2013. Berkeley MHAD: a comprehensive multimodal human action database. *Proceedings of the IEEE Workshop on Applications of Computer Vision (WACV)*. pp. 53–60.
- Peng, X., Schmid, C., 2016. Multi-region two-stream r-CNN for action detection. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 744–759.
- Pigou, L., Van Den Oord, A., Dieleman, S., Van Herreweghe, M., Dambre, J., 2016. Beyond temporal pooling: recurrence and temporal convolutions for gesture recognition in video. *Int. J. Comput. Vis.* 1–10.
- Poppe, R., 2010. A survey on vision-based human action recognition. *Image Vis. Comput.* 28 (6), 976–990.
- Presti, L.L., La Cascia, M., 2016. 3D skeleton-based human action classification: a survey. *Pattern Recognit.* 53, 130–147.
- Rahmani, H., Bennamoun, M., 2017. Learning action recognition model from depth and skeleton videos. *Proceedings of the International Conference on Computer Vision*. pp. 5832–5841.
- Rahmani, H., Mian, A., 2016. 3D action recognition from novel viewpoints. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1506–1515.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN: towards real-time object detection with region proposal networks. *Proceedings of the Advances in neural information processing systems*. pp. 91–99.
- Ryoo, M.S., 2011. Human activity prediction: early recognition of ongoing activities from streaming videos. *Proceedings of the 2011 IEEE International Conference on Computer Vision (ICCV)*. IEEE, pp. 1036–1043.
- Sadegh Aliakbarian, M., Sadat Saleh, F., Salzmann, M., Fernando, B., Petersson, L., Andersson, L., 2017. Encouraging LSTMs to anticipate actions very early. *Proceedings of the International Conference on Computer Vision*. pp. 280–289.
- Saha, S., Singh, G., Cuzzolin, F., 2017. Amtnet: Action-micro-tube regression by end-to-end trainable deep architecture. *Proceedings of the International Conference on Computer Vision*.
- Salakhutdinov, R., Tenenbaum, J.B., Torralba, A., 2013. Learning with hierarchical-deep models. *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (8), 1958–1971.
- Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J., 2013. Image classification with the fisher vector: theory and practice. *Int. J. Comput. Vis.* 105 (3), 222–245.
- Shahroudy, A., Liu, J., Ng, T.-T., Wang, G., 2016. NTU RGB + D: a large scale dataset for 3D human activity analysis. *Proceedings of the Computer Vision and Pattern Recognition*.
- Shahroudy, A., Ng, T.-T., Gong, Y., Wang, G., 2017. Deep multimodal feature analysis for action recognition in RGB + d videos. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 1045–1058.
- Sharma, S., Kiros, R., Salakhutdinov, R., 2016. Action recognition using visual attention. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Shi, Y., Tian, Y., Wang, Y., Zeng, W., Huang, T., 2017. Learning long-term dependencies for action recognition with a biologically-inspired deep network. *Proceedings of the*

- International Conference on Computer Vision. pp. 716–725.
- Shi, Z., Kim, T.-K., 2017. Learning and refining of privileged information-based RNNs for action recognition from depth sequences. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Shotton, J., Fitzgibbon, A., Cook, M., Sharp, T., Finocchio, M., Moore, R., Kipman, A., Blake, A., 2011. Real-time human pose recognition in parts from single depth images. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1297–1304.
- Shou, Z., Wang, D., Chang, S.-F., 2016. Temporal action localization in untrimmed videos via multi-stage CNNs. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1049–1058.
- Simonyan, K., Zisserman, A., 2014a. Two-stream convolutional networks for action recognition in videos. *Proceedings of the Neural Information Processing Systems*.
- Simonyan, K., Zisserman, A., 2014b. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*.
- Singh, B., Marks, T.K., Jones, M., Tuzel, O., Shao, M., 2016. A multi-stream bi-directional recurrent neural network for fine-grained action detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1961–1970.
- Singh, G., Saha, S., Sapienza, M., Torr, P., Cuzzolin, F., 2017. Online real-time multiple spatiotemporal action localisation and prediction. *Proceedings of the International Conference on Computer Vision*.
- Song, S., Lan, C., Xing, J., Zeng, W., Liu, J., 2017. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.
- Srivastava, N., Hinton, G.E., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15 (1), 1929–1958.
- Srivastava, N., Mansimov, E., Salakhutdinov, R., 2015. Unsupervised learning of video representations using LSTMs. *Proceedings of the International Conference on Machine Learning*. pp. 843–852.
- Sultani, W., Saleemi, I., 2014. Human action recognition across datasets by foreground-weighted histogram decomposition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 764–771.
- Sun, L., Jia, K., Chen, K., Yeung, D.-Y., Shi, B.E., Savarese, S., 2017. Lattice long short-term memory for human action recognition. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Sun, L., Jia, K., Yeung, D.-Y., Shi, B.E., 2015. Human action recognition using factorized spatio-temporal convolutional networks. *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4597–4605.
- Taylor, G.W., Fergus, R., LeCun, Y., Bregler, C., 2010. Convolutional learning of spatio-temporal features. *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 140–153.
- Teh, Y.W., Jordan, M.I., Beal, M.J., Blei, D.M., 2004. Sharing clusters among related groups: Hierarchical dirichlet processes. *Proceedings of the Neural Information Processing Systems*. pp. 1385–1392.
- Theis, F.J., 2007. Towards a general independent subspace analysis. *Proceedings of the Advances in Neural Information Processing Systems*. pp. 1361–1368.
- Toshev, A., Szegedy, C., 2014. DeepPose: human pose estimation via deep neural networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1653–1660.
- Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M., 2015. Learning spatiotemporal features with 3D convolutional networks. *Proceedings of the International Conference on Computer Vision*.
- Tran, L., Yin, X., Liu, X., 2017. Disentangled representation learning GAN for pose-invariant face recognition. *Proceedings of the Computer Vision and Pattern Recognition*. 4. pp. 7.
- Turaga, P., Chellappa, R., Subrahmanian, V.S., Udrea, O., 2008. Machine recognition of human activities: a survey. *IEEE Trans. Circ. Syst. Video Technol.* 18 (11), 1473–1488.
- Varol, G., Laptev, I., Schmid, C., 2018. Long-term temporal convolutions for action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 1–8.
- Veeriah, V., Zhuang, N., Qi, G.-J., 2015. Differential recurrent neural networks for action recognition. *Proceedings of the International Conference on Computer Vision*.
- Vondrick, C., Pirsivash, H., Torralba, A., 2016a. Anticipating visual representations from unlabeled video. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 98–106.
- Vondrick, C., Pirsivash, H., Torralba, A., 2016b. Generating videos with scene dynamics. *Proceedings of the Advances in Neural Information Processing Systems*. pp. 613–621.
- Vu, T.-H., Olsson, C., Laptev, I., Oliva, A., Sivic, J., 2014. Predicting actions from static scenes. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 421–436.
- Wan, J., Guo, G., Li, S.Z., 2016a. Explore efficient local features from RGB-D data for one-shot learning gesture recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 38 (8), 1626–1639.
- Wan, J., Li, S.Z., Zhao, Y., Zhou, S., Guyon, I., Escalera, S., 2016b. Chalearn looking at people RGB-D isolated and continuous datasets for gesture recognition. *Proceedings of the Computer Vision and Pattern Recognition Workshop*.
- Wan, J., Ruan, Q., Li, W., Deng, S., 2013. One-shot learning gesture recognition from RGB-D data using bag of features. *J. Mach. Learn. Res.* 14, 2549–2582.
- Wang, H., Kläser, A., Schmid, C., Liu, C.-L., 2013. Dense trajectories and motion boundary descriptors for action recognition. *Int. J. Comput. Vis.* 103 (1), 60–79.
- Wang, J., Liu, Z., Wu, Y., Yuan, J., 2012. Mining actionlet ensemble for action recognition with depth cameras. *Proceedings of the Computer Vision and Pattern Recognition*.
- Wang, J., Nie, X., Xia, Y., Wu, Y., Zhu, S.-C., 2014. Cross-view action modeling, learning and recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2649–2656.
- Wang, L., Hu, W., Tan, T., 2003. Recent developments in human motion analysis. *Pattern Recognit.* 36 (3), 585–601.
- Wang, L., Qiao, Y., Tang, X., 2015a. Action recognition with trajectory-pooled deep-convolutional descriptors. *Proceedings of the Computer Vision and Pattern Recognition*. pp. 4305–4314.
- Wang, L., Xiong, Y., Lin, D., Van Gool, L., 2017a. Untrimmednets for weakly supervised action recognition and detection. *Proceedings of the Computer Vision and Pattern Recognition*.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L., 2016a. Temporal segment networks: towards good practices for deep action recognition. *Proceedings of the European Conference on Computer Vision*. pp. 20–36.
- Wang, P., Li, W., Gao, Z., Tang, C., Ogunbona, P., 2018a. Depth pooling based large-scale 3d action recognition with convolutional neural networks. *IEEE Trans. Multim.* 20 (5), 1051–1061.
- Wang, P., Li, W., Gao, Z., Tang, C., Zhang, J., Ogunbona, P.O., 2015b. Convnets-based action recognition from depth maps through virtual cameras and pseudocoloring. *Proceedings of the ACM MM*.
- Wang, P., Li, W., Gao, Z., Zhang, J., Tang, C., Ogunbona, P., 2016b. Action recognition from depth maps using deep convolutional neural networks. *THMS* 46 (4), 498–509.
- Wang, P., Li, W., Gao, Z., Zhang, Y., Tang, C., Ogunbona, P., 2017b. Scene flow to action map: a new representation for RGB-D based action recognition with convolutional neural networks. *Proceedings of the Computer Vision and Pattern Recognition*.
- Wang, P., Li, W., Liu, S., Gao, Z., Tang, C., Ogunbona, P., 2016c. Large-scale isolated gesture recognition using convolutional neural networks. *Proceedings of the 2016 Twenty-third International Conference on Pattern Recognition (ICPR)*. IEEE, pp. 7–12.
- Wang, P., Li, W., Liu, S., Zhang, Y., Gao, Z., Ogunbona, P., 2016d. Large-scale continuous gesture recognition using convolutional neural networks. *Proceedings of the 2016 Twenty-third International Conference on Pattern Recognition (ICPR)*. pp. 13–18.
- Wang, P., Li, W., Wan, J., Ogunbona, P., Liu, X., 2018. Cooperative training of deep aggregation networks for RGB-D action recognition. *Proceedings of the AAAI*.
- Wang, P., Li, Z., Hou, Y., Li, W., 2016e. Action recognition based on joint trajectory maps using convolutional neural networks. *Proceedings of the ACM Multimedia Conference*.
- Wang, P., Wang, S., Gao, Z., Hou, Y., Li, W., 2017c. Structured images for RGB-D action recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1005–1014.
- Wang, X., Farhadi, A., Gupta, A., 2016f. Actions transformations. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2658–2667.
- Wang, Y., Song, J., Wang, L., Van Gool, L., Hilliges, O., 2016a. Two-stream SR-CNNs for action recognition in videos. *Proceedings of the British Machine Vision Conference*.
- Weinzaepfel, P., Harchaoui, Z., Schmid, C., 2015. Learning to track for spatio-temporal action localization. *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3164–3172.
- Wu, D., Pigou, L., Kindermans, P.-J., Le, N.D.-H., Shao, L., Dambre, J., Odobez, J.-M., 2016. Deep dynamic neural networks for multimodal gesture segmentation and recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 38 (8), 1583–1597.
- Wu, D., Shao, L., 2014. Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition. *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 724–731.
- Xia, L., Chen, C.-C., Aggarwal, J., 2012. View invariant human action recognition using histograms of 3D joints. *Proceedings of the Computer Vision and Pattern Recognition Workshop*.
- Xingjian, S., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., Woo, W.-c., 2015. Convolutional LSTM network: a machine learning approach for precipitation now-casting. *Proceedings of the Advances in Neural Information Processing Systems*. pp. 802–810.
- Yan, X., Chang, H., Shan, S., Chen, X., 2014. Modeling video dynamics with deep dyne-coder. *Proceedings of the European Conference on Computer Vision*. Springer, pp. 215–230.
- Yang, X., Zhang, C., Tian, Y., 2012. Recognizing actions using depth motion maps-based histograms of oriented gradients. *Proceedings of the ACM MultiMedia*.
- Ye, M., Zhang, Q., Wang, L., Zhu, J., Yang, R., Gall, J., 2013. A survey on human motion analysis from depth data. *Time-of-Flight and Depth Imaging. Sensors, Algorithms, and Applications*. Springer, pp. 149–187.
- Yeung, S., Russakovsky, O., Jin, N., Andriluka, M., Mori, G., Fei-Fei, L., 2017. Every moment counts: dense detailed labeling of actions in complex videos. *Int. J. Comput. Vis.* 126, 375–389.
- Yeung, S., Russakovsky, O., Mori, G., Fei-Fei, L., 2016. End-to-end learning of action detection from frame glimpses in videos. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2678–2687.
- Yun, K., Honorio, J., Chattopadhyay, D., Berg, T.L., Samaras, D., 2012. Two-person interaction detection using body-pose features and multiple instance learning. *Proceedings of the 2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, pp. 28–35.
- Zhang, B., Wang, L., Wang, Z., Qiao, Y., Wang, H., 2016a. Real-time action recognition with enhanced motion vector CNNs. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2718–2726.
- Zhang, J., Li, W., Ogunbona, P., 2017a. Joint geometrical and statistical alignment for visual domain adaptation. *Proceedings of the Computer Vision and Pattern Recognition*.
- Zhang, J., Li, W., Ogunbona, P.O., Wang, P., Tang, C., 2016b. RGB-D-based action recognition datasets: a survey. *Pattern Recognit.* 60, 86–105.
- Zhang, P., Lan, C., Xing, J., Zeng, W., Xue, J., Zheng, N., 2017b. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. *Proceedings of the International Conference on Computer Vision*.

- Zhang, S., Liu, X., Xiao, J., 2017c. On geometric features for skeleton-based action recognition using multilayer lstm networks. *Proceedings of the Winter Conference on Applications of Computer Vision*.
- Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Lin, D., Tang, X., 2017. Temporal action detection with structured segment networks. *Proceedings of the International Conference on Computer Vision*.
- Zhu, F., Shao, L., Xie, J., Fang, Y., 2016a. From handcrafted to learned representations for human action recognition: a survey. *Image Vis. Comput.* 55, 42–52.
- Zhu, G., Zhang, L., Mei, L., Shao, J., Song, J., Shen, P., 2016b. Large-scale isolated gesture recognition using pyramidal 3D convolutional networks. *Proceedings of the 2016 23rd International Conference on Pattern Recognition (ICPR)*. IEEE, pp. 19–24.
- Zhu, G., Zhang, L., Shen, P., Song, J., 2017a. Multimodal Gesture Recognition Using 3D Convolution and Convolutional LSTM. *IEEE Access*.
- Zhu, H., Vial, R., Lu, S., 2017b. Tornado: A spatio-temporal convolutional regression network for video action proposal. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5813–5821.
- Zhu, W., Hu, J., Sun, G., Cao, X., Qiao, Y., 2016c. A key volume mining deep framework for action recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1991–1999.
- Zhu, W., Lan, C., Xing, J., Zeng, W., Li, Y., Shen, L., Xie, X., 2016d. Co-occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks. *Proceedings of the AAAI*.
- Zolfaghari, M., Oliveira, G.L., Sedaghat, N., Brox, T., 2017. Chained multi-stream networks exploiting pose, motion, and appearance for action classification and detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.