



HAL
open science

SkeleMotion: A New Representation of Skeleton Joint Sequences Based on Motion Information for 3D Action Recognition

Carlos Caetano, Jessica Sena, Francois F Bremond, Jefersson Alex dos Santos,
William Robson Schwartz

► **To cite this version:**

Carlos Caetano, Jessica Sena, Francois F Bremond, Jefersson Alex dos Santos, William Robson Schwartz. SkeleMotion: A New Representation of Skeleton Joint Sequences Based on Motion Information for 3D Action Recognition. AVSS 2019 - 16th IEEE International Conference on Advanced Video and Signal-based Surveillance, Sep 2019, Taipei, Taiwan. hal-02422551

HAL Id: hal-02422551

<https://hal.science/hal-02422551v1>

Submitted on 22 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SkeleMotion: A New Representation of Skeleton Joint Sequences Based on Motion Information for 3D Action Recognition

Carlos Caetano¹ Jessica Sena¹ François Brémond²

Jefersson A. dos Santos¹ William Robson Schwartz¹

¹Smart Sense Laboratory, Department of Computer Science
Universidade Federal de Minas Gerais, Belo Horizonte, Brazil

²INRIA, Sophia Antipolis, France

{carlos.caetano, jessicasena, jefersson, william}@dcc.ufmg.br
francois.bremond@inria.fr

Abstract

Due to the availability of large-scale skeleton datasets, 3D human action recognition has recently called the attention of computer vision community. Many works have focused on encoding skeleton data as skeleton image representations based on spatial structure of the skeleton joints, in which the temporal dynamics of the sequence is encoded as variations in columns and the spatial structure of each frame is represented as rows of a matrix. To further improve such representations, we introduce a novel skeleton image representation to be used as input of Convolutional Neural Networks (CNNs), named SkeleMotion. The proposed approach encodes the temporal dynamics by explicitly computing the magnitude and orientation values of the skeleton joints. Different temporal scales are employed to compute motion values to aggregate more temporal dynamics to the representation making it able to capture long-range joint interactions involved in actions as well as filtering noisy motion values. Experimental results demonstrate the effectiveness of the proposed representation on 3D action recognition outperforming the state-of-the-art on NTU RGB+D 120 dataset.

1. Introduction

Human action recognition plays an important role in various applications, for instance surveillance systems can be used to detect and prevent abnormal or suspicious actions, health care systems can be used to monitor elderly people on their daily living activities and robot and human interactions.

Over the last decade, significant progress on the action recognition task has been achieved with the design of discriminative representations employed to the image and

video domains on RGB data or optical flow. Such information is based on appearance or motion analysis. Due to the development of low-cost RGB-D sensors (e.g., Kinect), it becomes possible to employ depth information as well as human skeleton joints to perform 3D action recognition. Compared to RGB and optical flow, skeleton data has the advantages of being computationally efficient since the data size is smaller. Moreover, skeleton data are robust to illumination changes, robust to background noise and invariant to camera views [6].

Many works for 3D action recognition have focused on designing handcrafted feature descriptors [29, 33, 35, 5, 31, 3] to encode skeleton data while adopting Dynamic Time Warping (DTW), Fourier Temporal Pyramid (FTP) or Hidden Markov Model (HMM) to model temporal dynamics in the sequences. Nowadays, large efforts have been directed to the employment of deep neural networks. These architectures learn hierarchical layers of representations to perform pattern recognition and have demonstrated impressive results on many pattern recognition tasks (e.g., image classification [12] and face recognition [24]). For instance, Recurrent Neural Networks (RNNs) with Long-Short Term Memory (LSTM) have been employed to model skeleton data for 3D action recognition [28, 25, 27, 36]. Although RNN approaches present excellent results in 3D action recognition task due to their power of modeling temporal sequences, such structures lack the ability to efficiently learn the spatial relations between the skeleton joints [34].

To take advantage of the spatial relations, a hierarchical structure was proposed by Du et al. [4]. The authors represent each skeleton sequence as 2D arrays, in which the temporal dynamics of the sequence is encoded as variations in columns and the spatial structure of each frame is represented as rows. Then, the representation is fed to a Convolutional Neural Network (CNN) which has the natural abil-

ity of learning structural information from 2D arrays. Such type of representations are very compact encoding the entire video sequence in one single image.

To further improve the representation of skeleton joints for 3D action recognition, in this paper we introduce a novel skeleton image representation to be used as input of CNNs named SkeleMotion. The proposed approach encodes temporal dynamics by explicitly using motion information computing the magnitude and orientation values of the skeleton joints. To that end, different temporal scales are used to filter noisy motion values as well as aggregating more temporal dynamics to the representation making it being able to capture long-range joint interactions involved in actions. Moreover, the method takes advantage of a structural organization of joints that preserves spatial relations of more relevant joint pairs. To perform action classification, we train a tiny CNN architecture with only three convolutional layers and two fully-connected layers. Since the network is shallow and takes as input a compact representation for each video, it is extremely fast to train.

In the literature, many works employed or improved the skeleton image representation for 3D action recognition [32, 21, 10, 13, 30, 34, 14, 2]. However, none of the methods model explicit motion information (i.e., magnitude and orientation) in multiple temporal scales, as the proposed approach does. Similar to our approach, the works of [13, 14] were the only ones that tried to encode motion on skeleton images, however they employed a naive approach by computing difference of motion joints on consecutive frames.

According to the experimental results, our proposed skeleton image representation can handle skeleton based 3D action recognition very well. Moreover, SkeleMotion representation achieves the state-of-the-art performance on the large scale NTU RGB+D 120 [15] dataset when combined with a spatial structural joint representation.

The code of our SkeleMotion representation is publicly available to facilitate future research ¹.

2. Related Work

In this section, we present a literature review of works that are close to the idea proposed in our approach by employing different representations based on skeleton images.

As the forerunner of skeleton image representations, Du et al. [4] represent the skeleton sequences as a matrix. Each row of such matrix corresponds to a chain of concatenated skeleton joint coordinates from the frame t . Hence, each column of the matrix corresponds to the temporal evolution of the joint j . At this point, the matrix size is $J \times T \times 3$, where J is the number of joints for each skeleton, T is the total frame number of the video sequence and 3 is the num-

ber coordinate axes (x, y, z) . The values of this matrix are quantified into an image (i.e., linearly rescaled to a $[0, 255]$) and normalized to handle the variable-length problem. In this way, the temporal dynamics of the skeleton sequence is encoded as variations in rows and the spatial structure of each frame is represented as columns. Finally, the authors use their representation as input to a CNN model composed by four convolutional layers and three max-pooling layers. After the feature extraction, a feed-forward neural network with two fully-connected layers is employed for classification.

Wang et al. [32, 30] present a skeleton representation to represent both spatial configuration and dynamics of joint trajectories into three texture images through color encoding, named Joint Trajectory Maps (JTM). The authors apply rotations to the skeleton data to mimicking multi-views and also for data enlargement to overcome the drawback of CNNs usually being not view invariant. JTMs are generated by projecting the trajectories onto the three orthogonal planes. To encode motion direction in the JTM, they use a hue colormap function to “color” the joint trajectories over the action period. They also encode the motion magnitude of joints into saturation and brightness claiming that changes in motion results in texture in the JMTs. Finally, the authors individually fine-tune three AlexNet [12] CNNs (one for each JTM) to perform classification.

To overcome the problem of the sparse data generated by skeleton sequence video, Ke et al. [10] represent the temporal dynamics of the skeleton sequence by generating four skeleton representation images. Their approach is closer to Du et al. [4] method, however they compute the relative positions of the joints to four reference joints by arranging them as a chain and concatenating the joints of each body part to the reference joints resulting onto four different skeleton representations. According to the authors, such structure incorporate different spatial relationships between the joints. Finally, the skeleton images are resized and each channel of the four representations is used as input to a VGG19 [26] pre-trained architecture for feature extraction.

To encode motion information on skeleton image representation, Li et al. [13] proposed the skeleton motion image. Their approach is created similar to Du et al. [4] skeleton image representation, however each matrix cell is composed by joint difference computation between two consecutive frames. To perform classification, the authors used Du et al. [4] approach and their proposed representation independently as input of a neural network with a two-stream paradigm. The CNN used was a small seven-layer network consisting of three convolution layers and four fully-connected layers.

Yang et al. [34] claim that the concatenation process of chaining all joints with a fixed order turn into lack of semantic meaning and leads to loss in skeleton structural informa-

¹<https://github.com/carloscaetano/skeleton-images>

tion. To that end, Yang et al. [34] proposed a representation named Tree Structure Skeleton Image (TSSI) to preserve spatial relations. Their method is created by traversing a skeleton tree with a depth-first order algorithm with the premise that the fewer edges there are, the more relevant the joint pair is. The generated representation are then quantified into an image and resized before being sent to a ResNet-50 [7] CNN architecture.

As it can be inferred from the reviewed methods, most of them are improved versions of Du et al. [4] skeleton image focusing on spatial structural of joint axes while the temporal dynamics of the sequence is encoded as variations in columns, or encode motion information in a naive manner (difference of motion joints on consecutive frames). Despite the aforementioned methods produce promising results, they do not explicit encode rich motion information. In view of that, to capture more motion information, our approach directly encodes it by using orientation and magnitude to provide information regarding the velocity of the movement in different temporal scales. In view of that, our SkeleMotion approach differs from the literature methods by capturing the temporal dynamics explicitly provided by magnitude and orientation motion information.

3. Proposed Approach

In this section, we introduce our proposed skeleton image representation based on magnitude and orientation motion information, named *SkeleMotion*.

3.1. SkeleMotion

As reviewed in Section 2, the majority of works that encode skeleton data as image representations are based on spatial structure encoding of the skeleton joints. According to Li et al. [14], temporal movements of joints can also be used as crucial cues for action recognition and although the temporal dynamics of the sequence can be implicitly learned by using a CNN, an explicit modeling can produce better recognition accuracies.

Recently, a new temporal stream for the two-stream networks called Magnitude-Orientation Stream (MOS) [1] was developed. The method is based on non-linear transformations and claim that motion information on a video sequence can be described by the spatial relationship contained on the local neighborhood of magnitude and orientation extracted from the optical flow and has shown excellent results on the 2D action recognition problem. Motivated by such results, in this paper we propose a novel skeleton image representation (named SkeleMotion), based on magnitude and orientation of the joints to explore the temporal dynamics. Our approach expresses the displacement information by using orientation encoding (direction of joints) and magnitude to provide information regarding the velocity of the move-

ment. Furthermore, due to the successful results achieved by the skeleton image representations, our approach follows the same fundamentals by representing the skeleton sequences as a matrix. First, we apply the depth-first tree traversal order [34] to the skeleton joints to generate a pre-defined chain order C that best preserves the spatial relations between joints in original skeleton structures². Afterwards, we compute matrix S that corresponds to a chain of concatenated skeleton joint coordinates from the frame t . In view of that, each column of the matrix corresponds to the temporal evolution of the arranged chain joint c . At this point, the size of matrix S is $C \times T \times 3$, where C is the number of joints of the chain, T is the total frame number of the video sequence and 3 is the number joint coordinate axes (x, y, z) . Then, we create the motion structure \mathcal{D} as

$$\mathcal{D}_{c,t} = S_{c,t+d} - S_c, \quad (1)$$

where each matrix cell is composed by the temporal difference computation of each joint between two frames of d distance, resulting in a $C \times T - d \times 3$ matrix.

By using the proposed motion structure \mathcal{D} , we build two different representations: one based on the magnitudes of joint motions and another one based the orientations of the joint motion. We compute both representations using

$$M_{c,t} = \sqrt{(\mathcal{D}_{c,t}^x)^2 + (\mathcal{D}_{c,t}^y)^2 + (\mathcal{D}_{c,t}^z)^2} \quad (2)$$

and

$$\theta_{c,t} = \text{stack}(\theta_{c,t}^{xy}, \theta_{c,t}^{yz}, \theta_{c,t}^{zx})$$

$$\theta_{c,t}^{xy} = \tan^{-1} \left(\frac{\mathcal{D}_{c,t}^y}{\mathcal{D}_{c,t}^x} \right),$$

$$\theta_{c,t}^{yz} = \tan^{-1} \left(\frac{\mathcal{D}_{c,t}^z}{\mathcal{D}_{c,t}^y} \right), \quad (3)$$

$$\theta_{c,t}^{zx} = \tan^{-1} \left(\frac{\mathcal{D}_{c,t}^x}{\mathcal{D}_{c,t}^z} \right),$$

where M is the magnitude skeleton representation of size $J \times T - d \times 1$ and θ is the orientation skeleton representation of size $J \times T - d \times 3$ (composed by 3 stacked channels).

Since the orientation values are estimated for every joint, it might generate noisy values for joints without any movement. Therefore, we perform a filtering on θ based on the values of M as

$$\theta'_{c,t} = \begin{cases} 0, & \text{if } M_{c,t} < m \\ \theta_{c,t}, & \text{otherwise} \end{cases}, \quad (4)$$

where m is a magnitude threshold value.

Finally, the generated matrices are normalized into $[0, 1]$ and empirically resized into a fixed size of $C \times 100$, since

²Chain C considering 25 Kinect joints: [2, 21, 3, 4, 3, 21, 5, 6, 7, 8, 22, 23, 22, 8, 7, 6, 5, 21, 9, 10, 11, 12, 24, 25, 24, 12, 11, 10, 9, 21, 2, 1, 13, 14, 15, 16, 15, 14, 13, 1, 17, 18, 19, 20, 19, 18, 17, 1, 2], as defined in [34].

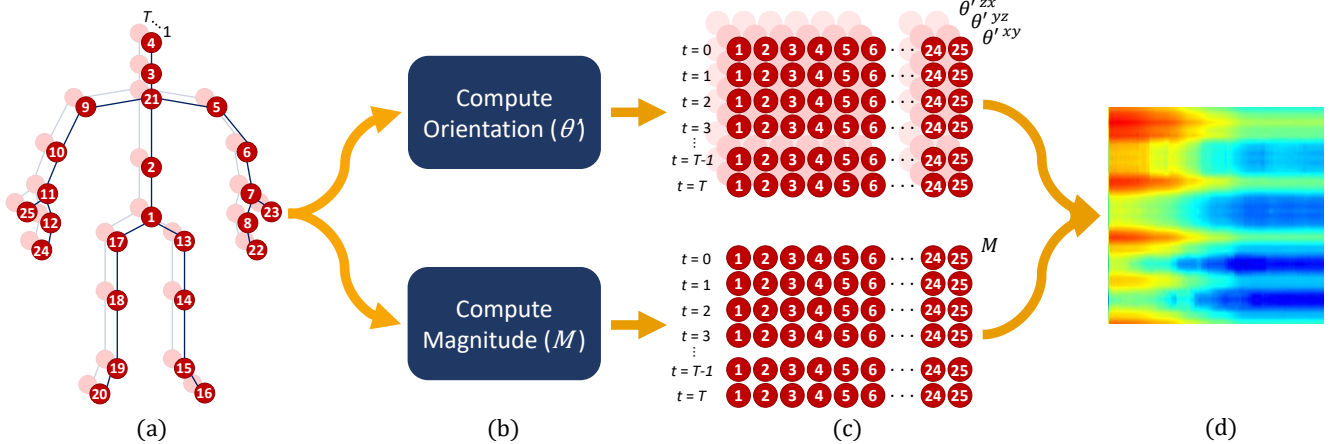


Figure 1. SkeleMotion representation. (a) Skeleton data sequence of T frames. (b) Computation of the magnitude and orientation from the joint movement. (c) θ' and M arrays: each row encodes the spatial information (relation between joint movements) while each column describes the temporal information for each joint movement. (d) Skeleton image after resizing and stacking of each axes.

number of frames may vary depending on the skeleton sequence of each video. Figure 1 gives an overview of our method for building the SkeleMotion representation.

3.1.1 Temporal Scale Aggregation (TSA)

Skeleton image representations in the literature basically encodes joint coordinates as channels. In view of that, it may cause a problem that the co-occurrence features are aggregated locally, being not able to capture long-range joint interactions involved in actions [14]. Moreover, one drawback of encoding motion values of joints is the noisy values that can be introduced to the representation due to small distance d between two frames. For instance, if the computation is performed considering two consecutive frames, it could add to the representation unnecessary motion of joints that are irrelevant to predict a specific action (e.g., motion of the head joint on a handshake action).

To overcome the aforementioned problems, we also propose a variation of our SkeleMotion representation by pre-computing the motion structure \mathcal{D} considering different d distances. For each of the motion structures \mathcal{D} , we compute its respective magnitude skeleton representation M and then stack them all into one single representation. The same is applied to compute the orientation skeleton representation θ , however a weighting scheme is applied during the filtering process explained before, as

$$\theta'_{c,t} = \begin{cases} 0, & \text{if } M_{c,t} < m \times d \\ \theta_{c,t}, & \text{otherwise} \end{cases} \quad (5)$$

Such technique adds more temporal dynamics to the representation by explicitly showing temporal scales to the network. In this way, the network can learn which movements are really important for the action learning and also being able to capture long-range joint interactions.

4. Experimental Results

In this section we present the experimental results obtained with the proposed skeleton image representation for the 3D action recognition problem. We compare it to other skeleton representations in the literature. Besides the classical skeleton image representation of Du et al. [4], we compare with other representations used by state-of-the-art approaches [32, 10, 13, 14, 34] as baselines on NTU RGB+D 60 [25]. We also compare our approach to state-of-the-art methods on the NTU RGB+D 120 [15]. To isolate only the contribution brought by SkeleMotion to the action recognition problem, all other representations were tested on the same datasets with the same split of training and testing data and using the same CNN.

4.1. Datasets

The **NTU RGB+D 60** [25] is publicly available 3D action recognition dataset. It consists of 56,880 videos from 60 action categories which are performed by 40 distinct subjects. The videos were collected by three Microsoft Kinect sensors. The dataset provides four different data information: (i) RGB frames; (ii) depth maps; (iii) 395 infrared sequences; and (iv) skeleton joints. There are two different evaluation protocols: cross-subject, which split the 40 subjects into training and testing; and cross-view, which uses samples from one camera for testing and the other two for training. The performance is evaluated by computing the average recognition across all classes.

The **NTU RGB+D 120** [15] is a large-scale 3D action recognition dataset captured under various environmental conditions. It consists of 114,480 RGB+D video samples captured using the Microsoft Kinect sensor. As in NTU RGB+D 60 [25], the dataset provides RGB frames, depth maps, infrared sequences and skeleton joints. It is com-




Figure 2. Network architecture employed for 3D action recognition.

posed by 120 action categories performed by 106 distinct subjects in a wide range of age distribution. There are two different evaluation protocols: cross-subject, which split the 106 subjects into training and testing; and cross-setup, which divides samples with even setup IDs for training (16 setups) and odd setup IDs for testing (16 setups). The performance is evaluated by computing the average recognition across all classes.

4.2. Implementation Details

To isolate only the contribution brought by the proposed representation to the action recognition problem, all compared skeleton image representations were implemented and tested on the same datasets and using the same network architecture. In view of that, we applied the same split of training and testing data, and employ the evaluation protocols and metrics proposed by the creators of the datasets.

The network architecture employed is a modified version of the CNN proposed by Li et al. [13]. They designed a small convolutional neural network which consists of three convolution layers and four fully-connected (FC) layers. However, here we modified it to a tiny version, employing the convolutional layers and only two FC layers. All convolutions have a kernel size of 3×3 , the first and second convolutional layers with a stride of 1 and the third one with a stride of 2. Max pooling and ReLU neuron are adopted and the dropout regularization ratio is set to 0.5. The learning rate is set to 0.001 and batch size is set to 1000. The training is stopped after 200 epochs. The loss function employed was the categorical cross-entropy. We opted for using such architecture since it demonstrated good performance and, according to the authors, it can be easily trained from scratch without any pre-training and is superior on its compact model size and fast inference speed as well. Figure 2 presents an overview of the employed architecture.

To cope with actions involving multi-person interaction (e.g., shaking hands), we apply a common choice in the literature which is to stack skeleton image representations of

different people as the network input.

To obtain the orientation skeleton image representation θ' we empirically set the parameter $m = 0.004$, as described in Section 3.

4.3. Evaluation

In this section, we present experiments for parameters optimization and report a comparison of our proposed skeleton representation. We used a subset of NTU RGB+D 60 [25] training set (considering cross-view protocol) to perform parameter setting and then used such parameter on the remaining experiments. We focused on the optimization of the number of temporal scales used on temporal scale aggregation (TSA).

Table 1. Action recognition accuracy (%) results on a subset of NTU RGB+D 60 [25] dataset by applying temporal scale aggregation (TSA) on our SkeleMotion representation.

	Temporal distances	Magnitude Acc. (%)	Orientation Acc. (%)
Two Temporal Scales	1, 5	64.9	62.4
	1, 10	67.4	62.9
	1, 15	66.0	64.1
	1, 20	66.1	63.5
Three Temporal Scales	1, 5, 10	68.6	64.6
	1, 10, 20	69.0	65.4
	5, 10, 15	70.1	65.1
Four Temporal Scales	1, 5, 10, 15	69.6	65.2
	5, 10, 15, 20	67.9	64.4

To set the number of temporal scales of our SkeleMotion approach, we empirically varied it from two to four temporal scales considering 20 frames in total. Table 1 shows the results obtained by such variation. We can see that the best result is obtained by using three temporal scales for both magnitude (5, 10, 15) and orientation (1, 10, 20). Moreover, we noticed that the performance tends to saturate or drop when considering four temporal scales.




Figure 3. The complementary between SkeleMotion (Magnitude-Orientation TSA) and Yang et al. [34] TSSI representation on NTU RGB+D 60 [25] dataset for cross-view protocol. Best viewed in color.

Table 2 presents a comparison of our approach with skeleton image representations of the literature. For the methods that have more than one “image” per representation ([32] and [10]), we stacked them to be used as input to the network. The same was performed for our SkeleMotion approach considering magnitude and orientation. Regarding the cross-subject protocol, the best result was obtained by Reference Joints representation from Ke et al., [10] achieving 70.8% of accuracy while our best result (SkeleMotion Magnitude (TSA)) achieves a competitive accuracy of 69.6%. It is worth noting that there is a considerable improvement of 12.8 (p.p.) obtained by SkeleMotion Magnitude (TSA) when compared to Li et al., [13] baseline, which also explicitly encode motion information. On the other side, the best result on cross-view protocol was obtained by our SkeleMotion Magnitude (TSA) approach achieving 80.1% of accuracy. There is an improvement of 4.5 percentage points (p.p.) when compared to the Tree Structure Skeleton Image (TSSI) from Yang et al., [34], which was the best baseline result. Again, there is a considerable improvement of 18.8 (p.p.) when compared to Li et al., [13] baseline.

Table 2. Action recognition accuracy (%) results on NTU RGB+D 60 [25] dataset. Results for the baselines were obtained running each method implementation.

	Approach	Cross-subject Acc. (%)	Cross-view Acc. (%)
Baselines	Du et al. [4]	68.7	73.0
	Wang et al. [32]	39.1	35.9
	Ke et al. [10]	70.8	75.5
	Li et al. [14]	56.8	61.3
	Yang et al. [34]	69.5	75.6
SkeleMotion results	Orientation	60.6	65.6
	Magnitude	58.4	64.2
	Orientation (TSA)	65.3	73.2
	Magnitude (TSA)	69.6	80.1

To exploit a possible complementarity of the tempo-

ral (our SkeleMotion) and spatial (Yang et al. [34]) skeleton image representations, we combined the different approaches by employing early and late fusion techniques. For the early fusion, we simply stacked the representations to be used as input to the network. On the other hand, the late fusion technique applied was a non-weighted linear combination of the prediction scores of each method. According to the results showed in Table 3, any type of combination performed with our SkeleMotion provides better results than their solo versions. Regarding cross-subject protocol, our best results achieves 73.5% of accuracy with early fusion technique against 76.5% of the late fusion approach. Furthermore, on cross-view protocol, our best results achieves 82.4% of accuracy with early fusion technique against 84.7% of the late fusion approach. Detailed improvements are shown in Figure 3.

Finally, Table 4 presents the experiments of our proposed skeleton image representation on the recent available NTU RGB+D 120 [15] dataset. Due to the results obtained on Table 3, here we employed the late fusion scheme for methods combination.

We obtained good results with our SkeleMotion representation outperforming many skeleton based methods [25, 9, 8, 18, 17, 20, 10, 16, 22, 19, 11, 23]. When combining our representation with Yang et al. [34] we achieve state-of-the-art results, outperforming the best reported method (Body Pose Evolution Map [23]) by up to 3.1 p.p. on cross-subject protocol and achieve competitive results on cross-setup protocol.

In comparison with LSTM approaches, we outperform the best reported method (Two-Stream Attention LSTM) by 1.7 p.p. using our skeleton image representation and 6.5 p.p. when combining it with Yang et al. [34] method on cross-subject protocol. Regarding the cross-setup protocol, we outperform them by 3.6 p.p. using our skeleton image representation fused with Yang et al. [34]. This indicates that our skeleton image representation approach used as input for CNNs leads to a better learning of temporal dynamics than the approaches that employs LSTM.

Table 3. Comparison between late and early fusion techniques on NTU RGB+D 60 [25] dataset.

Approach		Cross-subject Acc. (%)	Cross-view Acc. (%)
Early Fusion	Magnitude-Orientation	65.6	71.1
	Magnitude-Orientation + Yang et al.[34]	70.1	78.8
	Magnitude-Orientation (TSA)	70.5	78.7
	Magnitude (TSA) + Yang et al.[34]	71.7	82.4
	Orientation (TSA) + Yang et al.[34]	69.6	78.9
	Magnitude-Orientation (TSA) + Yang et al.[34]	73.5	82.1
Late Fusion	Magnitude-Orientation	65.6	71.1
	Magnitude-Orientation + Yang et al.[34]	73.2	79.5
	Magnitude-Orientation (TSA)	72.2	81.7
	Magnitude (TSA) + Yang et al.[34]	75.4	83.2
	Orientation (TSA) + Yang et al.[34]	73.6	80.6
	Magnitude-Orientation (TSA) + Yang et al.[34]	76.5	84.7

Table 4. Action recognition accuracy (%) results on NTU RGB+D 120 [15] dataset. Results for literature methods were obtained from [15].

Approach		Cross-subject Acc. (%)	Cross-setup Acc. (%)
Literature results	Part-Aware LSTM [25]	25.5	26.3
	Soft RNN [9]	36.3	44.9
	Dynamic Skeleton [8]	50.8	54.7
	Spatio-Temporal LSTM [18]	55.7	57.9
	Internal Feature Fusion [17]	58.2	60.9
	GCA-LSTM [20]	58.3	59.2
	Multi-Task Learning Network [10]	58.4	57.9
	FSNet [16]	59.9	62.4
	Skeleton Visualization (Single Stream) [22]	60.3	63.2
	Two-Stream Attention LSTM [19]	61.2	63.3
	Multi-Task CNN with RotClips [11]	62.2	61.8
	Body Pose Evolution Map [23]	64.6	66.9
	SkeleMotion results	Orientation (TSA)	52.2
Magnitude (TSA)		57.6	60.4
Magnitude-Orientation (TSA)		62.9	63.0
Magnitude-Orientation (TSA) + Yang et al.[34]		67.7	66.9

5. Conclusions and Future Works

In this work, we proposed a novel skeleton image representation to be used as input of CNNs, named SkeleMotion. The method is based on temporal dynamics encoding by explicitly using motion information (magnitude and orientation) of the skeleton joints. We further propose a variation of the magnitude skeleton representation considering different temporal scales in order to filter noisy motion values as well as aggregating more temporal dynamics to the representation. Experimental results on two publicly available datasets demonstrated the excellent performance of the proposed approach. Another interesting finding is that the combination of our representation with methods of the literature improves the 3D action recognition outperforming the state-of-the-art on NTU RGB+D 120 dataset.

Directions to future works include the evaluation of SkeleMotion with other distinct architectures. Moreover,

we intend to evaluate its behavior on 2D action datasets with skeletons estimated by methods of the literature.

Acknowledgments

The authors would like to thank the National Council for Scientific and Technological Development – CNPq (Grants 311053/2016-5, 204952/2017-4 and 438629/2018-3), the Minas Gerais Research Foundation – FAPEMIG (Grants APQ-00567-14 and PPM-00540-17) and the Coordination for the Improvement of Higher Education Personnel – CAPES (DeepEyes Project).

References

- [1] C. A. Caetano, V. H. C. D. Melo, J. A. dos Santos, and W. R. Schwartz. Activity recognition based on a magnitude-orientation stream network. In *SIBGRAP*, 2017.

- [2] V. Choutas, P. Weinzaepfel, J. Revaud, and C. Schmid. Position: Pose motion representation for action recognition. In *CVPR*, 2018.
- [3] M. Devanne, H. Wannous, S. Berretti, P. Pala, M. Daoudi, and A. Del Bimbo. 3-d human action recognition by shape analysis of motion trajectories on riemannian manifold. *IEEE Transactions on Cybernetics*, 2015.
- [4] Y. Du, Y. Fu, and L. Wang. Skeleton based action recognition with convolutional neural network. In *ACPR*, 2015.
- [5] M. A. Gowayyed, M. Torki, M. E. Hussein, and M. El-Saban. Histogram of oriented displacements (hod): Describing trajectories of human joints for action recognition. In *IJCAI*, 2013.
- [6] F. Han, B. Reily, W. Hoff, and H. Zhang. Space-time representation of people based on 3d skeletal data. *Computer Vision and Image Understanding*, 2017.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [8] J. Hu, W. Zheng, J. Lai, and J. Zhang. Jointly learning heterogeneous features for rgb-d activity recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [9] J. Hu, W. Zheng, L. Ma, G. Wang, J. Lai, and J. Zhang. Early action prediction by soft regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [10] Q. Ke, M. Bennamoun, S. An, F. Sohel, and F. Boussaid. A new representation of skeleton sequences for 3d action recognition. In *CVPR*, 2017.
- [11] Q. Ke, M. Bennamoun, S. An, F. Sohel, and F. Boussaid. Learning clip representations for skeleton-based 3d action recognition. *IEEE Transactions on Image Processing*, 2018.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [13] C. Li, Q. Zhong, D. Xie, and S. Pu. Skeleton-based action recognition with convolutional neural networks. In *ICMEW*, 2017.
- [14] C. Li, Q. Zhong, D. Xie, and S. Pu. Co-occurrence feature learning from skeleton data for action recognition and detection with hierarchical aggregation. In *IJCAI*, 2018.
- [15] J. Liu, A. Shahroudy, M. Perez, G. Wang, L.-Y. Duan, and A. C. Kot. Ntu rgb+d 120: A large-scale benchmark for 3d human activity understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [16] J. Liu, A. Shahroudy, G. Wang, L. Duan, and A. Kot Chichung. Skeleton-based online action prediction using scale selection network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [17] J. Liu, A. Shahroudy, D. Xu, A. C. Kot, and G. Wang. Skeleton-based action recognition using spatio-temporal lstm network with trust gates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [18] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *ECCV*, 2016.
- [19] J. Liu, G. Wang, L. Duan, K. Abdiyeva, and A. C. Kot. Skeleton-based human action recognition with global context-aware attention lstm networks. *IEEE Transactions on Image Processing*, 2018.
- [20] J. Liu, G. Wang, P. Hu, L. Duan, and A. C. Kot. Global context-aware attention lstm networks for 3d action recognition. In *CVPR*, 2017.
- [21] M. Liu, C. Chen, and H. Liu. 3d action recognition using data visualization and convolutional neural networks. In *ICME*, 2017.
- [22] M. Liu, H. Liu, and C. Chen. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recogn.*, 2017.
- [23] M. Liu and J. Yuan. Recognizing human actions as the evolution of pose estimation maps. In *CVPR*, 2018.
- [24] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- [25] A. Shahroudy, J. Liu, T. Ng, and G. Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016.
- [26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [27] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *AAAI Conference on Artificial Intelligence*, 2017.
- [28] V. Veeriah, N. Zhuang, and G.-J. Qi. Differential recurrent neural networks for action recognition. In *ICCV*, 2015.
- [29] J. Wang, Z. Liu, Y. Wu, and J. Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *CVPR*, 2012.
- [30] P. Wang, W. Li, C. Li, and Y. Hou. Action recognition based on joint trajectory maps with convolutional neural networks. *Knowledge-Based Systems*, 2018.
- [31] P. Wang, W. Li, P. Ogunbona, Z. Gao, and H. Zhang. Mining mid-level features for action recognition based on effective skeleton representation. In *DICTA*, 2014.
- [32] P. Wang, Z. Li, Y. Hou, and W. Li. Action recognition based on joint trajectory maps using convolutional neural networks. In *ACM International Conference on Multimedia (MM)*, 2016.
- [33] X. Yang and Y. L. Tian. Eigenjoints-based action recognition using naïve-bayes-nearest-neighbor. In *CVPRW*, 2012.
- [34] Z. Yang, Y. Li, J. Yang, and J. Luo. Action recognition with spatio-temporal visual attention on skeleton image sequences. *IEEE Transactions on Circuits and Systems for Video Technology*, 2018.
- [35] M. Zanfir, M. Leordeanu, and C. Sminchisescu. The moving pose: An efficient 3d kinematics descriptor for low-latency action recognition and detection. In *ICCV*, 2013.
- [36] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *ICCV*, 2017.