

Memory Attention Networks for Skeleton-based Action Recognition

Chunyu Xie^{1,†}, Ce Li^{2,†}, Baochang Zhang^{1,*}, Chen Chen³, Jungong Han⁴, Jianzhuang Liu⁵

¹ School of Automation Science and Electrical Engineering, Beihang University, Beijing, China

² China University of Mining & Technology, Beijing, China

³ Department of Electrical and Computer Engineering, University of North Carolina at Charlotte

⁴ School of Computing & Communications, Lancaster University, LA1 4YW, UK

⁵ Noah's Ark Lab, Huawei Technologies Co. Ltd., China

yuxie@buaa.edu.cn, celi@cumtb.edu.cn, bczhang@buaa.edu.cn, chenchen870713@gmail.com,

jungonghan77@gmail.com, liu.jianzhuang@huawei.com

Abstract

Skeleton-based action recognition task is entangled with complex spatio-temporal variations of skeleton joints, and remains challenging for Recurrent Neural Networks (RNNs). In this work, we propose a temporal-then-spatial recalibration scheme to alleviate such complex variations, resulting in an end-to-end Memory Attention Networks (MANs) which consist of a Temporal Attention Recalibration Module (TARM) and a Spatio-Temporal Convolution Module (STCM). Specifically, the TARM is deployed in a residual learning module that employs a novel attention learning network to recalibrate the temporal attention of frames in a skeleton sequence. The STCM treats the attention calibrated skeleton joint sequences as images and leverages the Convolution Neural Networks (CNNs) to further model the spatial and temporal information of skeleton data. These two modules (TARM and STCM) seamlessly form a single network architecture that can be trained in an end-to-end fashion. MANs significantly boost the performance of skeleton-based action recognition and achieve the best results on four challenging benchmark datasets: NTU RGB+D, HDM05, SYSU-3D and UT-Kinect.¹

1 Introduction

3D skeleton-based human action recognition has recently attracted a lot of research interests due to its high-level representation and robustness to variations of viewpoints, appearances and surrounding distractions [Han *et al.*, 2017; Presti and La Cascia, 2016; Ding and Fan, 2016]. It is motivated by the biological observations that human beings can

recognize actions from just the motion of a few joints of the human body, even without appearance information [Johansson, 1973]. To describe human actions, conventional recognition approaches use relative joint coordinates to overlook the absolute movements of skeleton joints and thus gain partial view-invariant transformation. They include aligned spherical coordinates with person's direction [Xia *et al.*, 2012], translated coordinates invariant to absolute position and orientation [Jiang *et al.*, 2015], and flexible view invariant transform with principal components [Raptis *et al.*, 2011].

Skeleton sequences are time series of joint coordinate positions. To learn the temporal context of sequences, Recurrent Neural Networks (RNNs) [Li *et al.*, 2017], Long Short-Term Memory (LSTM) [Zhu *et al.*, 2016], and Gated Recurrent Unit (GRU) [Cho *et al.*, 2014], have been successfully applied to skeleton based action recognition. But it still challenging to cope with the complex spatio-temporal variations of skeleton joints caused by a number of factors, such as action speed, jitters, and surrounding distractions. To handle these variations, attention mechanism is introduced in [Liu *et al.*, 2016; Liu *et al.*, 2017a; Zhu *et al.*, 2016; Song *et al.*, 2017] to provide a robust recognition system. For instance, STA-LSTM [Song *et al.*, 2017] allocates different attention weights for selecting key frames and discriminative joints within one frame. Similarly, GCA-LSTM [Liu *et al.*, 2017a] selects the global informative joints from a sequence.

A few works exploit CNNs to solve the skeleton based action recognition problem. In [Ke *et al.*, 2017], skeleton joints after being projected or encoded, are used as the input channels of CNNs, which causes temporal information loss during the conversion of 3D information (x, y, z joint coordinates) into 2D information (images). In [Liu *et al.*, 2017b; Lea *et al.*, 2016], skeleton joints in each frame are transformed and expressed as color heat maps in CNNs, where complex data preprocessing gives rise to the loss of distinct spatio-temporal information.

In this work, our goal is to bring these powerful tools (*e.g.* RNNs, CNNs and attention learning) under the same umbrella and develop an efficient framework to investigate a new

¹The code will be made publicly available at <https://github.com/memory-attention-networks>.

* Baochang Zhang is the corresponding author.

† Chunyu Xie and Ce Li have equal contribution to the paper.

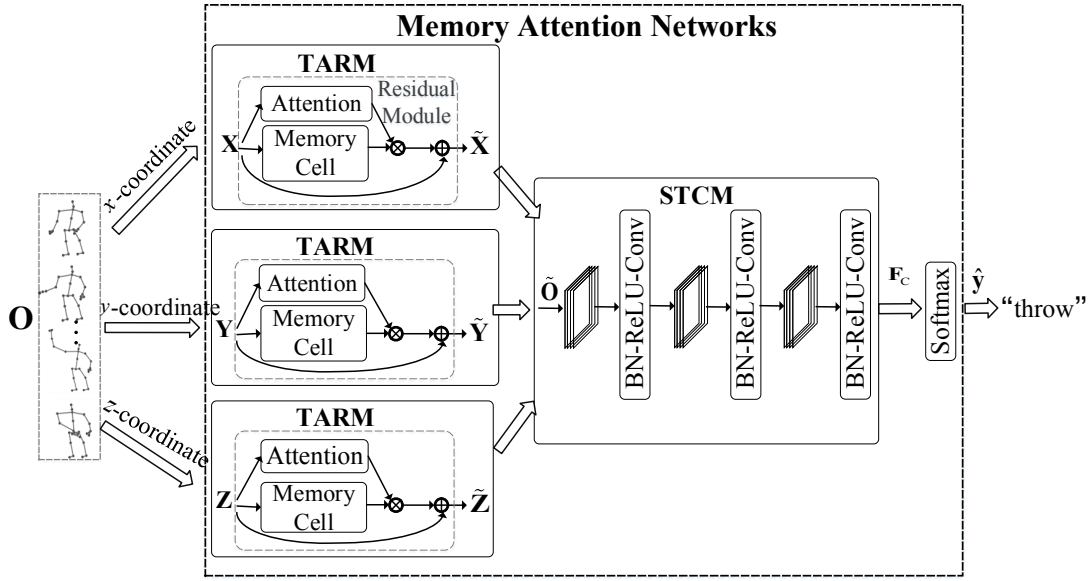


Figure 1: Memory Attention Networks use the temporal-then-spatial recalibration scheme. The (TARM) is deployed in the Residual Module (RM) to take advantage of the input features and learned attention information. The (STCM), which treats the skeleton sequences as images and leverages the CNNs, further models the spatial and temporal information of skeleton data to cope with complex spatio-temporal variations in skeleton joints.

Variable	Description
O	input of MANs
\tilde{O}	output of three TARMs
X	x -coordinate input of MANs
\tilde{X}	x -coordinate output feature map of a TARM
F_M	memory information in a TARM
F_A	attention weight in a TARM
F_C	output feature map of STCM
\hat{y}	predicted action label

Table 1: A brief description of notations used in the paper.

hypothesis of “memory attention + convolution network” for skeleton based action recognition. We propose an end-to-end deep network architecture, termed as Memory Attention Networks (MANs), to perform temporal-then-spatial feature recalibration. It can leverage the state-of-the-art CNNs to enhance the spatio-temporal features [Zhang *et al.*, 2016]. So far, CNNs particularly ResNets [He *et al.*, 2016] have been the most popular tools due to the unique residual module. Inspired by it, we design our temporal-then-spatial recalibration scheme in MANs based on the residual module as shown in Fig. 1. By doing so, both the original input features and the attention information can be fully exploited by subsequent CNNs in a unified framework, leading to a comprehensive and effective feature representation.

Distinctions between this work and prior art. (1) The state-of-the-art attention network [Song *et al.*, 2017] uses two LSTMs to model spatial and temporal attentions for each skeleton frame based on the input (frames) at time steps t and $t-1$. We also use RNN in TARM, but only to model the memory information of skeleton sequences. We design a new attention network to learn the attention weights and then make

use of the learned temporal attention to recalibrate the original skeleton sequence in a residual module, which facilitates efficient learning of attention features. (2) The existing CNNs based skeleton action recognition methods [Ke *et al.*, 2017; Liu *et al.*, 2017b] involve complicated pre-processing. For example, [Ke *et al.*, 2017] is based on clip generation (skeleton segmentation) and color images transformation; [Liu *et al.*, 2017b] performs skeleton coordinate transform and generates images with visual enhancement as the input. However, our proposed MANs directly operate on the skeleton sequences without bells and whistles, enabling an end-to-end training of network.

Contributions. The contributions of this paper are threefold.

1. We propose an end-to-end framework of Memory Attention Networks (MANs) to demonstrate the powerful capacity of a new “memory attention + convolution network” scheme for modeling the complex spatio-temporal variations in skeleton joints. It is the first time that a “RNNs + CNNs” framework has been developed for skeleton-based action recognition.
2. A new attention learning method is presented based on the residual module. It recalibrates temporal features to pay more attention to informative skeleton frames.
3. MANs achieve the state-of-the-art results on four challenging datasets. We also perform extensive ablation study to show the effectiveness of each unit in MANs.

2 Memory Attention Networks

In this section, we elaborate the two modules: Temporal Attention Recalibration Module (TARM) and Spatio-Temporal Convolution Module (STCM) in MANs. Table 1 summarizes the notations used in this paper.




Figure 2: The Residual Module (**left**), Temporal Attention Recalibration Module (**middle**), and BiGRU Memory Cell of TARM (**right**). TARM is designed based on the residual module as $\tilde{\mathbf{X}} = \mathbf{X} + \mathbf{F}(\mathbf{X})$, which incorporates the input and recalibrated features in a unified framework. $\mathbf{F}(\mathbf{X})$ is the recalibrated feature, *e.g.*, the output of our residual attention module, via $\mathbf{F}(\mathbf{X}) = \mathbf{F}_M(\mathbf{X}) \odot \mathbf{F}_A(\mathbf{X})$.

2.1 Temporal Attention Recalibration Module

The input skeleton data is a sequence of multi-frame 3D joint coordinates forming an action. Let $\mathbf{O} = \{\mathbf{X}, \mathbf{Y}, \mathbf{Z}\} \in \mathbb{R}^{T \times N \times 3}$, where $\mathbf{X} \in \mathbb{R}^{T \times N}$, $\mathbf{Y} \in \mathbb{R}^{T \times N}$, $\mathbf{Z} \in \mathbb{R}^{T \times N}$, denotes N joints along T frames with x , y , and z coordinates. For ease of explanation, \mathbf{X} is chosen as an example to describe TARM.

As shown in Fig. 2, given a 2D matrix \mathbf{X} , the learning of TARM pursues a specific attention based on the BiGRU in memory cell to capture the temporal information across the input action sequence. More specifically, inspired by the original RM in ResNets, we construct TARM via identity mapping with transformation from $\mathbf{X} \in \mathbb{R}^{T \times N}$ to $\tilde{\mathbf{X}} \in \mathbb{R}^{T \times N}$ to capture the richer temporal information, as

$$\tilde{\mathbf{X}} = \mathbf{X} + \mathbf{F}(\mathbf{X}), \quad (1)$$

where $\mathbf{F}(\mathbf{X})$ is the recalibrated feature, *e.g.*, the output of our residual attention module shown in Fig. 2, based on two branches: $\mathbf{F}_M(\mathbf{X})$ and $\mathbf{F}_A(\mathbf{X})$, which represent the memory information and attention weight, respectively.

$$\mathbf{F}(\mathbf{X}) = \mathbf{F}_M(\mathbf{X}) \odot \mathbf{F}_A(\mathbf{X}), \quad (2)$$

where \odot denotes the element-wise multiplication. For simplicity, $\mathbf{F}_M(\mathbf{X})$ and $\mathbf{F}_A(\mathbf{X})$ are denoted as $\mathbf{F}_M \in \mathbb{R}^{T \times K}$ and $\mathbf{F}_A \in \mathbb{R}^{T \times K}$, respectively. \mathbf{F}_A is the weight of \mathbf{F}_M to recalibrate temporal information. Obviously, for an action sequence, the importance of representative information in each frame is different, and only a few key frames containing important discriminative information deserves to be emphasized for action representation.

Calculating \mathbf{F}_M . We implement the memory cell via BiGRU. $\mathbf{X} \in \mathbb{R}^{T \times N}$ is resized and updated by the output of

a FC layer as $\mathbf{X} \leftarrow FC(\mathbf{X}) \in \mathbb{R}^{T \times K}$, as shown at the top of Fig. 2. In a slight abuse of notation, we still denote the resized output as \mathbf{X} . $\mathbf{F}_M \in \mathbb{R}^{T \times K}$ is the memory information made up of two directional combined hidden states in BiGRU, where K denotes the number of neuron units in BiGRU.

For simplicity, we still denote $\mathbf{X} \in \mathbb{R}^{T \times K}$ as

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}(1) \\ \vdots \\ \mathbf{x}(t) \\ \vdots \\ \mathbf{x}(T) \end{bmatrix} = \begin{bmatrix} x_1(1) & \cdots & x_k(1) & \cdots & x_K(1) \\ \vdots & & \ddots & & \vdots \\ x_1(t) & & x_k(t) & & x_K(t) \\ \vdots & & \ddots & & \vdots \\ x_1(T) & \cdots & x_k(T) & \cdots & x_K(T) \end{bmatrix}, \quad (3)$$

where $\mathbf{x}(t)$ is a row vector of \mathbf{X} to represent the sequence at t^{th} frame as $[x_1(t), \dots, x_K(t)]$, and $t \in (1, \dots, T)$.

As illustrated in Fig. 2, $\overrightarrow{\text{GRU}}(\mathbf{x}(t))$ and $\overleftarrow{\text{GRU}}(\mathbf{x}(t))$ are the output hidden states of the forward GRU and backward GRU, respectively. Combining these bidirectional hidden states, the informative vector $\mathbf{m}(t) \in \mathbb{R}^{1 \times K}$ at the t^{th} frame is denoted as

$$\mathbf{m}(t) = \overrightarrow{\text{GRU}}(\mathbf{x}(t)) + \overleftarrow{\text{GRU}}(\mathbf{x}(t)). \quad (4)$$

Finally, all the outputs are concatenated across T frames, and the memory information of skeleton joints is represented as

$$\mathbf{F}_M = [\mathbf{m}(1), \dots, \mathbf{m}(T)]^T \in \mathbb{R}^{T \times K}, \quad (5)$$

where \mathbf{F}_M summarizes the memory information in BiGRU for the skeleton joints across the sequence.

Calculating \mathbf{F}_A . To recalibrate the memory information \mathbf{F}_M , the attention weight \mathbf{F}_A is exploited as shown in Eq. (2).

In Fig. 2, following $\mathbf{X} \in \mathbb{R}^{T \times K}$, our recalibration scheme can capture global frame-wise dependence across T frames. We first aggregate each row vector of \mathbf{X} in Eq. (3) by the average pooling operation to produce a $T \times 1$ vector as

$$\mathbf{X}_p = \left[\frac{1}{K} \sum_{k=1}^K x_k(1), \dots, \frac{1}{K} \sum_{k=1}^K x_k(T) \right]^T \in \mathbb{R}^{T \times 1}. \quad (6)$$

We then duplicate it with K copies as

$$\mathbf{X} \leftarrow \mathbf{F}_d(\mathbf{X}_p, K) \in \mathbb{R}^{T \times K}, \quad (7)$$

where $\mathbf{F}_d(\mathbf{X}_p, K) = \underbrace{[\mathbf{X}_p, \dots, \mathbf{X}_p]}_K$.

In the right branch of TARM shown in Fig. 2, the attention mechanism is represented by a bottleneck with the two FC layers providing the non-linear interaction between frames. We introduce a dimensionality-reduction layer with parameters \mathbf{W}_1 and a ratio factor α (empirically set to be 16), followed by a ReLU function. And then, we introduce a dimensionality-increasing layer with parameters \mathbf{W}_2 and a sigmoid activation function. The dimensionality-reduction and dimensionality-increasing processing can be considered as the denoising and excitation operations respectively, and thus enhance the feature discriminability. Finally, the output of the attention branch \mathbf{F}_A is calculated as

$$\mathbf{F}_A = \sigma(\mathbf{W}_2 \theta(\mathbf{W}_1 \mathbf{X})) \in \mathbb{R}^{T \times K}, \quad (8)$$

where $\mathbf{W}_1 \in \mathbb{R}^{\frac{T}{\alpha} \times T}$ and $\mathbf{W}_2 \in \mathbb{R}^{T \times \frac{T}{\alpha}}$. To simplify the notation, the bias terms in Eq. (8) are omitted. $\theta(\cdot)$ refers to the ReLU function, and $\sigma(\cdot)$ denotes the sigmoid function. Finally, $\mathbf{F}(\mathbf{X})$ is obtained by the element-wise multiplication of \mathbf{F}_M and \mathbf{F}_A .

Furthermore, to calculate the output feature map of TARM, $\tilde{\mathbf{X}}, \mathbf{F}(\mathbf{X}) \in \mathbb{R}^{T \times K}$ is resized by a FC layer as $\mathbf{F}(\mathbf{X}) \leftarrow FC(\mathbf{F}(\mathbf{X})) \in \mathbb{R}^{T \times N}$, shown at the bottom of Fig. 2. As a result, the final $\mathbf{F}(\mathbf{X})$ describes the temporal information of the entire skeleton sequence. Similar to RM in ResNets, \mathbf{F}_M and \mathbf{F}_A in TARM can be jointly learned during training. In a similar way, we can obtain $\tilde{\mathbf{Y}}$ and $\tilde{\mathbf{Z}}$ based on \mathbf{Y} and \mathbf{Z} in their corresponding TARM, and $\tilde{\mathbf{O}} = \{\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}, \tilde{\mathbf{Z}}\} \in \mathbb{R}^{T \times N \times 3}$.

2.2 Spatio-Temporal Convolution Module

Conventional attention methods in skeleton action recognition are limited by the modeling capacity of RNNs [Liu *et al.*, 2016; Liu *et al.*, 2017a]. STCM is introduced based on CNNs to extract the enhanced spatio-temporal features from the output ($\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}$ and $\tilde{\mathbf{Z}}$) of the TARM. By leveraging the robustness to deformation of CNNs, STCM further extracts high-level feature representations to better cope with spatio-temporal variations of skeleton joints.

In principal, any CNNs can be used in STCM, *e.g.*, DenseNets [Huang *et al.*, 2017]. $\tilde{\mathbf{O}} \in \mathbb{R}^{T \times N \times 3}$ denotes the output of TARMs and also the input to STCM, and \mathbf{F}_C denotes the output of STCM for the softmax classifier. For example, in Fig. 1, the BN-ReLU-Conv blocks in STCM are used to interpret the high-level spatial structures of skeleton

joints as

$$\mathbf{F}_C = \text{Conv} \left(\text{ReLU} \left(\text{BN} \left(\dots \text{Conv} \left(\text{ReLU} \left(\text{BN} \left(\tilde{\mathbf{O}} \right) \right) \right) \right) \right) \right). \quad (9)$$

Afterwards, \mathbf{F}_C is fed to a softmax classifier to predict the class label as

$$\hat{\mathbf{y}} = \text{softmax}(\mathbf{W}_C, \mathbf{F}_C), \quad (10)$$

where \mathbf{W}_C and $\hat{\mathbf{y}}$ denote the weights in the softmax layer and the predicted action label, respectively. The cross-entropy loss function [Goodfellow *et al.*, 2016] is adopted to measure the difference between the true class label \mathbf{y} and the prediction result $\hat{\mathbf{y}}$.

3 Experiments

The proposed MANs are evaluated on four public skeleton action datasets: NTU RGB+D [Shahroudy *et al.*, 2016], HDM05 [Müller *et al.*, 2005], SYSU-3D [Hu *et al.*, 2015] and UT-Kinect [Xia *et al.*, 2012].

3.1 Implementation Details

For all the datasets, the matrices ($\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}$ and $\tilde{\mathbf{Z}}$) are generated with all the frames of a skeleton sequence. We use two different scales for the three input matrices of each sequence, *i.e.*, 224×224 and 50×50 , respectively. For the large scale, the number of hidden units of BiGRU in TARM is set to 2×128 ($K = 128$), where 2 indicates bidirectional GRU, 128 is the number of neurons. DenseNet-161 [Huang *et al.*, 2017] is used in STCM, leading to MANs (DenseNet-161). For the small scale, we set the hidden units of BiGRU to 2×64 , and stack multiple BN-ReLU-Conv blocks as STCM, resulting in MANs- n (*e.g.* $n = 9$) where n is the number of BN-ReLU-Conv blocks. The architectures of various MANs are illustrated in Table 2.

MANs are trained using the stochastic gradient descent algorithm, and the learning rate, decay, and momentum, are respectively set to 0.1, 0, and 0.9. The mini-batches of samples on NTU RGB+D, HDM05, SYSU-3D, and UT-Kinect are constructed by randomly sampling 40, 20, 8, and 8 samples from the training sets, respectively. The training stops after 100 epochs except for NTU RGB+D after 50 epochs. For a fair comparison, the performance of MANs on each dataset is compared with existing methods using the same evaluation protocol. All experiments are performed based on Keras² with Tensorflow backend using two NVIDIA Titan X Pascal GPUs.

3.2 Experiment Analysis

Ablation study. We conduct extensive ablation study of different units in MANs with the following settings: (A) STCM³ (*i.e.* applying CNN on the original skeleton images); (B) MANs (no attention)–MANs without attention in TARM (*i.e.* no \mathbf{F}_A); (C) MANs (other temporal attention)–MANs use the temporal attention scheme in [Song *et al.*, 2017] for the

²<http://keras.io>

³Here STCM takes $[\mathbf{X}, \mathbf{X}, \mathbf{Z}]$ as image input to perform classification via CNNs.

Module	Output Size	MANs-9		MANs-33		MANs-61	
TARM	50×50	64×2					
STCM	25×25	$5 \times 5, 64, \text{stride } 2$					
	25×25	$3 \times 3, 64$ $3 \times 3, 64$	$\times 2$	$3 \times 3, 64$ $3 \times 3, 64$	$\times 8$	$3 \times 3, 64$ $3 \times 3, 64$	$\times 15$
	13×13	$3 \times 3, 128$ $3 \times 3, 128$	$\times 2$	$3 \times 3, 128$ $3 \times 3, 128$	$\times 8$	$3 \times 3, 128$ $3 \times 3, 128$	$\times 15$
	1×1	Average Pooling, FC, Softmax					

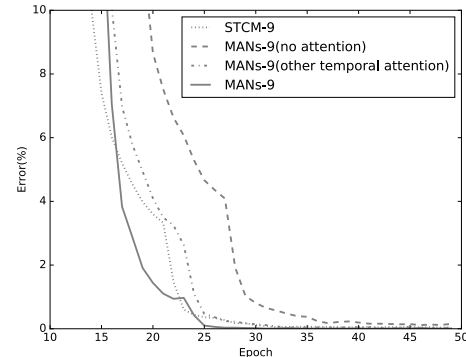
 Table 2: The architectures of various MANs- n (*i.e.*, $n = 9, 33, 61$).

Method	CS.	CV.
STCM-9 (CNNs)	81.31	89.78
MANs-9 (no attention)	81.41	89.84
MANs-9 (other temporal attention)	81.94	90.12
MANs-9	83.01	90.66
STCM (DenseNet-161)	81.56	90.24
MANs (DenseNet-161, no attention)	81.96	92.15
MANs (DenseNet-161, other temporal attention)	81.60	92.18
MANs (DenseNet-161)	82.67	93.22


Table 3: Performance of STCM, MANs (no attention), MANs (other temporal attention) and MANs on the NTU RGB+D dataset.

TARM. Table 3 shows the results of different architectures on the NTU RGB+D dataset. Note that the STCM of MANs uses DenseNet-161 for the input of 224×224 , and uses stacked BN-ReLU-Conv blocks with 9 layers for the input of 50×50 in Table 3, respectively. Comparing with our full MANs model, we have these findings. (1) Setting A yields much lower performance, indicating the importance of temporal information modeling for skeleton. (2) Setting B reveals the effectiveness of the proposed attention mechanism (*i.e.* learning \mathbf{F}_A). (3) Setting C substitutes the attention scheme in [Song *et al.*, 2017] for our residual attention module in MANs. The recognition accuracy is lower than that of our MANs, which again validates the superiority of our residual attention learning approach.

Learning convergence. We plot the training error and testing error curves of the four networks on a same trail of NTU RGB+D dataset in Fig. 3, including STCM-9, MANs-9 (no attention), MANs-9 (other temporal attention), and MANs-9. We can observe that MANs-9 in solid line converges at epoch #26 for training and obtains the best error rate of 16.99% for testing. For training error, STCM-9 stops decreasing at epoch #30, MANs-9 (no attention) converges at epoch #30, and MANs-9 (other temporal attention) converges at epoch #40. These curves show that MANs obviously converge faster and gain better performance than others, which proves that the “memory attention + convolution network” scheme in MANs can be used to improve the modeling ability of CNNs. MANs-9 converges much faster than MANs-9 (other temporal attention) (epoch #26 *v.s.* #40), due to the novel residual attention module which takes the input and attention information into account in the same framework. More specifically, the residual attention module not only uses



(a) Training error v.s. epoch



(b) Testing error v.s. epoch

Figure 3: Training and testing error curves of STCM-9, MANs-9 (no attention), MANs-9 (other temporal attention) and MANs-9 on the NTU RGB+D dataset (cross-subject setting).

the temporal attention recalibrated information, but also delivers the spatial structure information of the original input by the identity shortcut.

Various CNNs in STCM. The number of stacked layers of CNNs in STCM, *i.e.*, multiple BN-ReLU-Conv blocks used in MANs-9, MANs-33 and MANs-61, is evaluated in terms of the recognition accuracy on all the datasets (Table 4). We also include MANs (DenseNet-161) in the same table. The deeper CNNs have more learnable parameters than the shallower ones. We note that MANs-9 **with a similar parameter amount** as Deep LSTM [Zhu *et al.*, 2016] has much better performance than the state-of-the-arts. This reveals that our method is more effective if the network complexity should be considered. Interestingly, MANs-9 mostly achieve better per-

Method	#param	NTU RGB+D			HDM05	SYSU-3D	UT-Kinect
		Cross Subject	Cross View				
Hierarchical RNNs [Du <i>et al.</i> , 2015]	-	59.10	64.00	96.92	-	-	
Dynamic skeletons [Hu <i>et al.</i> , 2015]	-	60.23	65.22	-	75.50	-	
Deep LSTM [Zhu <i>et al.</i> , 2016]	0.6M	-	-	96.80	-	-	
ST-LSTM [Liu <i>et al.</i> , 2016]	-	61.70	75.50	-	76.50	97.00	
ST-LSTM + TG [Liu <i>et al.</i> , 2016]	-	69.20	77.70	-	76.80	97.50	
Two-stream RNNs [Wang and Wang, 2017]		71.30	79.50	-	-	-	
STA-LSTM [Song <i>et al.</i> , 2017]	0.5M	73.40	81.20	-	-	-	
Adaptive RNN-T [Li <i>et al.</i> , 2017]	-	74.60	83.20	-	-	-	
GCA-LSTM [Liu <i>et al.</i> , 2017a]	-	76.10	84.00	-	78.60	99.00	
Clips+CNN+MTLN [Ke <i>et al.</i> , 2017]	62M	79.57	84.83	-	-	-	
VA-LSTM [Zhang <i>et al.</i> , 2017]	-	79.40	87.60	-	77.50	-	
MANs-9	0.8M	83.01	90.66	98.46	87.04	100.0	
MANs-33	3.1M	82.40	90.94	98.85	86.81	100.0	
MANs-61	5.7M	82.42	90.97	98.76	87.63	99.50	
MANs (DenseNet-161)	27.6M	82.67	93.22	97.69	78.86	99.00	

Table 4: Comparison of the results of different units in MANs on four datasets.

formance than other deeper models, which is probably due to its compactness. However, the very deep MANs (DenseNet-161) obtain the best result on the challenging NTU RGB+D dataset in cross-view setting. It is worth noting that because of the “memory attention + convolution network” scheme of MANs, it is quite flexible to deploy different parameter amount in MANs by adjusting the number of CNN layers in STCM to balance between performance and network complexity.

3.3 Results and Comparisons

We show the performance comparison of various MANs architectures with other state-of-the-art approaches in Table 4 for the four datasets, respectively. Note that all our MANs under different parameter amounts are able to achieve better performance than the state-of-the-art RNNs-based approaches (*e.g.* VA-LSTM [Zhang *et al.*, 2017], GCA-LSTM [Liu *et al.*, 2017a]) and the state-of-the-art CNNs-based approaches (*e.g.* Clips+CNN+MTLN [Ke *et al.*, 2017]), demonstrating the superiority of our “memory attention + convolution network” architecture.

For the NTU RGB+D dataset, MANs perform significantly better than others in both the cross-subject and cross-view protocols. Comparing to other methods, MANs-9 increase the accuracy by 3.44% for cross-subject evaluation, and MANs (DenseNet-161) lead to a significant 5.62% improvement on this largest dataset in cross-view evaluation, which demonstrate that MANs can learn more discriminative spatio-temporal features to alleviate the spatial and temporal variations in skeleton joints.

For the HDM05 dataset, our MANs-33 achieves better result than the state-of-the-art multi-layer RNNs-based models. For the SYSU-3D dataset, four accuracies of our MANs are higher than all the RNNs-based methods [Hu *et al.*, 2015; Liu *et al.*, 2016; Liu *et al.*, 2017a; Zhang *et al.*, 2017]; es-

pecially MANs-61 outperform the previous best approach GCA-LSTM by **9.03%**. For the UT-Kinect dataset, MANs-9 and MANs-33 achieve a 100% accuracy, with 1.0% improvement in comparison with the state-of-the-art GCA-LSTM. The experiments show that compared with the existing RNNs-based methods, the deployed residual module with temporal attention and the temporal-then-spatial scheme can effectively improve the modeling ability of RNNs.

4 Conclusion

In this paper, we propose an end-to-end framework, termed Memory Attention Networks (MANs), to enhance the spatio-temporal features for skeleton-based action recognition. In MANs, TARM is designed to recalibrate the temporal attention to skeleton frames in action sequences, and STCM further models the spatial structure and temporal dependence of the skeleton sequence by leveraging the powerful CNNs. Through the unified framework, MANs significantly boost the performance for skeleton-based action recognition. The extensive experiments validate the superiority of MANs, which consistently perform the best on four benchmark datasets and contribute new state-of-the-art results.

Acknowledgements

The work was supported by the Natural Science Foundation of China under Contract 61672079, 61601466, and 61473086, the Open Projects Program of National Laboratory of Pattern Recognition, and Shenzhen Peacock Plan KQTD2016112515134654. Baochang Zhang is also with Shenzhen Academy of Aerospace Technology, Shenzhen, China.

References

- [Cho *et al.*, 2014] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Empirical Methods in Natural Language Processing*, pages 1724–1734, 2014. 1
- [Ding and Fan, 2016] Meng Ding and Guoliang Fan. Articulated and generalized gaussian kernel correlation for human pose estimation. *IEEE Transactions on Image Processing A Publication of the IEEE Signal Processing Society*, 25(2):776, 2016. 1
- [Du *et al.*, 2015] Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *CVPR*, pages 1110–1118, 2015. 6
- [Goodfellow *et al.*, 2016] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>. 4
- [Han *et al.*, 2017] Fei Han, Brain Reily, William Hoff, and Hao Zhang. Space-time representation of people based on 3D skeletal data: A review. *Computer Vision and Image Understanding*, 158:85–105, 2017. 1
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 2
- [Hu *et al.*, 2015] Jianfang Hu, Weishi Zheng, Jianhuang Lai, and Jianguo Zhang. Jointly learning heterogeneous features for RGB-D activity recognition. In *CVPR*, pages 5344–5352, 2015. 4, 6
- [Huang *et al.*, 2017] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 4
- [Jiang *et al.*, 2015] Min Jiang, Jun Kong, George Bebis, and Hongtao Huo. Informative joints based human action recognition using skeleton contexts. *Signal Processing: Image Communication*, 33:29–40, 2015. 1
- [Johansson, 1973] Gunnar Johansson. Visual perception of biological motion and a model for its analysis. *Perception and Psychophysic*, 14(2):201–211, 1973. 1
- [Ke *et al.*, 2017] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. A new representation of skeleton sequences for 3D action recognition. In *CVPR*, 2017. 1, 2, 6
- [Lea *et al.*, 2016] Colin Lea, Reiter Austin Vidal, René, and Gregory D. Hager. Temporal convolutional networks: A unified approach to action segmentation. *ECCV*, pages 47–54, 2016. 1
- [Li *et al.*, 2017] Wenbo Li, Longyin Wen, Ming-Ching Chang, Ser Nam Lim, and Siwei Lyu. Adaptive RNN tree for large-scale human action recognition. In *ICCV*, pages 1444–1452, 2017. 1, 6
- [Liu *et al.*, 2016] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal LSTM with trust gates for 3D human action recognition. In *ECCV*, pages 816–833, 2016. 1, 4, 6
- [Liu *et al.*, 2017a] Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, and Alex C. Kot. Global context-aware attention LSTM network for 3D action recognition. In *CVPR*, pages 1647–1656, 2017. 1, 4, 6
- [Liu *et al.*, 2017b] Mengyuan Liu, Hong Liu, and Chen Chen. Enhanced skeleton visualization for view invariant human action recognition. *PR*, 68:346–362, 2017. 1, 2
- [Müller *et al.*, 2005] M. Müller, T. Röder, and M. Clausen. Efficient content-based retrieval of motion capture data. *ACM Transactions on Graphic*, 24(3):677–685, 2005. 4
- [Presti and La Cascia, 2016] Liliana Lo Presti and Marco La Cascia. 3D skeleton-based human action classification. *PR*, 53:130–147, 2016. 1
- [Raptis *et al.*, 2011] Michalis Raptis, Darko Kirovski, and Hugues Hoppe. Real-time classification of dance gestures from skeleton animation. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 147–156. ACM, 2011. 1
- [Shahroudy *et al.*, 2016] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In *CVPR*, pages 1010–1019, 2016. 4
- [Song *et al.*, 2017] Sijie Song, Cuiling Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *AAAI*, pages 4263–4270, 2017. 1, 2, 4, 5, 6
- [Wang and Wang, 2017] Hongsong Wang and Liang Wang. Modeling temporal dynamics and spatial configuration of actions using two-stream recurrent neural networks. In *CVPR*, 2017. 6
- [Xia *et al.*, 2012] Lu Xia, Chiachih Chen, and Jake K Aggarwal. View invariant human action recognition using histograms of 3D joints. In *CVPRW*, pages 20–27, 2012. 1, 4
- [Zhang *et al.*, 2016] Bowen Zhang, Limin Wang, Zhe Wang, Yu Qiao, and HanLi Wang. Real-time action recognition with enhanced motion vector CNNs. In *CVPR*, pages 2718–2726, 2016. 2
- [Zhang *et al.*, 2017] Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *ICCV*, 2017. 6
- [Zhu *et al.*, 2016] Wentao Zhu, Cuiling Lan, Junliang Xing, Wenjun Zeng, Yanghao Li, Li Shen, and Xiaohui Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks. In *AAAI*, pages 3697–3704, 2016. 1, 5, 6