

# WEIGHTED MULTI-REGION CONVOLUTIONAL NEURAL NETWORK FOR ACTION RECOGNITION WITH LOW-LATENCY ONLINE PREDICTION

Yunfeng Wang<sup>\*</sup>, Wengang Zhou<sup>\*</sup>, Qilin Zhang<sup>†</sup>, Xiaotian Zhu<sup>\*</sup>, Houqiang Li<sup>\*</sup>

<sup>\*</sup>University of Science and Technology of China, Hefei, Anhui, China


<sup>†</sup>HERE Technologies, Chicago, Illinois, USA

## ABSTRACT

Spatio-temporal contexts are crucial in understanding human actions in videos. Recent state-of-the-art Convolutional Neural Network (ConvNet) based action recognition systems frequently involve 3D spatio-temporal ConvNet filters, chunking videos into fixed length clips and Long Short Term Memory (LSTM) networks. Such architectures are designed to take advantage of both short term and long term temporal contexts, but also requires the accumulation of a predefined number of video frames (e.g., to construct video clips for 3D ConvNet filters, to generate enough inputs for LSTMs). For applications that require low-latency online predictions of fast-changing action scenes, a new action recognition system is proposed in this paper. Termed “Weighted Multi-Region Convolutional Neural Network” (WMR ConvNet), the proposed system is LSTM-free, and is based on 2D ConvNet that does not require the accumulation of video frames for 3D ConvNet filtering. Unlike early 2D ConvNets that are based purely on RGB frames and optical flow frames, the WMR ConvNet is designed to simultaneously capture multiple spatial and short term temporal cues (e.g., human poses, occurrences of objects in the background) with both the primary region (foreground) and secondary regions (mostly background). On both the UCF101 and HMDB51 datasets, the proposed WMR ConvNet achieves the state-of-the-art performance among competing low-latency algorithms. Furthermore, WMR ConvNet even outperforms the 3D ConvNet based C3D algorithm that requires video frame accumulation. In an ablation study with the optical flow ConvNet stream removed, the ablated WMR ConvNet nevertheless outperforms competing algorithms.

**Index Terms**— Action Recognition, Low-Latency, Online Prediction, Multi-region, ConvNet

This work was supported in part to Dr. Houqiang Li by 973 Program under contract No. 2015CB351803 and NSFC under contract No. 61390514, and in part to Dr. Wengang Zhou by NSFC under contract No. 61472378 and No. 61632019, the Fundamental Research Funds for the Central Universities, and Young Elite Scientists Sponsorship Program By CAST (2016QNRC001).



**Fig. 1:** Background offers helpful auxiliary information in determining human actions. (a) Original frames. (b) Masked frame.

## 1. INTRODUCTION

Generally, there are two types of action recognition methods, one based on the handcrafted conventional features and the other based on deep neural networks [1, 2, 3, 4, 5]. Among the former type, the iDT [6] achieves the state-of-the-art performance but it is also excessively computationally expensive, due to its requirements of dense video trajectories. Thanks to the advancements in deep learning [7, 8], recent Convolutional Neural Network (ConvNet) based action recognition systems frequently involve 3D spatio-temporal ConvNet filters [9, 10] and chunking videos into fixed length clips. Such architectures are capable of exploiting even long term temporal contexts, but they require the accumulation of a predefined number of frames (e.g., to construct video clips for 3D ConvNet filters). For applications that require low-latency online predictions, alternative design of action recognition system need to be developed.

Alternatively, multi-stream ConvNet (typically with an optical flow stream) based methods have received tremendous attention in recent years, such as “Two-Stream” [11] and its derivatives [12, 13]. These methods are based on conventional 2D ConvNet, thus hardly requires video frame accumulation during prediction. Despite their lack of explicit 3D spatio-temporal ConvNet filters, they are still among the top performers on the UCF101 [14] and HMDB51 [15] datasets.

Despite the successes, these “Two-Stream” methods lack a holistic way of incorporating context information in video

frames. Traditionally viewed as a source of noise, the background itself arguably offers valuable auxiliary information [16, 17, 18] to recognize the characteristics of the human actions. A typical example is illustrated in Figure 1 (b). Even with the basketball player completely censored, the surrounding scene of the basketball court (especially the mid-air ball) provides substantial clues to recognize the “playing basketball” action.

Based on the aforementioned intuition, the “Weighted Multi-Region” (WMR) ConvNet is proposed to capture the context information from both the foreground and the background. Video frames are first fed into a fine-tuned Faster R-CNN [19] network to localize the “foreground” with a bounding box, which is termed the primary region. Subsequently, multiple secondary regions per video frame are obtained with the Selective Search algorithm [20] with varying intersection of Union (IoU) values with the primary region. Each testing image is fed into convolutional layers with the ROI-pooling [21] over the extracted primary region and secondary regions, as illustrated in Figure 2. After the final stages of the RGB image stream and optical flow image stream, a fusion layer is added to account for the contributions from both streams, based on which the final predictions are made.

In the testing phase, the proposed WMR ConvNet processes frame-by-frame in the spatial ConvNet (RGB input) and 10-frame-after-10-frame in the temporal ConvNet (Optical flow input). Therefore, a maximum delay in processing is about 0.3 second (assuming 30 fps). On contrary, conventional action recognition methods typically require 3D convolution and LSTM-style temporal postprocessing, which could incur much higher latency. For example, video clip-based methods need to process multiple clips (each clip last about 0.5 seconds) before generating a prediction, which leads to latency of several seconds to tens of seconds.

The major contributions of the paper are as follows.

- A new WMR ConvNet with “two-stream” 2D ConvNet architecture suitable for low-latency online prediction.
- An effective multi-region spatial sampling strategy that captures more informative spatial (RGB stream) and temporal (optical flow stream) contexts.
- The proposed WMR ConvNet achieves the state-of-the-art recognition accuracy on both the UCF101 and HMDB51 dataset, among competing low-latency algorithms.

## 2. RELATED WORK

Action recognition has been extensively studied in past few years [22, 6, 23, 11, 24, 9, 25, 10]. There are mainly two types of methods: conventional shallow feature based ones and those with deep neural network based ones. The improved Dense Trajectory (iDT) [6] method is the current state-of-the-art among the former type, in which the optical flow is utilized to estimate dense trajectories, followed by feature extractions around the trajectories. Inspired by the success

of deep learning in tasks of image understanding, many researchers have attempted to incorporate deep architectures for action recognition. A major difference between image based tasks and video based tasks is the extra temporal information in videos, which is critical for action recognition.

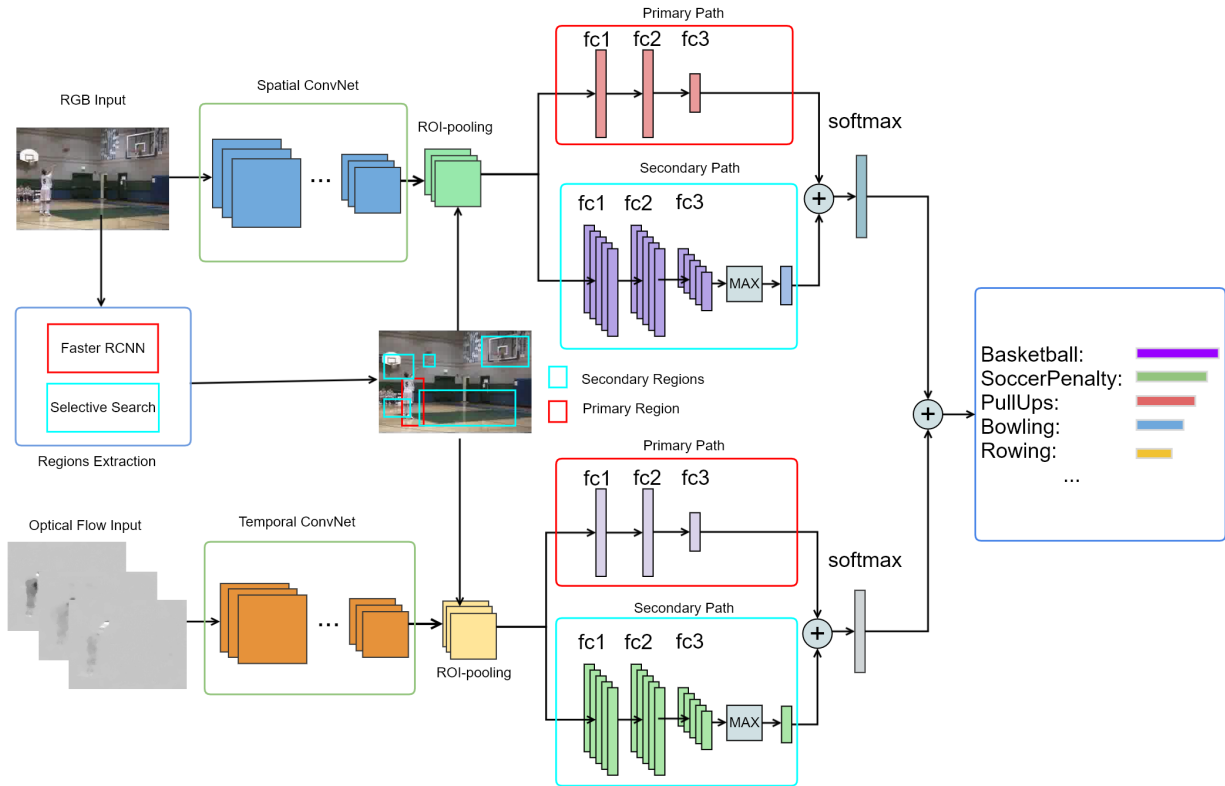
One of the most successful architectures in action recognition is two-stream network [11]. In that paper, the authors use one CNN stream with the RGB frame as input to capture spatial information and another CNN stream with the stacked optical flow frames as input to capture temporal information. At the end of the last softmax layer, a score fusion method is established to merge spatio-temporal information. Besides these 2D ConvNet based approaches (with low-latency online prediction capability), the C3D architecture proposed in [9] is proved to be a good tool of extracting features for videos. In that paper, the authors extend traditional two dimensional convolution to three dimensional convolution, which captures spatio-temporal information adequately. However, these 3D ConvNet [9, 13] based methods require video frame accumulation during prediction, thus unfit for action recognition applications with low-latency prediction requirements.

Since there are multiple cues in a still image, it’s feasible to run action recognition on image level. R\*CNN [26] is an efficient work in action recognition of still images, in which R-CNN [27] is adapted to multiple proposal regions. Different from this method, our proposed method is deployed on video action recognition task, and the UCF101 and HMDB51 video datasets have no ground truth bounding boxes.

Another work that is similar to our method is Dynamic Network [28], which summarizes the cues and information in video clips into a dynamic image, then feeds the dynamic image into a well-developed ConvNet to predict the label of the video.

## 3. PROPOSED WMR CONVNET

The proposed WMR ConvNet is illustrated in Figure 2. Given an image  $I$ , we first feed it into a fine-tuned Faster R-CNN [19] and Selective Search [20] pipeline to obtain primary region (object proposal region with the highest confidence score) and secondary regions (other object candidate proposal regions other than the primary region), respectively. Then we input these annotations and  $I$  to an adapted VGG-16 network [29]. An ROI Pooling layer is established to reuse the learnt feature maps in primary region and secondary regions. Then features of primary region and secondary regions are fed into three fully connected layers separately. In order to obtain the most informative secondary region for each action, we use max operation on the output of fully connected layers of the secondary region path. After obtaining maximum among secondary region scores, we add it to the primary score. Finally a softmax operation is established to transform scores into probabilities, which can be used to classify videos. Our full pipeline contains two paths of basic architecture: the one with



**Fig. 2:** WMR ConvNet architecture. The RGB input is fed into a spatial ConvNet and a regions extraction pipeline separately. We use Faster R-CNN to extract primary region, i.e., the region of human bounding box (red bounding boxes) and Selective Search to extract secondary region proposals, i.e., the regions with contexts and cues (cyan bounding box). After the ROI-pooling layer, primary region and selected secondary regions are fed into fc layers separately. Finally a softmax layer is set to fuse the scores from both regions. The stacked optical flow inputs reuse the regions of RGB input.

RGB input and the one with optical flow input. At the end of the pipeline, a fusion scheme is established to predict the final classification probability. Each part is detailed below.

**Region Proposal Generation.** Since the two main benchmarks of human action recognition in video have no bounding box annotations, we come up with a solution to use Faster R-CNN to annotate them. We first annotate bounding boxes manually for  $N_1$  and  $N_2$  RGB frames in training set of UCF101 and HMDB51, respectively. We only label the bounding box that contains the person in foreground. We then fine-tune on Faster R-CNN using only two classes: background and person. We choose the box with the max score of person as the annotation of the frame. In this way we want to train a model that circles out action performers in videos automatically. Compared with using pre-trained model directly, our method can annotate person in videos more accurately.

In order to obtain secondary regions, we first collect region proposals with Selective Search [20] on RGB images. Then we choose the region proposals which have an overlap with the primary region of images. Formally, these regions

can be defined as  $U(r; I)$ :

$$U(r; I) = \{s \in S(I) : \text{IoU}(s, r) \in [l, u]\}, \quad (1)$$

where  $r$  is the primary region of the image  $I$ ,  $S(I)$  is the selective search result of  $I$ ,  $\text{IoU}(s, r)$  denotes the overlap ratio between region  $s$  and  $r$ ,  $l$  and  $u$  is the lower and upper boundary of overlap range, respectively. In the experiments we use the faster approximate implementation of Selective Search in Dlib (<http://dlib.net/>). When coming to optical flow images, they have no distinct secondary regions since in the flow images the still objects and background information are wiped away. We find that during a small temporal interval, the displacement of background and objects is usually tiny. Based on this observation, we reuse the union of bounding boxes of secondary regions from contiguous RGB frames as the secondary regions of the optical flow image. Formally, we have

$$U(r; O_i) = U(r; I_i) \cup U(r; I_{i+1}), \quad (2)$$

where  $O_i$  is the  $i$ th optical flow frame of a video  $V$ ,  $I_i$  and  $I_{i+1}$  is the  $i$ th and  $i + 1$ th RGB frames of  $V$ , respectively. For the primary region of an optical flow image, there are

two ways to obtain it: 1) We get the rectangle region with the largest average magnitude as the bounding box. 2) We reuse the union of primary region of the two RGB frames where the optical flow is calculated. In a nutshell, by setting primary region and secondary regions reasonably for optical flow images, we extend the concept of cues and contexts to flow space.

**Learning Representations Using ConvNet.** After obtaining bounding boxes, we feed the training image into a convolutional network (ConvNet) to learn a representation of the image. In this work we use VGG-16 architecture because of its good performance and relatively low time-cost during training and testing. In order to train the secondary regions effectively, we adopt ROI pooling in our work as in [21]. Specifically, we replace the last pooling layer in VGG-16 with the ROI pooling layer, with the feature maps from previous layer and annotated region data as input. With it, each region generates a fixed-size feature map which can be fed into fully connected layers. Since the feature maps are all cropped from previous layer of ROI pooling (For example, the conv5-3 layer for VGG-16), no additional calculation is needed and this implementation is very efficient.

After the ROI pooling layer, we train the primary region and secondary regions separately to exploit potential cues. The primary region and secondary regions are separately fed to three fully connected layers to generate scores. Specifically, an extra max operation is carried out in the streams that process secondary regions. This max operation allows for the selection of the most informative secondary region for each action. Subsequently, this max score is fused with the primary region score with weights of 0.6 for score from primary region and 0.4 for score from secondary regions. At the very end, a softmax layer processes the scores to obtain the predicted probability of each class.

**Fusion of RGB Stream and Optical Flow Stream.** In the subsections above, we discussed the basic framework of our method. Although we demonstrate our framework using RGB data, we claim here that the input of the framework is also applicable to optical flow data. Inspired by the popular Two-Stream architecture in action recognition task, we feed RGB data and optical flow data into our basic framework separately, then combine the outputs of these two networks, as shown in Figure 2. For each video, 1 frame of RGB data and 10 frames of optical flow data are fed into the basic architecture, respectively. After that, we merge extracted spatio-temporal features using a fusion scheme to get final prediction with weights of 0.4 and 0.6 for RGB stream and optical flow stream, respectively.

#### 4. EXPERIMENTS

We evaluate our method on UCF101 [14] and HMDB51 [15]. UCF101 contains 13,320 video clips of 101 action classes collected from YouTube. HMDB51 [15] contains 6766 video

clips from movies and YouTube, annotated into 51 action classes.

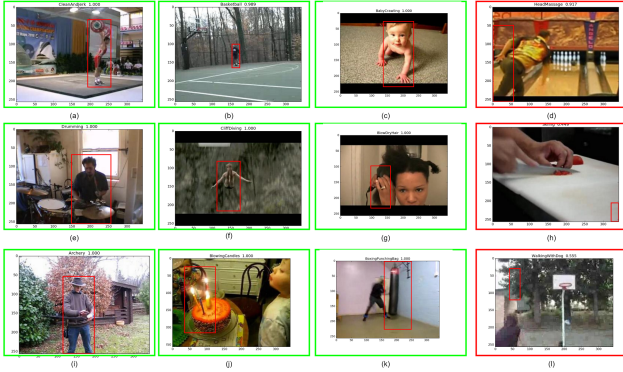
**ConvNet Configuration.** Our implementation is built on R-CNN [26] pretrained on ImageNet [30]. An ROI pooling layer is used to classify multiple ROI regions and two paths of fully connected layers for the primary region and secondary regions, respectively. We set initial learning rate to 0.0001 and divide it by 10 after every 50000 iterations. We set the dropout ratio to 0.6 according to ablation experiments (omitted here due to limited space). We set the batch size to 256 and use 2 images per batch. In total, our model is trained with 200K iterations. For more details on network configurations and ablation studies, please refer to the extended version [31]. During the training phase, we randomly choose 1 RGB frame from each training video as input, following [11]. For optical flow data, we choose 10 frames from each video. During the testing phase, we sample 25 frames of RGB data and flow data for each video.

**Evaluation.** In this subsection we conduct the human body detection and evaluate its impact on the final result. We label  $N_1 = 2525$  RGB images from UCF101,  $N_2 = 2177$  RGB images from HMDB51 to run transfer learning on Faster R-CNN. After fine-tuning the Faster R-CNN network, we input a testing image to the network and select the bounding box with the highest score of person category as the primary region, which is shown with a red box in each sub figure in Figure 3. In each row, the first three samples are correctly annotated (with a green box outside) and the fourth sample is wrongly annotated (with a red box outside).

From Figure 3 we find that the accuracy of annotation affects the final results heavily. For example, in (d), a hand of a person playing the bowling (the true label of this frame) is detected and its pose is similar to pose in “HeadMessage” videos the ConvNet has seen, thus the testing label is set to it. In (l), a small part of a green tree is detected, which is similar to green grasses in “WalkingWithDogs” videos. In (h), a white chopping board similar to the snowfield is detected so the ConvNet classifies this frame to “Skiing”. These examples show that the region detected as the primary region is vital for the action recognition task.

In order to obtain secondary regions, we run Selective Search firstly to get all region proposals. Then select regions among them using Equation (1) by setting  $l$  to 0.1 and  $u$  to 0.9. We use two approaches to get the primary regions of optical flow images. The results are shown in Table 1. The approach reusing bounding boxes of RGB image is better, which indicates that the cues in RGB data are more useful than magnitude of optical flow.

**Feature Fusion Schemes.** In addition to weighted sum, we explore another classifier to fuse features of the two-path networks. We propose two alternatives here: 1) For features extracted from fully connected layers, we concatenate features from primary region and secondary regions. After that, we use a linear SVM to classify them. 2) For scores obtained



**Fig. 3:** The human detection results. Green boxes denote correct detection; while red boxes indicate wrong detections.

**Table 1:** Optical flow ( $u$  and  $v$  channel) results with 2 primary region annotation methods. 1), a rectangle with max average magnitude is selected; 2), the bounding box of corresponding RGB frame is selected. We conduct experiments on  $u$  channel and  $v$  channel separately.

Method	Accuracy ( $u$ channel)	Accuracy ( $v$ channel)
1)	75.3%	75.3%
2)	<b>76.2%</b>	<b>77.4%</b>

after softmax layer, we use either weighted sum or a linear SVM. In Table 2, we find that results with weighted sum operation are the best. The feature of fc1 layer and fc2 layer are worse than score layer features and weighted sum outperforms SVM on score layer features.

Here we compare the proposed ablated WMR ConvNet with the dynamic image networks (MDI, MDI + static RGB). Only RGB is used for fair comparison. Table 3 shows the results of our method and dynamic image networks. We find that our method outperforms the dynamic image networks. This indicates the necessity of encoding the cues in video for our method.

From Table 4, we find that our method is comparable with iDT and C3D on UCF101 dataset, and outperforms other methods on HMDB51 dataset. Note that C3D [9] requires video frame accumulation during prediction, thus it is not directly compatible with low-latency online prediction requirements. This indicates that our proposed WMR ConvNet is effective and spatio-temporal cues and contexts are crucial for human action recognition in videos. Here we compare our results on spatial stream ConvNet with other published results. Table 5 shows that our method outperforms the compared methods in spatial domain. This superior performance demonstrates the effectiveness of exploiting multi-region based cues in spatial space and indicates that our method is more suitable for low-latency usage than compared methods.

**Table 2:** Fusion methods comparison with RGB data of UCF101 split-1.

Fusion method	Accuracy
fc1 + SVM	61.1%
fc2 + SVM	66.5%
score + SVM	73.7%
score + weighted sum	<b>75.9%</b>

**Table 3:** Comparison with the Dynamic Network [28]. Only RGB is used for fair comparison.

Method	UCF101
static RGB	70.1%
MDI	70.9%
MDI + static RGB	76.9%
ablated WMR ConvNet (RGB only)	<b>78.8%</b>

## 5. CONCLUSIONS

In this paper, we propose the WMR ConvNet for action recognition with low-latency online prediction requirements. An efficient multi-region spatial sampling strategy is proposed to explicitly capture both human related information (e.g., poses) and context information, via primary region and secondary regions, respectively. Such regions are also shared by the optical flow pipeline, which offers additional temporal information. Even with the optical flow pipeline removed, the ablated WMR ConvNet achieves a high recognition accuracy of 78.8% on the UCF101 dataset. The complete WMR ConvNet outperforms all competing 2D ConvNet based algorithms (with low-latency online prediction) and even a 3D ConvNet based C3D algorithm on the UCF101 dataset.

## 6. REFERENCES

- [1] Jie Huang, Wengang Zhou, Qilin Zhang, Houqiang Li, and Weiping Li, "Video-based sign language recognition without temporal segmentation," in *AAAI, New Orleans, Louisiana, USA, 2018*.
- [2] Yunfeng Wang, Wengang Zhou, Qilin Zhang, and Houqiang Li, "Visual attribute-augmented three-dimensional convolutional neural network for enhanced human action recognition," *arXiv preprint arXiv:1805.02860*, 2018.
- [3] Xuhuan Duan, Le Wang, Changbo Zhai, Qilin Zhang, Zhenzong Niu, Nanning Zheng, and Gang Hua, "Joint spatio-temporal action localization in untrimmed videos with per-frame segmentation," in *ICIP, Athens, Greece, 2018*.
- [4] Xin Lv, Le Wang, Qilin Zhang, Zhenzong Niu, Nanning Zheng, and Gang Hua, "Video object co-segmentation from noisy videos by a multi-level hypergraph model," in *ICIP, Athens, Greece, 2018*.
- [5] Jinliang Zang, Le Wang, Ziyi Liu, Qilin Zhang, Zhenxing Niu, Gang Hua, and Nanning Zheng, "Attention-based temporal

**Table 4:** Comparison of proposed WMR ConvNet with competing methods on UCF101 and HMDB51 datasets.

Method	UCF101	HMDB51
iDT [6]	85.9%	57.2%
Slow Fusion [32]	65.4%	-
Two Stream [11]	<b>88.0%</b>	59.4%
C3D [9]	85.2%	-
WMR ConvNet	85.7%	<b>66.7%</b>

**Table 5:** Spatial stream (RGB) ConvNet on UCF101.

Method	Accuracy
Two-Stream [11]	73.0%
DIN [28]	76.9%
SR-CNN [12]	78.3%
RGB-only WMR ConvNet	<b>78.8%</b>

weighted convolutional neural network for action recognition,” in *AIAI, Rhodes, Greece*, 2018.

- [6] Heng Wang and Cordelia Schmid, “Action recognition with improved trajectories,” in *ICCV*, 2013.
- [7] Lingyan Ran, Yanning Zhang, Wei Wei, and Qilin Zhang, “A hyperspectral image classification framework with spatial pixel pair features,” *Sensors*, vol. 17, no. 10, pp. 2421, 2017.
- [8] Lingyan Ran, Yanning Zhang, Qilin Zhang, and Tao Yang, “Convolutional neural network-based robot navigation using uncalibrated spherical images,” *Sensors*, vol. 17, no. 6, pp. 1341, 2017.
- [9] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri, “Learning spatiotemporal features with 3d convolutional networks,” in *ICCV*, 2015.
- [10] Joao Carreira and Andrew Zisserman, “Quo vadis, action recognition? a new model and the kinetics dataset,” *arXiv preprint arXiv:1705.07750*, 2017.
- [11] Karen Simonyan and Andrew Zisserman, “Two-stream convolutional networks for action recognition in videos,” in *NIPS*, 2014.
- [12] Yifan Wang, Jie Song, Limin Wang, Luc Van Gool, and Otmar Hilliges, “Two-stream sr-cnns for action recognition in videos,” *BMVC*, 2016.
- [13] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool, “Temporal segment networks: towards good practices for deep action recognition,” in *ECCV*, 2016.
- [14] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah, “Ucf101: A dataset of 101 human actions classes from videos in the wild,” *arXiv:1212.0402*, 2012.
- [15] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre, “HMDB: a large video database for human motion recognition,” in *ICCV*, 2011.
- [16] Qilin Zhang, Gang Hua, Wei Liu, Zicheng Liu, and Zhengyou Zhang, “Can visual recognition benefit from auxiliary information in training?,” in *ACCV, Singapore*, 1-5 November, 2014, pp. 65–80.
- [17] Qilin Zhang, Gang Hua, Wei Liu, Zicheng Liu, and Zhengyou Zhang, “Auxiliary training information assisted visual recognition,” *IPSP Trans. Comput. Vis. and Appl.*, vol. 7, pp. 138–150, 2015.
- [18] Qilin Zhang and Gang Hua, “Multi-view visual recognition of imperfect testing data,” in *ACM Multimedia, Brisbane, Australia*, 2015, pp. 561–570.
- [19] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *NIPS*, 2015.
- [20] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders, “Selective search for object recognition,” *International Journal of Computer Vision*, vol. 104, no. 2, pp. 154–171, 2013.
- [21] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik, “Region-based convolutional networks for accurate object detection and segmentation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 1, pp. 142–158, 2016.
- [22] Heng Wang, Alexander Kläser, Cordelia Schmid, and Cheng-Lin Liu, “Action recognition by dense trajectories,” in *CVPR*, 2011.
- [23] Xingyang Cai, Wengang Zhou, and Houqiang Li, “An effective representation for action recognition with human skeleton joints,” in *Optoelectronic Imaging and Multimedia Technology III*. International Society for Optics and Photonics, 2014, vol. 9273, p. 92731R.
- [24] Xingyang Cai, Wengang Zhou, and Houqiang Li, “Attribute mining for scalable 3d human action recognition,” in *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 2015, pp. 1075–1078.
- [25] Xingyang Cai, Wengang Zhou, Lei Wu, Jiebo Luo, and Houqiang Li, “Effective active skeleton representation for low latency human action recognition,” *IEEE Transactions on Multimedia*, vol. 18, no. 2, pp. 141–154, 2016.
- [26] Georgia Gkioxari, Ross Girshick, and Jitendra Malik, “Contextual action recognition with r\*cnn,” in *ICCV*, 2015.
- [27] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” in *CVPR*, 2014.
- [28] Hakan Bilen, Basura Fernando, Efstratios Gavves, Andrea Vedaldi, and Stephen Gould, “Dynamic image networks for action recognition,” in *CVPR*, 2016.
- [29] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [30] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR*, 2009.
- [31] Yunfeng Wang, Wengang Zhou, Qilin Zhang, Xiaotian Zhu, and Houqiang Li, “Low-latency human action recognition with weighted multi-region convolutional neural network,” *arXiv preprint arXiv:1805.02877*, 2018.
- [32] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei, “Large-scale video classification with convolutional neural networks,” in *CVPR*, 2014.