



HAL
open science

A Deep Learning Approach for Real-Time 3D Human Action Recognition from Skeletal Data

Huy-Hieu Pham, Houssam Salmane, Louahdi Khoudour, Alain Crouzil, Pablo Zegers, Sergio A. Velastin

► **To cite this version:**

Huy-Hieu Pham, Houssam Salmane, Louahdi Khoudour, Alain Crouzil, Pablo Zegers, et al.. A Deep Learning Approach for Real-Time 3D Human Action Recognition from Skeletal Data. 16th International Conference on Image Analysis and Recognition (ICIAR 2019can), Aug 2019, Waterloo, Canada. pp.18-32, 10.1007/978-3-030-27202-9_2 . hal-02883879

HAL Id: hal-02883879

<https://hal.science/hal-02883879>

Submitted on 29 Jun 2020

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.



Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in:
<http://oatao.univ-toulouse.fr/26211>

Official URL

https://doi.org/10.1007/978-3-030-27202-9_2

To cite this version: Pham, Huy-Hieu and Salmane, Houssam and Khoudour, Louahdi and Crouzil, Alain and Zegers, Pablo and Velastin, Sergio A. *A Deep Learning Approach for Real-Time 3D Human Action Recognition from Skeletal Data*. (2019) In: 16th International Conference on Image Analysis and Recognition (ICIAR 2019can), 27 August 2019 - 29 August 2019 (Waterloo, Canada).

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

A Deep Learning Approach for Real-Time 3D Human Action Recognition from Skeletal Data

Huy Hieu Pham^{1,2}(✉) , Houssam Salmane¹ , Louahdi Khoudour¹ ,
Alain Crouzil² , Pablo Zegers³ , and Sergio A. Velastin^{4,5}

¹ Cerema, Equipe-projet STI, 1 Avenue du Colonel Roche, 31400 Toulouse, France
{huy-hieu.pham,houssam.salmane,louahdi.khoudour}@cerema.fr

² Université Toulouse III - Paul Sabatier, Institut de Recherche en Informatique de
Toulouse, 31062 Cedex 9, Toulouse, France
alain.crouzil@irit.fr

³ Aparnix, La Gioconda 4355, Santiago, Chile
pablozegers@gmail.com

⁴ Cortexica Vision Systems Ltd., London, UK
sergio.velastin@ieee.org

⁵ Queen Mary University of London and Department of Computer Science,
University Carlos III of Madrid, Madrid, Spain

Abstract. We present a new deep learning approach for real-time 3D human action recognition from skeletal data and apply it to develop a vision-based intelligent surveillance system. Given a skeleton sequence, we propose to encode skeleton poses and their motions into a single RGB image. An Adaptive Histogram Equalization (AHE) algorithm is then applied on the color images to enhance their local patterns and generate more discriminative features. For learning and classification tasks, we design Deep Neural Networks based on the Densely Connected Convolutional Architecture (DenseNet) to extract features from enhanced-color images and classify them into classes. Experimental results on two challenging datasets show that the proposed method reaches state-of-the-art accuracy, whilst requiring low computational time for training and inference. This paper also introduces CEMEST, a new RGB-D dataset depicting passenger behaviors in public transport. It consists of 203 untrimmed real-world surveillance videos of realistic normal and anomalous events. We achieve promising results on real conditions of this dataset with the support of data augmentation and transfer learning techniques. This enables the construction of real-world applications based on deep learning for enhancing monitoring and security in public transport.

Keywords: Action recognition · Skeletal data · Enhanced-SPMF · DenseNet

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-27202-9_2) contains supplementary material, which is available to authorized users.

1 Introduction

Human Action Recognition or HAR for short, plays a crucial role in many computer vision applications such as intelligent surveillance, human-computer interaction or robotics. Although significant progress has been achieved, detecting accurately what humans do in unknown videos is still a challenging task due to numerous challenges, e.g. viewpoint changes, intra-class variation, or surrounding distractions [31]. At present, depth sensor-based HAR is considered as one of the best available methods for overcoming the above obstacles. Cost-effective depth sensors are able to provide 3D structural information of the human body, which is suitable for HAR task. In particular, most of these devices have integrated the real-time skeleton estimation algorithms [34] that are robust to surrounding distractions as well as invariant to camera viewpoints. Therefore, exploiting skeletal data for HAR opens up opportunities for addressing the limitations of RGB and depth modalities. In the literature of skeleton-based action recognition, there are two main issues that need to be solved. The first challenge is how to transform the raw skeleton sequences into an effective representation, which is able to capture the spatio-temporal dynamics of human motions. The second is to model and recognize actions using the motion representation obtained from skeletons. Previous works on this topic can be divided into two main groups: HAR based on hand-crafted features and HAR using deep learning models. The first group of methods extracts hand-crafted local features from skeleton joints and uses probabilistic graphical models such as Hidden Markov Model (HMM) [26], Conditional Random Field (CRF) [8], and Fourier Temporal Pyramid (FTP) [39] to model and classify actions. For instance, since the first work on 3D HAR from depth data was introduced [20], many methods for skeleton-based action recognition have been proposed [8, 25, 39, 41–43, 46, 47]. The common characteristic of these approaches is that, they extract the geometric features from the 3D coordinates of the skeleton joints and model their temporal information by a generative model. Although promising results have been achieved, most of these approaches are shallow, data-dependent and require a lot of feature engineering. E.g., they require pre-processing input data in which the skeleton sequences need to be segmented or aligned. In contrast, we propose a skeleton-based representation and a learning framework for 3D HAR that learns to recognize actions from the raw skeletons in an end-to-end manner, without dependence on the length of actions.

The second group considers skeleton-based action recognition as a time-series problem and proposes to use Recurrent Neural Networks with Long Short-Term Memory units (RNN-LSTMs) [12] to analyze the temporal evolutions of skeletons. They are considered as the most popular deep learning based approach for the HAR task from skeletons and have achieved high-level performance [6, 21, 23, 33, 35, 38, 51]. The temporal evolutions of skeletons are in fact spatio-temporal patterns. Thus, they can be modeled by memory cells in the structure of RNN-LSTMs. However, RNN-LSTM based methods tend to overemphasize the temporal information and lose spatial information of skeletons [6] – an important characteristic for 3D HAR. Another limitation of RNN-LSTM networks is that

they just model the overall temporal dynamics of actions without considering the detailed temporal dynamics of them [19]. Additionally, this approach considers skeletons as a kind of low-level feature by feeding raw skeletal data directly into the network. The huge number of input features makes RNN-LSTMs complex, time-consuming and easily lead to overfitting. Furthermore in many cases, RNN-LSTMs act as a classifier, which cannot extract high-level features for the HAR problem [32]. In this paper, we propose a CNN-based method to extract rich geometric motion features from skeleton sequences and model various temporal dynamics, including both short-term and long-term actions.

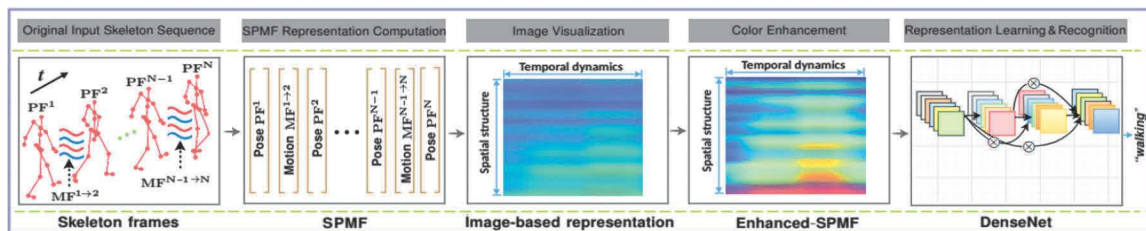


Fig. 1. Overview of the proposed approach for real-time 3D action recognition from skeletal data. Each skeleton sequence is encoded as a single RGB image via a skeleton-based representation called SPMF [27]. The SPMF is built from *Pose Feature* vectors (PFs) and *Motion Feature* vectors (MFs), which are estimated from the 3D coordinates of skeletons. A color enhancement technique [30] is adopted to enhance the local textures of SPMF to form the enhanced motion maps, namely Enhanced-SPMF. Finally, they are fed to a deep network for learning image features and performing action classification.

In contrast to the existing approaches, we aim to build an end-to-end deep learning framework for real-time action recognition from skeleton sequences. We believe that an effective motion representation is the key factor influencing recognition performance. Therefore, we propose to encode human poses and motions extracted from the 3D coordinates of skeleton joints into color images. These color-coded images are then enhanced in their local textures by an Adaptive Histogram Equalization (AHE) algorithm [30] before feeding into Deep Convolutional Neural Networks (D-CNNs), which are built based on the DenseNet architecture [14]. Before that, a smoothing filter is applied to reduce the effects of noise on the input skeletal data. The overview of the proposed method is illustrated in Fig. 1. Generally speaking, four hypotheses that motivate us to build a skeleton-based representation and design DenseNets for 3D HAR include: (1) human actions can be correctly represented via movements of the skeleton [16]; (2) spatio-temporal evolutions of skeletons can be transformed into color images – a kind of 3D tensor that can be effectively learned by D-CNNs [1, 3, 5]. This hypothesis was proved in our previous studies [27–29]; (3) compared to RGB and depth modalities, skeletal data has high-level information with much less complexity. This makes the learning model much simpler and requiring less computation, allowing us to build real-time deep learning framework for HAR

task; (4) DenseNet is currently one of the most effective CNN architecture for image recognition. It has a densely connected structure allowing maximal information flow and facilitates features reuse as each layer in its architecture has direct access to the features from previous layers. This helps DenseNet to improve its learning performance. Therefore, we explore and optimise this architecture for learning and recognizing human actions on the proposed image-based representation.

The main contributions of this work are three-fold. **First**, we introduce Enhanced-SPMF (Enhanced Skeleton Pose-Motion Feature) – a 3D motion representation for HAR tasks (Sect. 2). This is an extended representation of SPMF, which was presented in our previous work [27]. The new representation aims to improve the efficiency of the SPMF by using a smoothing filter on input skeleton sequences and a color enhancement technique that could make the proposed Enhanced-SPMF more robust and discriminative. An ablation study on the Enhanced-SPMF demonstrates that the new representation leads to better overall action recognition performance than the SPMF. **Second**, we introduce an end-to-end deep framework based on D-CNNs¹ for learning and recognizing actions from the Enhanced-SPMFs (Sect. 3). This approach is general in the sense that it can be applied to other data modalities, e.g. mocap data or the output of 3D pose estimation algorithms. The proposed method is evaluated on two highly competitive benchmark datasets and achieved state-of-the-art performance on both these two benchmark tasks with high computational efficiency (Sect. 4). **Finally**, we collect and introduce a new RGB-D dataset consisting of real-world surveillance videos for analyzing anomalous and normal events in public transport. Experimental results show that the proposed method achieves promising performance in realistic conditions (Sect. 4).

The rest of this paper is organized as follows: Sect. 2 presents the proposed skeleton-based representation. The proposed deep learning framework is presented in Sect. 3. Datasets and experiments are provided in Sect. 4, including a description of our dataset and the obtained results. Section 5 concludes the paper.

2 Enhanced Skeleton Pose-Motion Feature

One of the major challenges in exploiting D-CNNs for skeleton-based action recognition is how a skeleton sequence could be effectively represented and fed to the deep networks. As D-CNNs work well on still images [18], our idea therefore is to encode the spatial and temporal dynamics of skeletons into 2D images [28, 29]. Two essential elements for describing an action are static poses and their temporal dynamics. As shown by Zhang et al. in [50], the combination of too many geometric features will lead to lower performance than using only a single feature or several main features. Moreover, joint features such as joint distance and joint motion are stronger features than others [49]. Hence, we decide to transform these two important elements into the static spatial structure of

¹ Codes and models are available on our GitHub project at <https://bit.ly/2EC9vj9>.

a color image. The details of this idea are explained in our previous work [27], in which the spatio-temporal patterns of a skeleton sequence can be encoded into a single color image as a global representation, namely SPMF, via pose and motion feature vectors. Due to the limited space available, detailed description of the SPMF is not included. Instead, we refer the interested readers to [27] for further technical details. Figure 2 shows some SPMF representations in form of an image-based representation obtained from MSR Action3D dataset [20].




Fig. 2. Some SPMF representations obtained from the MSR Action3D dataset [20]. The change in color reflects the change of distance and orientation between the joints.

The color images obtained after the process of encoding mainly reflect the spatio-temporal distribution of skeleton joints. We observe that these images are represented by close contrast values, as can be seen in as Fig. 2. In this case, a color enhancement method can be useful for increasing the contrast of these representations and highlighting the texture and edges of motion maps. This helps to better distinguish similar actions. Therefore, it is necessary to enhance the local features on the generated color images. The Adaptive Histogram Equalization (AHE) [30] is a common approach for this task. This technique is capable of enhancing the local features of an image. Mathematically, let \mathbf{I} be a given image, represented as a r -by- c matrix of integer pixels with intensity levels in the range $[0, L - 1]$. The histogram of image \mathbf{I} will be defined by $H_k = \mathbf{n}_k$, where \mathbf{n}_k is the number of pixels with intensity k in \mathbf{I} . Hence, the probability of occurrence of intensity level k in \mathbf{I} is

$$p_k = \frac{\mathbf{n}_k}{r \times c}, \quad (k = 0, 1, 2, \dots, L - 1). \quad (1)$$

The histogram equalized image will be formed by transforming the pixel intensities, n , of \mathbf{I} by the function

$$T(n) = \text{floor}((L - 1) \sum_{k=0}^n p_k), \quad (n = 0, 1, 2, \dots, L - 1), \quad (2)$$

The Histogram Equalization (HE) method is used for increasing the global contrast of the image. However, it cannot solve the problem of increasing the local contrast. To do this, the image needs to be divided into \mathcal{R} regions and the HE is then applied in each region. This technique is called the Adaptive Histogram Equalization (AHE). Figure 3 shows samples of the enhanced motion maps with $\mathcal{R} = 8$, which we refer to it as Enhanced-SPMF for some actions from the MSR Action 3D dataset [20].

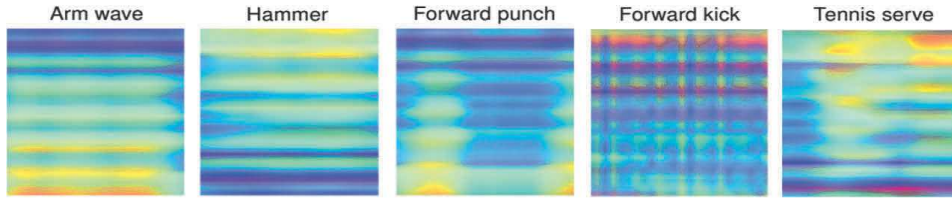


Fig. 3. The corresponding Enhanced-SPMF representations after applying the AHE algorithm [30]. This color enhancement step could make the proposed Enhanced-SPMF more robust and discriminative for the representation learning phase with D-CNNs later.

3 Deep Learning Model

This section reviews the key ideas behind the DenseNet architecture and presents the proposed deep networks for recognizing actions from the Enhanced-SPMFs.

3.1 DenseNet Review

DenseNet [14], a recently proposed CNN model, has some interesting properties. Each layer is connected to all the others within a dense block and all layers can access feature maps from their preceding layers. Besides, each layer receives direct information flow from the loss function through shortcut connections. These properties make DenseNet less prone to overfitting for supervised learning problems. Traditional CNN architectures use the output feature maps \mathbf{x}_{l-1} of the $(l-1)^{\text{th}}$ layer as input to the l^{th} layer and learn a mapping function $\mathbf{x}_l = \mathcal{H}_l(\mathbf{x}_{l-1})$. Here, $\mathcal{H}_l(\cdot)$ is a non-linear transformation that is usually implemented by a series of operations such as Convolution (**Conv.**), Rectified Linear Unit (**ReLU**) [7], Pooling, and Batch Normalization (**BN**) [15]. When increasing the depth of the network, the problem of optimization becomes complex due to the vanishing-gradient problem and the degradation phenomenon [9]. To solve these problems, [11] introduced ResNet. The key idea behind the ResNet architecture is to add shortcut connections that bypass the non-linear transformations $\mathcal{H}_l(\cdot)$ with an identity function $id(\mathbf{x}) = \mathbf{x}$. Inspired by the philosophy of ResNet, to maximize information flow through layers, Huang et al. [14] proposed DenseNet in which the l^{th} layer in a dense block receives the feature maps of all preceding layers as inputs. That means

$$\mathbf{x}_l = \mathcal{H}_l(\text{concat}[\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{l-1}]), \quad (3)$$

where $\text{concat}[\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{l-1}]$ is a single tensor constructed by concatenation of the previous layer output feature maps. All layers receive direct supervision signal from the loss function through the shortcut connections. Therefore, DenseNets are easy to optimize and resistant to overfitting. In DenseNet, multiple dense blocks are connected via transition layers. Each block with its transition layer produces k feature maps and the parameter k is called as the “growth rate” of the network. The function $\mathcal{H}_l(\cdot)$ in the original work [14] is a composite function of three consecutive layers: **BN-ReLU-Conv.**

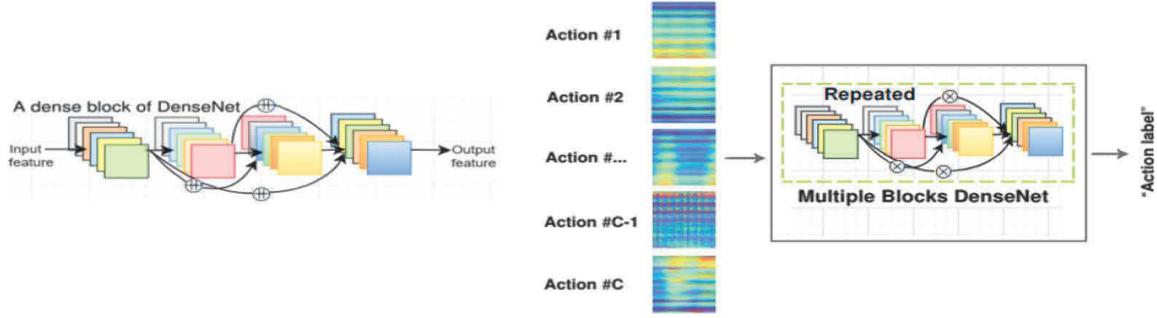


Fig. 4. A DenseNet block (**left**). The symbols \oplus denotes the concatenation operator. We explore and optimize this architecture for learning and recognizing human actions on the proposed image-based representation (**right**).

3.2 Network Design

We design D-CNNs based on the DenseNet architecture [14] to learn and classify actions on the Enhanced-SPMF. To study how performance varies with architecture size, we test three different configurations of DenseNet: {DenseNet-16, $k = 12$ }; {DenseNet-28, $k = 12$ }; and {DenseNet-40, $k = 12$ }. Here, the numbers 16, 28, 40 refer to the depth of the network and k is the network growth rate. For computational efficiency, we use three dense blocks on 32×32 input images. The $\mathcal{H}_l(\cdot)$ function is implemented by a sequence of layers: Batch Normalization (**BN**), advanced activation layer named Exponential Linear Units (**ELU**) [4] and 3×3 Convolution (**Conv**). Dropout [4] with a rate of 0.2 is used after each **Conv**. to prevent overfitting. The proposed networks can be trained in an end-to-end manner by gradient descent using Adam update rule [17]. During training, we minimize a cross-entropy loss function between the true action label \mathbf{y} and the predicted action $\hat{\mathbf{y}}$ over the training samples \mathcal{X} , by solving the following optimization problem

$$\text{Arg min}_{\mathcal{W}}(\mathcal{L}_{\mathcal{X}}(\mathbf{y}, \hat{\mathbf{y}})) = \text{Arg min}_{\mathcal{W}} \left(-\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^C \mathbf{y}_{ij} \log \hat{\mathbf{y}}_{ij} \right), \quad (4)$$

where \mathcal{W} is the set of weights that will be learned by the model, M denotes the number of samples in training set \mathcal{X} and C is the number of action classes.

4 Experiments

The proposed method is first evaluated on two challenging datasets: the MSR Action3D and NTU RGB+D (Sect. 4.1). We then introduce the CEMEST dataset² and report experimental results on this dataset. The implementation details of the proposed D-CNNs are also provided in this section (Sect. 4.2).

² Created by Cerema and Tisséo public transport in France and available for research purposes from <https://bit.ly/2SNbrdE>.

4.1 Datasets and Settings

MSR Action3D Dataset [20]: This dataset contains 20 actions performed by 10 subjects. Each skeleton is composed of 20 key joints. Experiments were conducted on 557 action sequences. We follow the protocol proposed by [20]. Specifically, the whole dataset is divided into three subsets: AS1, AS2 and AS3. For each subset, five subjects are selected for training and the rest are used for testing (see Supplemental Material). Data augmentation techniques including random cropping, vertically flipping, and Gaussian filtering have been applied on this dataset.

NTU RGB+D Dataset [33]: This Kinect 2 captured dataset is a very large-scale RGB+D dataset. It is currently the largest and state-of-the-art dataset that provides skeletal data for 3D HAR. The NTU RGB+D has more than 56 thousand video samples, 4 millions frames, collected from 40 distinct subjects for 60 different action classes. Each skeleton contains the 3D coordinates of 25 body joints. The authors of this dataset suggested two different evaluation criteria, including Cross-Subject and Cross-View evaluations. For the Cross-Subject setting, the sequences performed by 20 subjects are used for training and the rest sequences are used for testing. In Cross-View setting, the sequences provided by cameras 2 and 3 are used for training while sequences from camera 1 are used for testing (see Supplemental Material). Due to the very large-scale nature of the NTU RGB+D dataset, we do not apply any data augmentation technique on this dataset.

CEMEST Dataset: We have collected a new RGB-D dataset, called CEMEST (CErema METro STation dataset) using Kinect v2 sensor and carried out experiments on this dataset to verify the effectiveness of the proposed method on a




Fig. 5. Some samples from the CEMEST dataset: (a), (b) *crossing over the barriers*, (c), (d) *jumping over the ticket barriers* and (e), (f) *sneaking under ticket barriers*.

real-world dataset. The CEMEST was made at a metro station in France without any control of the passenger behavior as well as illumination. It contains three actions including both normal and abnormal behaviors: *crossing normally over the barriers*, *jumping over the barriers*, and *sneaking under the barriers*. These three behaviors are taken into account for acquisition because they have a significant impact on monitoring and management in public transport. As an example, the French National Railway Company (SNCF) reported that they lost €500 million every year through people trying to cheat the ticket system [37]. In summary, this dataset provides RGB, depth and skeletal data. The skeleton sequences are extracted by Kinect SDK with 25 key joints for each subject, at a frame rate of 30 FPS. All recorded sequences are manually segmented and labeled. Figure 5 shows some samples from the CEMEST. We carried out two experimental evaluations on this dataset. In the first setting, we randomly chose 67% of the data as training set and the remaining 33% is used for testing. In the second setting, the proposed networks are trained on a combination dataset, which is created from a portion of the MSR Action3D [20] and NTU RGB+D [33] datasets (see Supplementary Material for more details). To ensure the number of samples in each action class is balanced, we augmented samples in the MSR Action3D to match the size of the larger dataset. The pre-trained model is then deployed on the CEMEST dataset in the hope that transfer learning will help to solve overfitting problem when training on small dataset. In both experiments, data augmentation (i.e. cropping, flipping, Gaussian filtering) has been used.

4.2 Implementation Details

The Enhanced-SPMFs are computed directly from skeletons without using a fixed number of frames. The proposed DenseNets were implemented in Python using Keras. For training, we use mini-batches of 64 images. The weights are initialized as [10]. Adam optimizer [17] is used with an initial learning rate $\eta = 3e-4$. All networks are trained for 250 epochs from scratch.

4.3 Experimental Results and Evaluation

Experimental results and comparison of the proposed method with existing approaches on the MSR Action3D dataset are summarized in Table 1. The DenseNet-40 achieves an average accuracy of 99.10% over three subsets, which outperformed previous approaches by [2, 6, 20, 39, 42, 44, 48] and surpassed our previous work on SPMF [27]. Figure 6 (left) shows an example of the learning curves of the network during training on this dataset.

For the NTU RGB+D dataset, as shown in Table 2, the proposed DenseNet-40 achieves an accuracy of 79.95% on the Cross-Subject and 87.52% on Cross-View evaluations, respectively. These results demonstrate the effectiveness of the proposed representation and deep learning framework since they surpassed previous state-of-the-art approaches reported in [6, 13, 21, 22, 24, 33, 39, 40, 50] as well as a higher level of performance than SPMF [27]. Figure 6 (right) shows the training loss and test accuracy of the proposed DenseNet-40 on the NTU RGB+D

Table 1. Experimental results and comparison with the state-the-art approaches on the MSR Action3D dataset [20]. The best accuracies are in *italics*. Results that surpass previous works are in **bold**.

Method (protocol of [20])	AS1	AS2	AS3	Aver.
Bag of 3D Points [20]	72.90%	71.90%	71.90%	74.70%
Depth Motion Maps [2]	96.20%	83.20%	92.00%	90.47%
Bi-LSTM [36]	92.72%	84.93%	97.89%	91.84%
Lie Group [39]	95.29%	83.87%	98.22%	92.46%
Hierarchical RNN [6]	99.33%	94.64%	95.50%	94.49%
Graph-Based Motion [42]	93.60%	95.50%	95.10%	94.80%
ST-NBNN [44]	91.50%	95.60%	97.30%	94.80%
ST-NBMIM [45]	92.50%	95.60%	98.20%	95.30%
S-T Pyramid [48]	<i>99.10%</i>	<i>92.90%</i>	96.40%	96.10%
SPMF [27]	97.54%	98.73%	99.41%	98.56%
Enhanced-SPMF DenseNet-16 (ours)	98.05%	98.38%	98.80%	98.41%
Enhanced-SPMF DenseNet-28 (ours)	98.44%	98.47%	99.18%	98.70%
Enhanced-SPMF DenseNet-40 (ours)	98.88%	<i>99.05%</i>	99.24%	<i>99.10%</i>

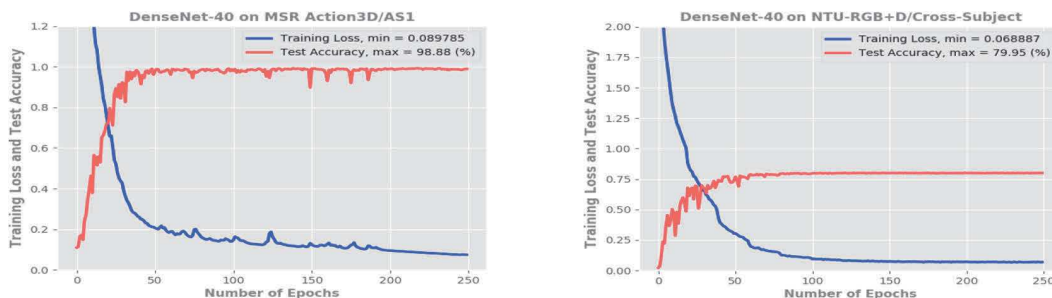


Fig. 6. Learning curves of the proposed DenseNet-40 network on the Enhanced-SPMFs obtained from the MSR Action3D and NTU RGB+D

dataset. On the CEMEST dataset, an accuracy of 91.18% has been made by the DenseNet-40 in the first setting. In the second setting, transfer learning is used. The experimental results show that the proposed method reached an accuracy of 95.70%, increasing the performance by nearly 5% compared to the first experiment. This could be explained by the fact that since the CEMEST dataset is quite small, it benefits from the knowledge transfer coming from larger datasets such as the MSR Action3D and NTU RGB+D datasets. This result indicates that the use of data augmentation and transfer learning is crucial to address the small amount of samples in real-world datasets. Figure 7 shows learning curves of the proposed deep learning networks on the CEMEST dataset from scratch (Fig. 7a–c), pre-training on the combined dataset (Fig. 7d–f) and fine-tuning on CEMEST dataset (Fig. 7g–i).

Table 2. Experimental results and comparison with the state-the-art approaches on the NTU RGB+D dataset [33]. The best accuracies are in *italics*. Results that surpass previous works are in **bold**.

Method (protocol of [33])	Cross-Subject	Cross-View
Lie Group [39]	50.10%	52.80%
Hierarchical RNN [6]	59.07%	63.97%
Dynamic Skeletons [13]	60.20%	65.20%
Two-Layer P-LSTM [33]	62.93%	70.27%
ST-LSTM Trust Gates [21]	69.20%	77.70%
Geometric Features [50]	70.26%	82.39%
Two-Stream RNN [40]	71.30%	79.50%
Enhanced Skeleton [24]	75.97%	82.56%
GCA-LSTM [22]	76.10%	84.00%
SPMF [27]	78.89%	86.15%
Enhanced-SPMF DenseNet-16 (ours)	77.89%	86.55%
Enhanced-SPMF DenseNet-28 (ours)	79.07%	86.82%
Enhanced-SPMF DenseNet-40 (ours)	<i>79.95%</i>	<i>87.52%</i>

4.4 An Ablation Study on Enhanced-SPMF

We believe that the use of the smoothing filter and the AHE algorithm [30] helps the proposed representation to be more discriminative, which improves recognition accuracy. To verify this hypothesis, we carried out an ablation study on the proposed representation by removing the color enhancement module and seeing how that affects performance. We observed that this kind of transformation is needed for improving learning performance of deep neural networks. Specifically, we trained the proposed DenseNet-40 on both the SPMFs and Enhanced-SPMFs provided by MSR Action3D dataset [20]. During training, the same hyper-parameters and training methodology were applied. The experimental results indicate that the proposed deep network achieves better recognition accuracy when trained on the Enhanced-SPMF (+1.42%). This result validates our hypothesis above.

4.5 Computational Efficiency Evaluation

The proposed learning framework comprises three main stages: **(1)** the computation of Enhanced-SPMF; **(2)** the training stage; and **(3)** the inference stage. To evaluate the computational efficiency of this method, we measure the execution time of each stage on the AS1 subset/MSR Action3D dataset with the proposed DenseNet-40 network, which only has 1.0M parameters. With the implementation in Python using Keras and training on a single GTX Ti 1080 GPU, the

training process takes less than one hour to reach convergence. While the inference stage, including the stage (1) that is executed on a CPU and the stage (3), takes an average of 0.175s per sequence without parallel processing. This result verifies the appropriateness of our method in terms of computational cost. Additionally, the computation of the Enhanced-SPMF can be implemented on a GPU for real-time applications.

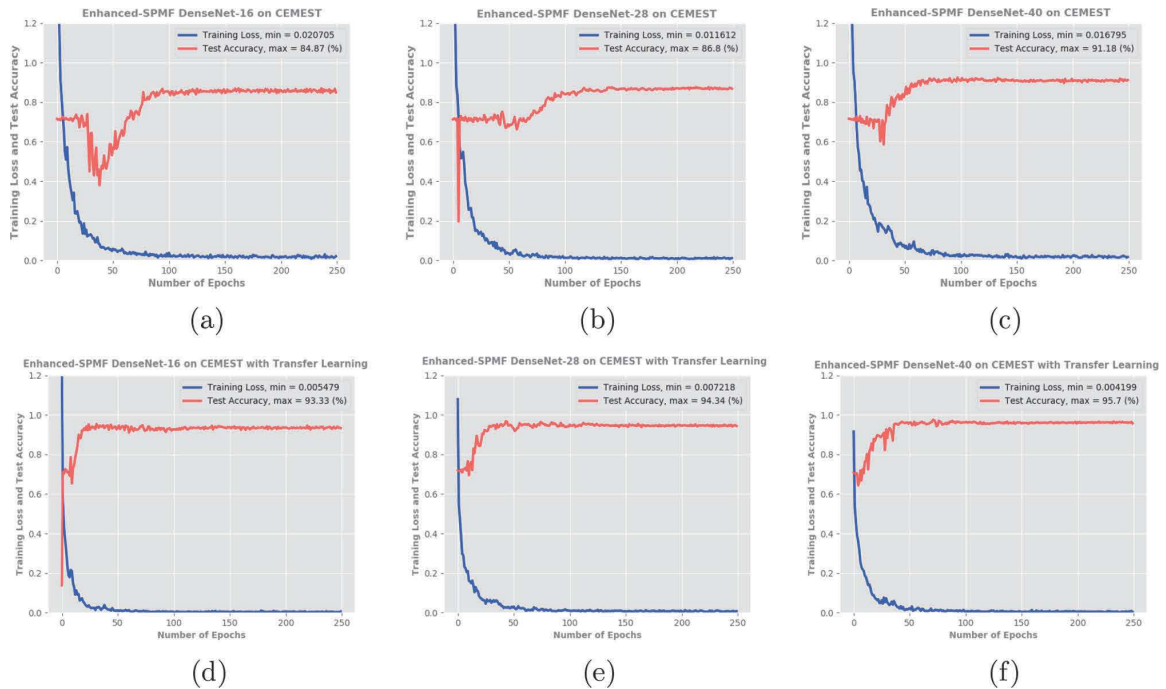


Fig. 7. Learning curves of the three proposed deep networks (DenseNet-16, DenseNet-28, DenseNet-40) on CEMEST dataset when trained from scratch (a)-(b)-(c) and fine-tuned on CEMEST dataset (d)-(e)-(f). Our best configuration (DenseNet-40) achieved an accuracy of 91.18% when trained on the CEMEST dataset from scratch. With the support of transfer learning, the proposed method reached an accuracy of 95.70%, increasing the recognition accuracy by nearly 5%.

5 Conclusions

We introduced a deep learning framework for 3D action recognition from skeletal data. A new motion representation that captures the spatio-temporal patterns of skeleton movements and encodes them into color images has been proposed. Densely connected networks have been designed to learn and recognize actions from the proposed representation in an end-to-end manner. Experiments on two public datasets have demonstrated the effectiveness of our method, both in terms of accuracy as well as computational time. We also introduced CEMEST, a new real-world surveillance dataset containing both normal and anomalous events for studying human behaviors in public transport. Experimental results on this dataset show that the proposed deep learning based approach achieved promising

results. We are currently expanding this study by adding more visual evidence to the network in order to further gains in performance. A new approach for 3D pose estimation will also be studying to replace depth sensors. The preliminary results are encouraging.

Acknowledgements. This research was supported by the Cerema, France. Sergio A. Velastin is grateful for funding from the Universidad Carlos III de Madrid, the EU’s 7th Framework Programme for Research, Technological Development and demonstration (grant 600371), Ministerio de Economía, Industria y Competitividad (COFUND2013-51509), Ministerio de Educación, cultura y Deporte (CEI-15-17) and Banco Santander.

References

1. Bilen, H., Fernando, B., Gavves, E., Vedaldi, A.: Action recognition with dynamic image networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(12), 2799–2813 (2018)
2. Chen, C., Liu, K., Kehtarnavaz, N.: Real-time human action recognition based on depth motion maps. *Journal of Real-Time Image Processing* **12**(1), 155–163 (2016)
3. Choutas, V., Weinzaepfel, P., Revaud, J., Schmid, C.: Potion: Pose motion representation for action recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7024–7033 (2018)
4. Clevert, D.A., Unterthiner, T., Hochreiter, S.: Fast and accurate deep network learning by Exponential Linear Units (ELUs). *arXiv preprint arXiv:1511.07289* (2015)
5. Ding, Z., Wang, P., Ogunbona, P.O., Li, W.: Investigation of different skeleton features for cnn-based 3d action recognition. In: *2017 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. pp. 617–622. IEEE (2017)
6. Du, Y., Wang, W., Wang, L.: Hierarchical recurrent neural network for skeleton based action recognition. In: *IEEE CVPR*. pp. 1110–1118 (2015)
7. Glorot, X., Bordes, A., Bengio, Y.: Deep sparse rectifier neural networks. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. pp. 315–323 (2011)
8. Han, L., Wu, X., Liang, W., Hou, G., Jia, Y.: Discriminative human action recognition in the learned hierarchical manifold space. *Image and Vision Computing* **28**(5), 836–849 (2010)
9. He, K., Sun, J.: Convolutional neural networks at constrained time cost. In: *IEEE CVPR*. pp. 5353–5360 (2015)
10. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In: *IEEE ICCV*. pp. 1026–1034 (2015)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE CVPR*. pp. 770–778 (2016)
12. Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural Computation* **9**(8), 1735–1780 (1997)
13. Hu, J., Zheng, W.S., Lai, J.H., Jianguo, Z.: Jointly learning heterogeneous features for RGB-D activity recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**, 2186–2200 (2015)
14. Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L.: Densely connected convolutional networks. In: *IEEE CVPR*. p. 3 (2017)

15. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML. pp. 448–456 (2015)
16. Johansson, G.: Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics* **14**(2), 201–211 (1973)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
18. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436 (2015)
19. Lee, I., Kim, D., Kang, S., Lee, S.: Ensemble deep learning for skeleton-based action recognition using temporal sliding lstm networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1012–1020 (2017)
20. Li, W., Zhang, Z., Liu, Z.: Action recognition based on a bag of 3D points. In: IEEE CVPR. pp. 9–14 (2010)
21. Liu, J., Shahroudy, A., Xu, D., Wang, G.: Spatio-temporal LSTM with trust gates for 3D human action recognition. In: ECCV. pp. 816–833 (2016)
22. Liu, J., Wang, G., Duan, L.Y., Abdiyeva, K., Kot, A.C.: Skeleton-based human action recognition with global context-aware attention LSTM networks. *IEEE Transactions on Image Processing* **27**(4), 1586–1599 (2018)
23. Liu, J., Wang, G., Hu, P., Duan, L.Y., Kot, A.C.: Global context-aware attention LSTM networks for 3D action recognition. In: IEEE CVPR. pp. 3671–3680 (2017)
24. Liu, M., Liu, H., Chen, C.: Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recognition* **68**, 346–362 (2017)
25. Luo, J., Wang, W., Qi, H.: Group sparsity and geometry constrained dictionary learning for action recognition from depth maps. In: IEEE ICCV. pp. 1809–1816 (2013)
26. Lv, F., Nevatia, R.: Recognition and segmentation of 3D human action using HMM and multi-class Adaboost. In: ECCV. pp. 359–372 (2006)
27. Pham, H., Khoudour, L., Crouzil, A., Zegers, P., Velastin, S.A.: Skeletal movement to color map: A novel representation for 3D action recognition with Inception Residual networks. In: IEEE International Conference on Image Processing (ICIP). pp. 3483–3487 (2018)
28. Pham, H.H., Khoudour, L., Crouzil, A., Zegers, P., Velastin, S.: Learning to Recognize 3D Human Action from A New Skeleton-based Representation Using Deep Convolutional Neural Networks. *IET Computer Vision* (2018)
29. Pham, H.H., Khoudour, L., Crouzil, A., Zegers, P., Velastin, S.A.: Exploiting deep residual networks for human action recognition from skeletal data. *Computer Vision and Image Understanding* **170**, 51–66 (2018)
30. Pizer, S.M., Amburn, E.P., Austin, J.D., Cromartie, R., Geselowitz, A., Greer, T., ter Haar Romeny, B., Zimmerman, J.B., Zuiderveld, K.: Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing* **39**(3), 355–368 (1987)
31. Poppe, R.: A survey on vision-based human action recognition. *Image and Vision Computing* **28**(6), 976–990 (2010)
32. Sainath, T.N., Vinyals, O., Senior, A., Sak, H.: Convolutional, Long Short-Term Memory, Fully Connected Deep Neural Networks. In: IEEE ICASSP. pp. 4580–4584 (2015)
33. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: NTU RGB+D: A large scale dataset for 3D human activity analysis. In: IEEE CVPR. pp. 1010–1019 (2016)
34. Shotton, J., Sharp, T., Kipman, A., Fitzgibbon, A., Finocchio, M., Blake, A., Cook, M., Moore, R.: Real-time human pose recognition in parts from single depth images. *Communications of the ACM* **56**(1), 116–124 (2013)

35. Si, C., Jing, Y., Wang, W., Wang, L., Tan, T.: Skeleton-based action recognition with spatial reasoning and temporal stack learning. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 103–118 (2018)
36. Tanfous, A.B., Drira, H., Amor, B.B.: Coding Kendall’s shape trajectories for 3D action recognition. In: IEEE CVPR. pp. 2840–2849 (2018)
37. The Local: SNCF increases fines for ticket dodgers. <https://bit.ly/2mYaJwW> (2015), published 20 February 2015. Accessed 10 July 2018
38. Veeriah, V., Zhuang, N., Qi, G.J.: Differential recurrent neural networks for action recognition. In: IEEE ICCV. pp. 4041–4049 (2015)
39. Vemulapalli, R., Arrate, F., Chellappa, R.: Human action recognition by representing 3D skeletons as points in a lie group. In: IEEE CVPR. pp. 588–595 (2014)
40. Wang, H., Wang, L.: Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks. In: IEEE CVPR. pp. 3633–3642 (2017)
41. Wang, J., Liu, Z., Wu, Y., Yuan, J.: Mining actionlet ensemble for action recognition with depth cameras. In: IEEE CVPR. pp. 1290–1297 (2012)
42. Wang, P., Yuan, C., Hu, W., Li, B., Zhang, Y.: Graph based skeleton motion representation and similarity measurement for action recognition. In: ECCV. pp. 370–385 (2016)
43. Wang, P., Li, W., Ogunbona, P., Gao, Z., Zhang, H.: Mining mid-level features for action recognition based on effective skeleton representation. In: International Conference on Digital Image Computing: Techniques and Applications (DICTA). pp. 1–8 (2014)
44. Weng, J., Weng, C., Yuan, J.: Spatio-temporal Naive-Bayes Nearest-Neighbor (ST-NBNN) for skeleton-based action recognition. In: IEEE CVPR. pp. 4171–4180 (2017)
45. Weng, J., Weng, C., Yuan, J., Liu, Z.: Discriminative spatio-temporal pattern discovery for 3D action recognition. IEEE Transactions on Circuits and Systems for Video Technology pp. 1–1 (2018)
46. Wu, D., Shao, L.: Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition. In: IEEE CVPR. pp. 724–731 (2014)
47. Xia, L., Chen, C.C., Aggarwal, J.: View invariant human action recognition using histograms of 3D joints. In: IEEE CVPR. pp. 20–27 (2012)
48. Xu, H., Chen, E., Liang, C., Qi, L., Guan, L.: Spatio-temporal pyramid model based on depth maps for action recognition. In: IEEE International Workshop on Multimedia Signal Processing (MMSP). pp. 1–6 (2015)
49. Yun, K., Honorio, J., Chattopadhyay, D., Berg, T.L., Samaras, D.: Two-person interaction detection using body-pose features and multiple instance learning. In: IEEE CVPR. pp. 28–35 (2012)
50. Zhang, S., Liu, X., Xiao, J.: On geometric features for skeleton-based action recognition using multilayer LSTM networks. In: IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 148–157 (2017)
51. Zhu, W., et al.: Co-occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks. In: AAAI. p. 8 (2016)