

Viewpoint Invariant Action Recognition using RGB-D Videos

Jian Liu^a, Naveed Akhtar^{a,*}, Ajmal Mian^a

^a*School of Computer Science and Software Engineering, The University of Western
Australia, 35 Stirling Highway, Crawley, 6009 WA, Australia.*

Abstract

In video-based action recognition, viewpoint variations often pose major challenges because the same actions can appear different from different views. We use the complementary RGB and Depth information from the RGB-D cameras to address this problem. The proposed technique capitalizes on the spatio-temporal information available in the two data streams to extract action features that are largely insensitive to the viewpoint variations. We use the RGB data to compute dense trajectories that are translated to viewpoint insensitive deep features under a non-linear knowledge transfer model. Similarly, the Depth stream is used to extract CNN-based view invariant features on which Fourier Temporal Pyramid is computed to incorporate the temporal information. The heterogeneous features from the two streams are combined and used as a dictionary to predict the label of the test samples. To that end, we propose a sparse-dense collaborative representation classification scheme that strikes a balance between the discriminative abilities of the dense and the sparse representations of the samples over the extracted heterogeneous dictionary. To establish the effectiveness of our approach, we benchmark it on three standard datasets and compare its performance with twelve existing methods. Experiments show that the proposed approach achieves up to 7.7% improvement in the accuracy over its nearest competitor.

*Corresponding author

Email address: naveed.akhtar@research.uwa.edu.au (Naveed Akhtar)

Keywords: Action recognition, viewpoint invariance, RGB-D videos.

1. Introduction

Video-based human action recognition is a challenging problem because of the large intra-action variations resulting from different illumination conditions, scales of videos, object textures and different scene backgrounds. In the real-world scenarios, the task becomes even more arduous due to the camera viewpoint variations. Recent advances in the Depth video technology has resulted in mitigating the problems arising from texture and illumination variations in the scenes because the Depth data is largely insensitive to such diversities. Nevertheless, the Depth videos still remain easily influenced by the camera viewpoint variations. To incorporate viewpoint invariance in video-based human action recognition, we propose to fuse view-invariant deep features from RGB and Depth data streams and use the resulting heterogeneous features for action recognition.

It is a common practice in the RGB video based action recognition [1, 2, 3, 4] to extract the optical flow and dense trajectory features from the data and use them for predicting the action labels. RGB videos generally contain high-fidelity appearance information about the scenes and the extracted dense trajectories provide important temporal cues that are useful in classifying an action. However, dense trajectories are sensitive to the camera viewpoints [5]. Moreover, such features do not contain any explicit information regarding the human poses in the videos. It is easy to imagine that human pose information can be useful in a more accurate action recognition in realistic conditions.

Rahmani and Mian [6] recently proposed a Non-linear Knowledge Transfer Model (NKTm) for cross-view action recognition. Their model maps dense trajectory features of RGB videos to a single canonical viewpoint, that makes their approach less sensitive to the viewpoint variations. Nevertheless, their technique is unable to take full advantage of the human pose information because it is harder to extract such details using only the RGB data. Comparatively, the

Depth data is much more suitable for extracting the human poses. Thus, their approach becomes a sub-optimal choice for the cases where the Depth data is also easily available. To take advantage of the human pose information in videos, Rahmani and Mian [5] separately proposed a Human Pose Model (HPM), that extracts view-invariant features using the Depth stream only. Accurate action recognition using the HPM features [5] substantiates the usefulness of the pose information in action recognition.

With the recent easy availability of the RGB-D sensors (e.g. Kinect), it is intuitive to jointly exploit both RGB and Depth data streams for action recognition. Existing approaches [7, 8, 9, 10, 11, 12] also report reasonable performance gain with the joint usage of the two data streams. However, none of these approaches are tailored to incorporate viewpoint invariance in action recognition. Whereas RGB and Depth data streams of RGB-D videos naturally complement each other, specialized processing of each stream is imperative for successful viewpoint invariant action recognition. Moreover, to take full advantage of both data streams it is important to carefully combine them to amplify the view-invariant character of the resulting data.

In this work, we propose a view-invariant action recognition approach, illustrated in Fig. 1, that accounts for the aforementioned observations. The proposed approach, that builds on our recent work [13], simultaneously extracts view-invariant features from the RGB and the Depth data streams and meticulously combines these heterogeneous features for better classification. For the Depth features, it exploits the CNN-based Human Pose Model [5] and performs temporal encoding of the extracted CNN features using the Fourier Temporal Pyramid. For the RGB stream, we first extract dense trajectory features and encode them with a representative code book. The resulting codes are passed through another deep network model to extract the desired view invariant features. The network is implemented using the Caffe library [14] and it has been learned from scratch for extracting the view-invariant features.

Our approach carefully combines the RGB and the Depth features and uses the joint features as a dictionary to classify the test samples. To that end, we

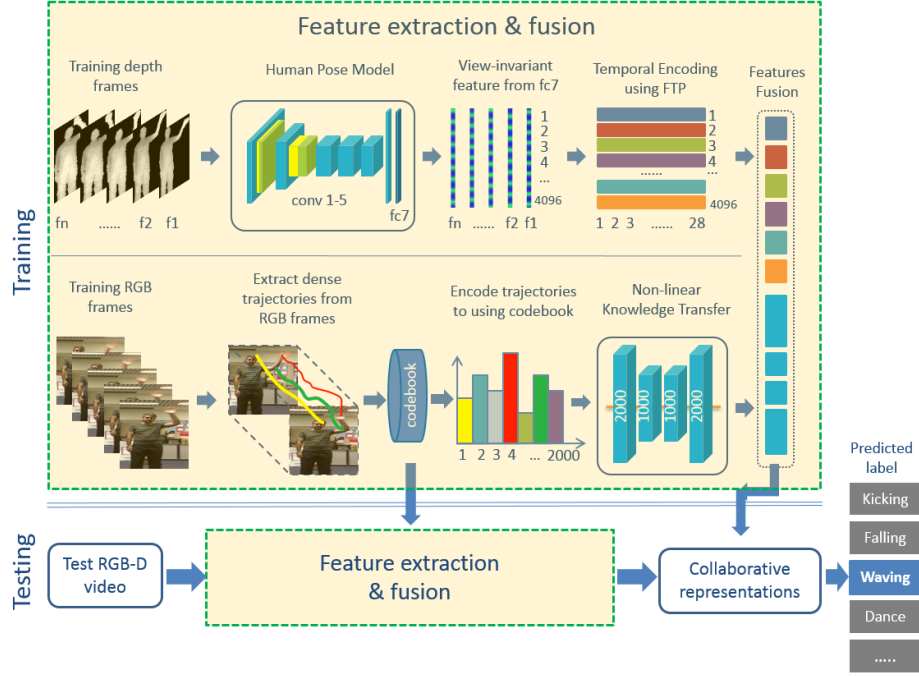


Figure 1: Schematics of the proposed approach: During training, the Depth stream extracts CNN features using the Human Pose Model [5] and performs temporal encoding using the FTP. The RGB stream is used to extract dense trajectories and a deep network, implementing non-linear knowledge transfer, is employed to extract the view invariant features. The resulting heterogeneous features are fused together and used as a set of basis vectors in collaborative representation based classification of the test samples.

propose to take advantage of both dense and sparse representations of a test feature over the extracted heterogeneous dictionary. We propose a classification scheme that balances between the discriminative abilities of dense and sparse representations to achieve the optimal performance under the collaborative representation based classification framework [15]. We evaluate our approach on three standard multi-view action recognition datasets and compare its performance with twelve existing techniques. Results of the performed experiments clearly establish the effectiveness of the proposed approach.

The remaining article is organized as follows. In Section 2, we discuss the related existing literature in RGB, Depth and RGB-D action recognition. We

explain the proposed technique in Section 3 and report detailed experimental results in Section 4. The paper is concluded in Section 5.

2. Related work

RGB and Depth (D) are the two major video data formats used in human action recognition. In this section, we discuss the related work relevant to each of these formats as well as the combined RGB-D format.

2.1. RGB Video Human Action Recognition

In RGB video based action recognition, few exiting approaches [16, 17, 18] directly use geometric transformations to incorporate the much needed view-point invariance. However, to achieve the desired performance level, it is critical for these methods to accurately estimate the skeleton joints. In practical conditions, it is often challenging to achieve high level of accuracy in skeleton joint estimation, which makes these methods less appealing for the practical purpose. Another stream of techniques [19, 20, 21, 22] exploits spatio-temporal features in the RGB videos to incorporate the viewpoint invariance. Nevertheless, the action recognition performance of these approaches is generally limited by the structure of the extracted features [5].

Another popular framework in RGB video based view-invariant action recognition is to find a latent space where the features are insensitive to viewpoint variations and classify the actions employing that latent space [23, 24, 25, 26, 27, 28, 29]. A combination of hand-crafted features and deep-learned features was also proposed by Wang et al. [30] for the RGB action recognition. In their approach, trajectory pooling was used for one stream of the data and deep learning framework was used for the other. The two feature streams were combined to form trajectory-pooled deep-convolutional descriptors. Nevertheless, only RGB data is used in their approach and the problem of viewpoint variance in action recognition is not directly addressed.

2.2. Depth Video Human Action Recognition

With the easy availability of the Depth data through the Microsoft Kinect sensor, the Depth video based action recognition became much popular in the last decade. In [31] and [32], 3D data points are used at silhouettes of the Depth images and 3D joint positions to extract features for action recognition. A binary range-sample feature was proposed for the Depth videos by Lu *et al* [33], that demonstrated significant improvement in achieving viewpoint invariance. Rahmani *et al* [34] proposed the Histogram of Oriented Principal Components (HOPC) to detect interest points in Depth videos and extracting their spatio-temporal descriptors. HOPC extracts local features in an object-centered local coordinate basis, thereby making them viewpoint invariant. Nevertheless, these features must be extracted at a large number of interest points that makes the overall approach computationally expansive. In another work, Yang *et al.* [35] clustered hypersurface normals in the Depth sequences to characterize the local motions and the shape information. An adaptive spatio-temporal pyramid is used in their approach to divide the Depth data into a set of space-time grids and low-level polynomials are aggregated to form a Super Normal Vector (SNV). This vector is eventually employed in action recognition.

2.3. RGB-D Video Human Action Recognition

Most of the Depth sensors also provide simultaneous RGB videos. This fact has lead to a significant interest of the scientific community to jointly exploit the two data streams for various tasks, including action recognition [7, 10, 11, 12, 9, 8, 36]. For instance, a restricted graph-based genetic programming approach is proposed by Liu and Shao [9] for the fusion of the Depth and the RGB data streams for improved RGB-D data based classification. In another approach, Hu *et al.* [8] proposed to learn heterogeneous features for the RGB-D video based action recognition. They proposed a joint heterogeneous features learning model (JOULE) to take advantage of both shared and action-specific components in the RGB-D videos. Kong and Fu [36] also projected and compressed both Depth

and RGB features to a shared feature space, in which the decision boundaries are learned for the classification purpose.

Whereas one of the major advantages of the Depth videos is in the easy availability of the information useful for the problem of viewpoint invariant action recognition, none of the aforementioned approaches directly address this problem. Moreover, the Depth and the RGB video frames are mainly combined in those approaches by either projecting them to a common feature space [8, 36] or by using the same filtering-pooling operation for both modalities [9]. We empirically verified that on the used action recognition datasets that involve multiple camera views, these techniques achieve no more than 4% improvement (RGB-D combined) over the single (RGB or Depth) modality. On the other hand, the technique proposed in this work achieves up to 7.7% average improvement over the single modality while dealing with the viewpoint variations. The strength of our approach resides in processing the RGB and the Depth streams individually to fully capitalize on their individual characteristics and then fusing the two modalities at a latter stage of the pipeline. This strategy has proven much more beneficial than combining the two data streams earlier in data processing.

3. Proposed approach

The proposed RGB-D based human action recognition approach is illustrated in Fig. 1. We use a Deep Learning based Human Pose Model [37] to extract features from the Depth data stream. These features are post processed by computing Fourier Temporal Pyramids. For the RGB stream, we extract dense trajectory features and encode them over a codebook. The resulting codes are passed through a network implementing a non-linear knowledge transfer model. The two types of features are combined and used as a dictionary in our approach. To classify, we first extract the joint features of the test samples and then collaboratively represent them over the extracted dictionary. These representations are used in predicting the class labels.

3.1. Depth Feature Extraction

We extract CNN-based features from the Depth data stream using a pre-trained Human Pose Model (HPM) [5]. HPM uses an architecture structure similar to the AlexNet [38] but it has been trained using multi-viewpoint action data that is generated synthetically by fitting human models to the CMU motion capture data [39]. Using data from 180 different viewpoints for each action makes the network training largely insensitive to the view variations in the actions. HPM essentially extracts features using Depth *images* with static human poses. Thus, to incorporate the temporal dimension we use the Fourier Temporal Pyramid (FTP) [40] on the HPM features.

Let us denote the t^{th} frame of the i^{th} Depth video by \mathbf{V}_t^i , where $t \in \{1, 2, \dots, f\}$, such that f represents the total number of frames. We first crop a Depth frame $\mathbf{V}^i \in \{\mathbf{V}_1^i, \mathbf{V}_2^i, \dots, \mathbf{V}_f^i\}$ to resize it to a 227×227 image in order to match the input dimensions of the HPM. These frames are passed through the HPM and the *fc7* layer output activation vector $\mathbf{a}_t^i \in \mathbb{R}^{4096}$ is used as the feature for the Depth frame \mathbf{V}_t^i . The features from the i^{th} video are combined into a matrix $\mathbf{A}^i = [\mathbf{a}_1^i, \mathbf{a}_2^i, \dots, \mathbf{a}_f^i] \in \mathbb{R}^{4096 \times f}$ and the Fourier Temporal Pyramid of the matrix is computed for the temporal encoding. More specifically, the computed pyramid has 3 levels, and by dividing the feature set \mathbf{A}^i in half at each level, we get $1 + 2 + 4 = 7$ feature groups. Short Fourier Transform is applied to each feature group, and the first 4 low-frequency coefficients are concatenated (i.e. $4 \times 7 = 28$) to form a spatial-temporal descriptor matrix $\mathbf{S}^i \in \mathbb{R}^{4096 \times 28}$ for an action. Finally, the matrix \mathbf{S}^i is vectorized as $\mathbf{d}^i \in \mathbb{R}^{114688}$ to get the Depth feature for the i -th Depth video.

3.2. RGB Feature Extraction

For our approach, we employ another neural network to extract the RGB features. The network model is implemented using the Caffe library [14] and it has been trained from the scratch for the proposed approach. Momentarily, we defer the discussion on the model details and describe the preprocessing of the data for its training. Similar to the HPM, we use the CMU motion capture




Figure 2: Caffe implementation of the used network. Dropout and sigmoid activation are applied on the first 3 inner product layers.

data [39] for training this network. The data pre-processing is conducted as follows. We first fit cylinders to the available skeleton data in order to approximate the human limbs and the torso. The points on those cylinders are then rendered from eighteen different viewpoints using their orthographic projections onto the x-y plane. We extract the dense trajectory features [1] from the resulting data. In order to do that, we follow Gupta et al. [1] and fix the frame length to $L = 15$. This results in thirty-dimensional trajectory features. Notice that we again use the data generated from multiple viewpoints, which is important to introduce viewpoint invariance in the RGB features, similar to the Depth features. Once the trajectory features are extracted, we perform clustering over these features using the K-Means algorithm, where we empirically choose $K = 2000$. We store the cluster centroids in $\mathbf{C} \in \mathbb{R}^{30 \times 2000}$. The matrix \mathbf{C} represents a codebook, containing the 2000 most representative trajectories for the data.

To train the network, a video is first encoded over the matrix \mathbf{C} . The coding is performed such that the resulting Bag of word $\boldsymbol{\xi} \in \mathbb{R}^{2000}$ represents a histogram of video trajectories closest to each column of the matrix \mathbf{C} in the Euclidean space. Our neural network implements the non-linear knowledge transfer model [6] that projects the vectors $\boldsymbol{\xi}$ of all the viewpoints of an action to a single canonical (i.e. frontal) viewpoint. The model essentially regresses for the problem of projecting multiple features to a single feature, thereby not requiring any explicit label during the training. The architecture of the network




Figure 3: Three network architectures trained using the Caffe library [14].

based on our Caffe implementation is shown in Fig. 2. The network comprises four layers with drop-outs and sigmoid activations applied to the first three layers. The main intuition behind the chosen architecture is that since we must keep the regression flexible, in the sense that it is able to project different views to a single view, we first reduce the layer size to drop redundant information in the features. In the latter layer, the size is again increased to discriminate between the view specific high level details of the inputs.

In light of the original proposal of the non-linear knowledge transfer approach by Rahmani and Mian [6], we tested multiple architectural variations of the above mentioned network. In Fig 3, we show the three variations for which we also present the experimental results in Section 4 to illustrate the sensitivity of our approach to the architecture of the network. The chosen network is labeled NKTM #1 in Fig. 3 (for the Non-linear Knowledge Transfer Model). For training, we initialize the network with the weight filler “xavier” and variance norm “AVERAGE”, and configure the loss layer as the “EuclideanLoss”. We use the initial learning rate of 0.001 for all the layers and decrease it by a factor of 10 after every 1000 iterations. The weight decay of the network is configured to 0.0005. We train our network using the back-propagation algorithm with six thousand iterations. These parameters are chosen empirically in our approach.

Once the network is trained using the CMU mocap data we use it to extract features of the action recognition training data. For that, we first extract their dense trajectories following [2, 3]. These trajectories are coded over the already learned codebook \mathbf{C} and the resulting codes are passed through the trained

network. In order to finally arrive at the view invariant features, we concatenate the outputs of each layer of the network. It was empirically verified for the used network that its earlier layers are more informative for the small variations in the viewpoints whereas the latter layers are more informative for the larger viewpoint variations. Therefore, we concatenate the outputs from all the layers to construct the RGB feature vector $\mathbf{r} \in \mathbb{R}^{6000}$.

3.3. Feature Fusion

The RGB and the Depth features are extracted from two different data streams in our approach. However, they contain complementary information due to the following reasons. Firstly, the data streams originate from the same actions. Secondly, the dense trajectories model the temporal information from the RGB frames using only the motion (without explicitly encoding the poses), whereas the Depth features are extracted by modeling the human poses from the raw Depth images (and using the FTP). Thus, the latter is specifically meant to complement the former. It is also worth noting that the dense trajectories are better extracted from the RGB frames due to the presence of texture. On the other hand, human poses are more accurately computed using the Depth images because they contain the shape information.

To integrate the strengths of the RGB and the Depth features, we propose to fuse them into a joint feature. We use these heterogeneous joint features as the basis vectors to collaboratively represent the test sample features and classify the resulting representations. The proposed feature fusion scheme works as follows. Let $\mathbf{D} = [\mathbf{d}^1, \mathbf{d}^2, \dots, \mathbf{d}^n] \in \mathbb{R}^{114688 \times n}$ and $\mathbf{R} = [\mathbf{r}^1, \mathbf{r}^2, \dots, \mathbf{r}^n] \in \mathbb{R}^{6000 \times n}$ denote the feature sets obtained from the Depth and the RGB streams respectively. We first transform \mathbf{D} and \mathbf{R} into $\mathbf{D}_z \in \mathbb{R}^{114688 \times n}$ and $\mathbf{R}_z \in \mathbb{R}^{6000 \times n}$, where the latter matrices are computed as the Z-scores of the formers. Then, the columns of \mathbf{D}_z and \mathbf{R}_z are rescaled to the range $[0, 1]$. Finally, the rescaled features are row-wise concatenated to form the heterogeneous feature set $\mathbf{X} \in \mathbb{R}^{120688 \times n}$. The procedure is also illustrated in Fig. 4.

The feature fusion procedure is specifically designed keeping in view the




Figure 4: Illustration of feature fusion process. First, the Z-scores of the Depth and the RGB features are computed. Then, sample-wise scaling is performed to restrict features in range $[0,1]$. Finally, a row-wise concatenation is performed.

requirements of the proposed approach. In our framework, the obtained Depth and RGB features generally vary greatly in terms of scale and dimensionality. However, we still aim at using the combined features as the joint basis vectors to represent the test sample features. Therefore, we first transform the features to a dimensionality independent measure of Z-score, which represents each coefficient of a feature vector in terms of the standard deviation of the vector components. A further rescaling is performed to remove the scale differences of the extracted features. Once the rescaled features are combined, they form an effective set of basis vectors for the subspace in which the test sample features reside.

3.4. Classification

Our label prediction stage is inspired by the Collaborative Representation (CR) based classification framework [41]. To predict the label of a test sample, we first extract its joint feature $\mathbf{y} \in \mathbb{R}^{114688}$ following the feature extraction procedure for the training data. Then, we compute two separate collaborative representations of the test feature in terms of the training features by solving the following optimization problems:

$$\check{\alpha} = \min_{\alpha} \|\mathbf{y} - \mathbf{X}\alpha\|_2 + \lambda \|\alpha\|_2 \quad (1)$$

$$\hat{\alpha} = \min_{\alpha} \|\mathbf{y} - \mathbf{X}\alpha\|_2, \text{ s.t. } \|\alpha\|_0 \leq k, \quad (2)$$

where, $\|\cdot\|_p$ denotes the ℓ_p -norm of a vector and k is the sparsity threshold.

The vector $\check{\alpha} \in \mathbb{R}^n$ is computed by regularizing the ℓ_2 -norm of the coefficient vector, that allows most of the components of $\check{\alpha}$ to be non-zero in the representation. This makes $\check{\alpha}$ the *dense* representation of \mathbf{y} over the matrix \mathbf{X} . On the other hand, the sparsity constraint in (2) forces most of the coefficients in the representation vector to become zero, thereby making $\hat{\alpha} \in \mathbb{R}^n$ the *sparse* representation of \mathbf{y} . Since the basis vectors in \mathbf{X} have heterogeneous composition in our approach, CR-based classification framework can be expected to favor sparse representations in few instances and dense ones for the other. Therefore, we propose to combine the two representations in a convex manner and use this combination for predicting the class labels of the test samples. More precisely, we compute $\overset{\circ}{\alpha} = \lambda_1 \hat{\alpha} + (1 - \lambda_1) \check{\alpha}$, where $\lambda_1 \leftarrow [0, 1]$ and later employ $\overset{\circ}{\alpha} \in \mathbb{R}^n$ to predict the class label. We note that this representation combination strategy differs from [41] that augments a dense collaborative representation with a sparse one by simply adding the two and normalizing the resulting vector. Convex combination of the two types of representations ascertain global optimality of the solution with respect to the parameter λ_1 . This guarantee is not provided by the data augmentation method proposed in [41].

In CR-based classification, it is generally the case that the class label is predicted either by maximizing the reconstruction fidelity of the collaborative representation coefficients [15] or by using a multi-class classifier in conjunction with the computed representation [42], [43]. However, these methods are computationally expensive. We recently showed [41] that collaborative representations can be used for predicting the class labels much more efficiently by integrating over their coefficients for each class and assigning the class label to the largest integrated value. Here, we implement this strategy by multiplying $\overset{\circ}{\alpha}$ with a binary matrix $\mathbf{B} \in \mathbb{R}^{C \times n}$, where C denotes the total number of classes. This matrix is constructed by assigning 1 to the i^{th} row of its j^{th} column if the j^{th} column of \mathbf{X} belongs to the i^{th} class of the training features. All other coefficients of the matrix are kept 0. With a matrix thus constructed, each coefficient of the vector $\mathbf{q} = \mathbf{B} \overset{\circ}{\alpha} \in \mathbb{R}^C$ integrates the coefficients of $\overset{\circ}{\alpha}$ for a single class. We maximize over the coefficients of \mathbf{q} to predict the class label of the test sample.

The proposed classification procedure is also summarized as Algorithm 1.

Algorithm 1 Sparse-Dense CR-based Classification

Input: (a) Training features \mathbf{X} , normalized in ℓ_2 -norm. (b) Test sample feature \mathbf{y} . (c) Regularization parameters λ, λ_1 . (d) Sparsity threshold k . (e) Binary matrix \mathbf{B} .

1: *Optimization:*

a) Solve (1) as $\check{\alpha} = \mathbf{P}\mathbf{y}$, where, $\mathbf{P} = (\mathbf{X}^T\mathbf{X} + \lambda\mathbf{I}_n)^{-1}\mathbf{X}^T$.

b) Solve (2) using the Orthogonal Matching Pursuit algorithm [44]:

2: *Convex combination:* Compute $\hat{\alpha} = \lambda_1\check{\alpha} + (1 - \lambda_1)\check{\alpha}$.

3: *Labeling:* $\text{label}(\mathbf{y}) = \arg \max_i \{q_i\}$, where q_i denotes the i^{th} coefficient of $\mathbf{q} = \mathbf{B}\hat{\alpha}$.

Output: $\text{label}(\mathbf{y})$.

4. Experiments

The proposed approach has been evaluate on three multiview RGB-D datasets: UWA 3D Multiview Activity II Dataset [45, 34], Northwestern-UCLA Multiview Action 3D Dataset [46] and the NTU RGB+D Human Activity Dataset [47]. We compare the performance of the proposed approach with the existing state-of-the-art methods. The results of the existing methods are taken directly from the original papers where applicable, otherwise, these results are taken from the best reported results in the literature. For the proposed approach, we used $\lambda = 0.01$, $\lambda_1 = 0.35$ and the sparsity threshold $k = 50$. These parameters were optimized empirically using cross-validation.

4.1. UWA3D-II Dataset

This dataset is composed of 30 human actions that have been performed by 10 subjects, where each action is recorded from 4 different viewpoints. These actions include: (1) one hand waving, (2) one hand punching, (3) two hands waving, (4) two hands punching, (5) sitting down, (6) standing up, (7) vibrating, (8) falling down, (9) holding chest, (10) holding head, (11) holding back,




Figure 5: RGB and Depth pairs from the UWA 3D Multiview Activity II Dataset [34].

(12) walking, (13) irregular walking, (14) lying down, (15) turning around, (16) drinking, (17) phone answering, (18) bending, (19) jumping jack, (20) running, (21) picking up, (22) putting down, (23) kicking, (24) jumping, (25) dancing, (26) moping floor, (27) sneezing, (28) sitting down (chair), (29) squatting, and (30) coughing. The four viewpoints are: (a) front, (b) left, (c) right, (d) top. The dataset is challenging because its action classes are reasonably large and they additionally containing subject scale and viewpoint variations. Furthermore, human-object interaction and self-occlusion in few videos makes action recognition on this dataset even more challenging. Figure 5 shows a representative example of RGB and Depth pair for one of the actions from different viewpoints.

For the evaluation, we follow Wang et al. [46] and use videos from two views for training and the remaining views for testing, creating 12 different view combinations. We first compare the results of the proposed approach employing three variants of the non-linear knowledge transfer model network architectures (see Fig 3). Note that, the original knowledge transfer model [6] only deals with the RGB frames. We implement the network variants using the Caffe library and train our models from the scratch. Table 1 summarizes the results of this comparison. Based on these results we chose Network #1 for our approach. Henceforth, unless otherwise mentioned, the results of the proposed approach are based on this architecture.

Table 1: Action recognition accuracies (%) on UWA 3D-II dataset using variants of NKTM architectures shown in Fig 3. The architectures are implemented using Caffe library and trained from the scratch.

Training view	V1 & V2		V1 & V3		V1 & V4		V2 & V3		V2 & V4		V3 & V4		Mean
Testing view	V3	V4	V2	V4	V2	V3	V1	V4	V1	V3	V1	V2	
Input: RGB + Depth images													
NKTM #1	86.9	89.8	81.9	89.5	76.7	83.6	83.6	79.0	89.6	82.1	89.2	83.8	84.6
NKTM #2	85.1	88.8	82.0	87.6	76.3	82.1	83.3	78.3	90.3	81.0	89.2	79.3	83.6
NKTM #3	85.1	88.4	80.8	89.9	74.4	82.1	83.6	76.4	90.0	81.3	90.0	82.3	83.7




Figure 6: Comparison of the proposed method with the best performing RGB-only and Depth-only methods on the UWA 3D-II dataset. V12-3 indicates that view 1 and 2 are used for training while view 3 is used for testing.

Table 2, summarizes the quantitative comparison of the proposed approach with the existing methods. In the table, the proposed method achieves 84.6% average recognition accuracy, which is 7.7% higher than its nearest competitor HPM+TM [5]. In Fig. 6 we compare our approach to the the best performing RGB-only method [6] and the best performing Depth-only method [5]. Note that, for all of the train-test view combinations, the proposed method provides a significant improvement in the recognition accuracy. The maximum reduction in the error rate, i.e. 36.4%, is achieved when view 2 and 3 are used for training and view 4 is used for testing. Based on these results, we can argue that our method effectively integrates the advantages of RGB and Depth video streams to enhance the recognition accuracy, especially for the large viewpoint variations.

Table 2: Comparison of action recognition accuracy (%) on the UWA3D-II dataset. Each time two views are used for training and the remaining two views are individually used for testing.

Training views	V1 & V2		V1 & V3		V1 & V4		V2 & V3		V2 & V4		V3 & V4		Mean
Testing view	V3	V4	V2	V4	V2	V3	V1	V4	V1	V3	V1	V2	
Input: RGB images													
AOG [46]	47.3	39.7	43.0	30.5	35.0	42.2	50.7	28.6	51.0	43.2	51.6	44.2	42.3
Action Tube [48]	49.1	18.2	39.6	17.8	35.1	39.0	52.0	15.2	47.2	44.6	49.1	36.9	37.0
LRCN [49]	53.9	20.6	43.6	18.6	37.2	43.6	56.0	20.0	50.5	44.8	53.3	41.6	40.3
NKTM [6]	60.1	61.3	57.1	65.1	61.6	66.8	70.6	59.5	73.2	59.3	72.5	54.5	63.5
Input: Depth images													
DVV [26]	35.4	33.1	30.3	40.0	31.7	30.9	30.0	36.2	31.1	32.5	40.6	32.0	33.7
CVP [28]	36.0	34.7	35.0	43.5	33.9	35.2	40.4	36.3	36.3	38.0	40.6	37.7	37.3
HON4D [50]	31.1	23.0	21.9	10.0	36.6	32.6	47.0	22.7	36.6	16.5	41.4	26.8	28.9
SNV [35]	31.9	25.7	23.0	13.1	38.4	34.0	43.3	24.2	36.9	20.3	38.6	29.0	29.9
HOPC [34]	52.7	51.8	59.0	57.5	42.8	44.2	58.1	38.4	63.2	43.8	66.3	48.0	52.2
HPM+TM [5]	80.6	80.5	75.2	82.0	65.4	72.0	77.3	67.0	83.6	81.0	83.6	74.1	76.9
Input: RGB + Depth images													
Ours (RGB-D)	86.9	89.8	81.9	89.5	76.7	83.6	83.6	79.0	89.6	82.1	89.2	83.8	84.6

4.2. Northwestern-UCLA Dataset

The Northwestern-UCLA Multiview Action 3D Dataset [46] contains RGB-D videos captured simultaneously by 3 Kinect cameras from 3 different viewpoints. There are 10 action classes: (1) pick up with one hand, (2) pick up with two hands, (3) drop trash, (4) walk around, (5) sit down, (6) stand up, (7) donning, (8) doffing, (9) throw, and (10) carry. The three viewpoints available in the dataset are: (a) left, (b) front, and (c) right. Every action in this dataset is performed by 10 subjects. In Fig. 7, we show representative sample RGB and Depth image pairs from the three viewpoints. This dataset is challenging for two reasons. Firstly, many action categories share the same “walking” pattern before and after the actual occurrence of the action of interest. Secondly, some actions such as “pick up with one hand” and “pick up with two hands” are very similar, which makes them to distinguished in the presence of the viewpoint




Figure 7: RGB and Depth pairs of a single action from three different views in the Northwestern-UCLA Multiview Action 3D Dataset [46].

Table 3: Comparison of action recognition accuracy (%) on the Northwestern-UCLA dataset using different NKTM networks from Figure 3.

Training view	V1&V2	V1&V3	V2&V3	Mean
Testing view	V3	V2	V1	
Input: RGB + Depth images				
NKTM #1	92.9	82.8	72.5	82.7
NKTM #2	92.5	82.4	71.9	82.3
NKTM #3	92.9	82.6	71.3	82.3

variations.

For this dataset, we also use videos captured from two different views for training and the third view for testing. Again, we first provide the relative performance comparison of the proposed approach using the three network architectures shown in Fig 3 in Table 3. We can see that NKTM #1 is also the most suited architecture for this dataset. In Table 4, we summarize the performance comparison of the proposed method with the existing methods. The proposed RGB-D method achieves an average accuracy of 82.7%, which is 3.0% higher than the nearest competitor HPM+TM.

Table 4: Comparison of action recognition accuracy (%) on the Northwestern-UCLA Multi-view Action 3D Dataset.

Training view	V1&V2	V1&V3	V2&V3	Mean
Testing view	V3	V2	V1	
Input: RGB images				
Hankelets [19]	45.2	-	-	-
DVV [26]	58.5	55.2	39.3	51.0
CVP [28]	60.6	55.8	39.5	52.0
AOG [46]	73.3	-	-	-
nCTE [1]	68.6	68.3	52.1	63.0
NKTM [6]	75.8	73.3	59.1	69.4
Input: Depth images				
DVV [26]	52.1	-	-	-
CVP [28]	53.5	-	-	-
HON4D [50]	39.9	-	-	-
SNV [35]	42.8	-	-	-
HOPC [34]	80.0	-	-	-
HPM+TM [5]	92.2	78.5	68.5	79.7
Input: RGB + Depth images				
Ours (RGB-D)	92.9	82.8	72.5	82.7

4.3. NTU RGB+D Human Activity Dataset

The NTU RGB+D Human Activity Dataset [47] is a large-scale RGB+D dataset for human activity analysis. It was collected using the Kinect v2 sensor and it includes 56,880 action samples each for the RGB videos, depth videos, skeleton sequences and the infra-red videos. In our experiments, we only make use of the RGB and the depth videos from the dataset. In these videos, there are 40 human subjects performing 60 types of actions including 50 single person actions and 10 two-person interactions. Three sensors were used to capture the data simultaneously from three horizontal angles: $-45^\circ, 0^\circ, 45^\circ$, and every action performer performed the action twice, facing the left or the right sensor. Moreover, the height of the sensors and their distances to the action performers were also adjusted to get further variations in the viewpoints. The NTU RGB+D dataset is one of the largest and the most complex cross-view action




Figure 8: RGB and depth samples from the NTU RGB+D Human Activity Dataset [47]. Three sensors C-1, C-2 and C-3 are used to record the data. The left group of images shows the actions recorded with the performer facing the sensor C-3, and the right group of images are recorded when the action performer faces the sensor C-2.

dataset of its kind to date. Figure 8 illustrates the RGB and the depth sample frames in the NTU RGB+D dataset.

We follow the standard evaluation protocol [47], which includes cross-subject and cross-view evaluations. For the cross-subject protocol, 40 subjects are partitioned into the training and the testing groups, where each group consists of 20 subjects. For the cross-view protocol, the videos captured by the sensor C-2 and C-3 are used as the training samples, and the videos captured by the sensor C-1 are used as the testing samples.

Table 5 summarizes the comparison of our method with the existing approaches on the NTU dataset. The second column of the table indicates the data type used by the approaches from this dataset. As can be seen, most of the existing approaches reporting result on this challenging dataset mainly exploit skeletal information. Using only the RGB frames or the Depth frames generally does not result in reasonable performance on this dataset. Similar to our approach, the DSSCA-SSLM [51] also uses both RGB and Depth information in their approach. However, the proposed method is able to achieve a significant

improvement over the performance of that DSSCA-SSLM.

5. Conclusion

We proposed an RGB-D human action recognition method that capitalizes on the view invariant characteristics of both Depth and RGB data streams, thereby making action recognition largely insensitive to the viewpoint variations in the videos. The proposed method processes the RGB and Depth streams separately to fully exploit the individual modalities. We extract dense action trajectories using the RGB frames to encode motion information, and then pass them through a deep network to get the viewpoint invariant features. For the Depth frames, we exploit the human pose model [5] to extract the appearance information. The CNN fc_7 layer viewpoint invariant features are encoded by Fourier Temporal Pyramid to incorporate the temporal dimension. Spatio-temporal features from both RGB and Depth streams are normalized and combined to form a set of heterogeneous features that are used to collaboratively represent the test features. A convex combination of dense and sparse collaborative representations is eventually used to predict the label of the test feature. Experiments on three standard multi-view RGB-D dataset and comparison to twelve existing methods demonstrate the effectiveness of the proposed approach.

Acknowledgement

This research was supported by ARC grant DP160101458.

Table 5: Action recognition accuracy (%) on the NTU RGB+D Human Activity Dataset.

Method	Data type	Cross Subject	Cross View
Baseline			
HON4D [50]	Depth	30.6	7.3
SNV [52]	Depth	31.8	13.6
HOG-2 [53]	Depth	32.4	22.3
Skeletal Quads [54]	Joints	38.6	41.4
Lie Group [55]	Joints	50.1	52.8
Deep RNN [47]	Joints	56.3	64.1
HBRNN-L [56]	Joints	59.1	64.0
Dynamic Skeletons [57]	Joints	60.2	65.2
Deep LSTM [47]	Joints	60.7	67.3
LieNet [58]	Joints	61.4	67.0
P-LSTM [47]	Joints	62.9	70.3
LTMD [59]	Depth	66.2	-
ST-LSTM [60]	Joints	69.2	77.7
Two-stream RNN [61]	Joints	71.3	79.5
Res-TCN [62]	Joints	74.3	83.1
DSSCA-SSLM [51]	RGB-D	74.9	-
Proposed			
Ours (RGB-D)	RGB-D	77.5	84.5

References

- [1] A. Gupta, J. Martinez, J. J. Little, R. J. Woodham, 3d pose from motion for cross-view action recognition via nonlinear circulant temporal encoding, in: CVPR, 2014.
- [2] H. Wang, A. Klser, C. Schmid, C. Liu, Action recognition by dense trajectories, in: CVPR, 2011.
- [3] H. Wang, A. Klser, C. Schmid, C. Liu, Dense trajectories and motion boundary descriptors for action recognition, in: IJCV, 2013.
- [4] H. Wang, C. Schmid, Action recognition with improved trajectories, in: ICCV, 2013.
- [5] H. Rahmani, A. Mian, 3d action recognition from novel viewpoints, in: CVPR, 2016.
- [6] H. Rahmani, A. Mian, Learning a non-linear knowledge transfer model for cross-view action recognition, in: CVPR, 2015.
- [7] A. Chaaoui, J. Padilla-Lopez, F. Flórez-Revuelta, Fusion of skeletal and silhouette-based features for human action recognition with rgb-d devices, in: ICCVW, 2013.
- [8] J.-F. Hu, W.-S. Zheng, J. Lai, J. Zhang, Jointly learning heterogeneous features for rgb-d activity recognition, in: CVPR, 2015.
- [9] L. Liu, L. Shao, Learning discriminative representations from rgb-d video data, in: IJCAI, 2013.
- [10] A. Shahroudy, G. Wang, T.-T. Ng, Multi-modal feature fusion for action recognition in rgb-d sequences, in: ISCCSP, 2014.
- [11] H. S. Koppula, A. Saxena, Learning spatio-temporal structure from rgb-d videos for human activity detection and anticipation, in: ICML, 2013.

- [12] P. Wei, Y. Zhao, N. Zheng, S.-C. Zhu, Modeling 4d human-object interactions for event and object recognition, in: ICCV, 2013.
- [13] J. Liu, N. Akhtar, A. Mian, View point invariant rgb-d human action recognition, in: International Conference on Digital Image Computing: Techniques and Applications (DICTA), 2017.
- [14] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell, Caffe: Convolutional architecture for fast feature embedding, arXiv preprint arXiv:1408.5093.
- [15] L. Zhang, M. Yang, X. Feng, Sparse representation or collaborative representation: Which helps face recognition?, in: International Conference on Computer Vision, 2011.
- [16] F. Lv, R. Nevatia, Single view human action recognition using key pose matching and viterbi path searching, in: CVPR, 2007.
- [17] T. Syeda-Mahmood, A. Vasilescu, S. Sethi, Action recognition from arbitrary views using 3d exemplars, in: ICCV, 2007.
- [18] A. Yilmaz, M. Shah, Action sketch: a novel action representation, in: CVPR, 2005.
- [19] B. Li, O. Camps, M. Sznajder, Cross-view activity recognition using skeletons, in: CVPR, 2012.
- [20] V. Parameswaran, R. Chellappa, View invariance for human action recognition, in: IJCV, 2006.
- [21] C. Rao, A. Yilmaz, M. Shah, View-invariant representation and recognition of actions, in: IJCV, 2002.
- [22] D. Weinland, R. Ronfard, E. Boyer, Free viewpoint action recognition using motion history volumes, in: CVIU, 2006.

- [23] A. Farhadi, M. K. Tabrizi, Learning to recognize activities from the wrong view point, in: ECCV, 2008.
- [24] A. Farhadi, M. K. Tabrizi, I. Endres, D. A. Forsyth, A latent model of discriminative aspect, in: ICCV, 2009.
- [25] R. Gopalan, R. Li, R. Chellapa, Domain adaption for object recognition: An unsupervised approach, in: ICCV, 2011.
- [26] R. Li, T. Zickler, Discriminative virtual views for crossview action recognition, in: CVPR, 2012.
- [27] J. Liu, M. Shah, B. Kuipersy, S. Savarese, Cross-view action recognition via view knowledge transfer, in: CVPR, 2011.
- [28] Z. Zhang, C. Wang, B. Xiao, W. Zhou, S. Liu, C. Shi, Cross-view action recognition via a continuous virtual path, in: CVPR, 2013.
- [29] J. Zheng, Z. Jiang, Learning view-invariant sparse representations for cross-view action recognition, in: ICCV, 2013.
- [30] L. Wang, Y. Qiao, X. Tang, Action recognition with trajectory-pooled deep-convolutional descriptors, in: CVPR, 2015.
- [31] W. Li, Z. Zhang, Z. Liu, Action recognition based on a bag of 3d points, in: CVPRW, 2010.
- [32] F. Lv, R. Nevatia, Recognition and segmentation of 3-d human action using hmm and multi-class adaboost, in: ECCV, 2006.
- [33] C. Lu, J. Jia, C.-K. Tang, Range-sample depth feature for action recognition, in: CVPR, 2014.
- [34] H. Rahmani, A. Mahmood, D. Q. Huynh, A. Mian, Hopc: Histogram of oriented principal components of 3d pointclouds for action recognition, in: ECCV, 2014.

- [35] X. Yang, Y. Tian, Super normal vector for activity recognition using depth sequences, in: CVPR, 2014.
- [36] Y. Kong, Y. Fu, Bilinear heterogeneous information machine for rgb-d action recognition, in: CVPR, 2015.
- [37] Pretrained human pose model, <http://staffhome.ecm.uwa.edu.au/~00053650/code.html>.
- [38] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: NIPS, 2012.
- [39] Cmu motion capture database, <http://mocap.cs.cmu.edu/>.
- [40] J. Wang, Z. Liu, Y. Wu, J. Yuan, Learning actionlet ensemble for 3d human action recognition, in: PAMI, 2013.
- [41] N. Akhtar, F. Shafait, A. Mian, Efficient classification with sparsity augmented collaborative representation, *Pattern Recognition* 65 (2017) 136 – 145.
- [42] N. Akhtar, F. Shafait, A. Mian, Discriminative bayesian dictionary learning for classification, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38 (12) (2016) 2374–2388.
- [43] N. Akhtar, A. Mian, F. Porikli, Joint discriminative bayesian dictionary and classifier learning, in: *IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [44] J. A. Tropp, A. C. Gilbert, Signal recovery from random measurements via orthogonal matching pursuit, *IEEE Transactions on Information Technology* 53 (12) (2007) 4655–4666.
- [45] H. Rahmani, A. Mahmood, D. Huynh, A. Mian, Histogram of oriented principal components for cross-view action recognition, in: PAMI, 2016.

- [46] J. Wang, X. N. nad Y. Xia, Y. Wu, S. Zhu, Cross-view action modeling, learning and recognition, in: CVPR, 2014.
- [47] A. Shahroudy, J. Liu, T.-T. Ng, G. Wang, Ntu rgb+d: A large scale dataset for 3d human activity analysis, in: IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 1010–1019.
- [48] G. Gkioxari, J. Malik, Finding action tubes, in: CVPR, 2015.
- [49] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, in: CVPR, 2015.
- [50] O. Oreifej, Z. Liu, Hon4d: histogram of oriented 4d normals for activity recognition from depth sequences, in: CVPR, 2013.
- [51] A. Shahroudy, T.-T. Ng, Y. Gong, G. Wang, Deep multimodal feature analysis for action recognition in rgb+ d videos, IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [52] X. Yang, Y. Tian, Super normal vector for activity recognition using depth sequences, in: IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 804–811.
- [53] E. Ohn-Bar, M. Trivedi, Joint angles similarities and hog2 for action recognition, in: IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2013, pp. 465–470.
- [54] G. Evangelidis, G. Singh, R. Horaud, Skeletal quads: Human action recognition using joint quadruples, in: International Conference on Pattern Recognition, 2014, pp. 4513–4518.
- [55] R. Vemulapalli, F. Arrate, R. Chellappa, Human action recognition by representing 3d skeletons as points in a lie group, in: IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 588–595.

- [56] Y. Du, W. Wang, L. Wang, Hierarchical recurrent neural network for skeleton based action recognition, in: IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 1110–1118.
- [57] J.-F. Hu, W.-S. Zheng, J. Lai, J. Zhang, Jointly learning heterogeneous features for rgb-d activity recognition, in: IEEE conference on Computer Vision and Pattern Recognition, 2015, pp. 5344–5352.
- [58] Z. Huang, C. Wan, T. Probst, L. Van Gool, Deep learning on lie groups for skeleton-based action recognition, in: IEEE Conference on Computer Vision and Pattern Recognition, 2016.
- [59] Z. Luo, B. Peng, D.-A. Huang, A. Alahi, L. Fei-Fei, Unsupervised learning of long-term motion dynamics for videos, in: IEEE Conference on Computer Vision and Pattern Recognition, 2017.
- [60] J. Liu, A. Shahroudy, D. Xu, G. Wang, Spatio-temporal lstm with trust gates for 3d human action recognition, in: European Conference on Computer Vision, 2016, pp. 816–833.
- [61] H. Wang, L. Wang, Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks, in: IEEE Conference on Computer Vision and Pattern Recognition, 2017.
- [62] T. S. Kim, A. Reiter, Interpretable 3d human action analysis with temporal convolutional networks, in: IEEE Conference on Computer Vision and Pattern Recognition Workshop, 2017.