



Human Motion Analysis with Deep Metric Learning

Huseyin Coskun^{1(✉)}, David Joseph Tan^{1,2}, Sailesh Conjeti¹, Nassir Navab^{1,2},
and Federico Tombari^{1,2}

¹ Technische Universität München, Munich, Germany
huseyin.coskun@tum.de

² Pointu3D GmbH, Munich, Germany

Abstract. Effectively measuring the similarity between two human motions is necessary for several computer vision tasks such as gait analysis, person identification and action retrieval. Nevertheless, we believe that traditional approaches such as L2 distance or Dynamic Time Warping based on hand-crafted local pose metrics fail to appropriately capture the semantic relationship across motions and, as such, are not suitable for being employed as metrics within these tasks. This work addresses this limitation by means of a triplet-based deep metric learning specifically tailored to deal with human motion data, in particular with the problem of varying input size and computationally expensive hard negative mining due to motion pair alignment. Specifically, we propose (1) a novel metric learning objective based on a triplet architecture and Maximum Mean Discrepancy; as well as, (2) a novel deep architecture based on attentive recurrent neural networks. One benefit of our objective function is that it enforces a better separation within the learned embedding space of the different motion categories by means of the associated distribution moments. At the same time, our attentive recurrent neural network allows processing varying input sizes to a fixed size of embedding while learning to focus on those motion parts that are semantically distinctive. Our experiments on two different datasets demonstrate significant improvements over conventional human motion metrics.

1 Introduction

In image-based human pose estimation, the similarity between two predicted poses can be precisely assessed through conventional approaches that either evaluate the distance between corresponding joint locations [8, 28, 43] or the average difference of corresponding joint angles [24, 37]. Nevertheless, when human poses have to be compared across a temporal set of frames, the assessment of the similarity between two sequences of poses or motion becomes a non-trivial problem. Indeed, human motion typically evolves in a different manner on different sequences, which means that specific pose patterns tend to appear at different

H. Coskun and D.J. Tan—Equal contribution.




Fig. 1. When asked to measure the similarity to a query sequence (“Walking”, top), both the L2 and the DTW measures judge the unrelated sequence (“Standing”, bottom) as notably more similar compared to a semantically correlated one (“Walking”, middle). Conversely, our learned metric is able to capture the contextual information and measure the similarity correctly with respect to the given labels.

time instants on sequences representing the same human motion: see, *e.g.*, the first two sequences in Fig. 1, which depict two actions belonging to the same class. Moreover, these sequences result also in varying length (*i.e.*, a different number of frames), this making the definition of a general similarity measure more complicated. Nevertheless, albeit challenging, estimating the similarity between human poses across a sequence is a required step in human motion analysis tasks such as action retrieval and recognition, gait analysis and motion-based person identification.

Conventional approaches deployed to compare human motion sequences are based on estimating the L2 displacement error [23] or Dynamic Time Warping (DTW) [42]. Specifically, the former computes the squared distance between corresponding joints in the two sequences at a specific time t . As shown by Martinez *et al.* [23], such measure tends to disregard the specific motion characteristics, since a constant pose repeated over a sequence might turn out to be a better match to a reference sequence than a visually similar motion with a different temporal evolution. On the other hand, DTW tries to alleviate this problem by warping the two sequences via compressions or expansions so to maximize the matching between local poses. Nevertheless, DTW can easily fail in appropriately estimating the similarity when the motion dynamic in terms of peaks and plateaus exhibits small temporal variations, as shown in [18]. As an example, Fig. 1 illustrates a typical failure case of DTW when measuring the similarity among three human motions. Although the first two motions are visually similar to each other while the third one is unrelated to them, DTW estimates a smaller distance between the first and the third sequence. In general, neither the DTW nor the L2 metrics can comprehensively capture the semantic relationship between two sequences since they disregard the contextual information (in the temporal sense), this limiting their application in the aforementioned scenarios.

The goal of this work is to introduce a novel metric for estimating the similarity between two human motion sequences. Our approach relies on deep metric learning that uses a neural network to map high-dimensional data to a low-dimensional embedding [31, 33, 35, 45]. In particular, our first contribution is to design an approach so to map semantically similar motions over nearby loca-

tions in the learned embedding space. This allows the network to express a similarity measure that strongly relies on the motion’s semantic and contextual information. To this end, we employ a novel objective function based on the Maximum Mean Discrepancy (MMD) [14], which enforces motions to be embedded based on their distribution moments. The main advantage with respect to standard triplet loss learning is represented by the fact that our approach, being based on distributions and not samples, does not require hard negative mining to converge, which is computationally expensive since finding hard negatives in a human motion datasets requires the alignment of sequence pairs, which has an $O(n^2)$ complexity (n being the sequence length). As our second main contribution, we design a novel deep learning architecture based on attentive recurrent neural networks (RNNs) which exploits attention mechanisms to map an arbitrary input size to a fixed sized embedding while selectively focusing on the semantically descriptive parts of the motion.

One advantage of our approach is that, unlike DTW, we do not need any explicit synchronization or alignment of the motion patterns appearing on the two sequences, since motion patterns are implicitly and semantically matched via deep metric learning. In addition, our approach can naturally deal with varied size input thanks to the use of the recurrent model, while retaining the distinctive motion patterns by means of the attention mechanism. An example is shown in Fig. 1, comparing our similarity measure to DTW and L2. We validate the usefulness of our approach for the tasks of action retrieval and motion-based person identification on two publicly available benchmark datasets. The proposed experiments demonstrate significant improvements over conventional human motion similarity metrics.

2 Related Work

In recent literature, image-based deep metric learning has been extensively studied. However, just a few works focused on metric learning for time-series data, in particular human motion. Here, we first review metric learning approaches for human motion, then follow up with recent improvements in deep metric learning.

Metric Learning for Time Series and Human Motion. We first review metric learning approaches for time series, then focus only on works related on human motion analysis. Early works on metric learning for time series approaches measure the similarity in a two steps process [4, 9, 30]. First, the model determines the best alignment between two time series, then it computes the distance based on the aligned series. Usually, the model finds the best alignment by means of the DTW measure, first by considering all possible alignments, then ranking them based on hand-crafted local metric. These approaches have two main drawbacks: first, the model yields an $O(n^2)$ complexity; secondly, and most importantly, the local metric can hardly capture relationship in high dimensional data. In order to overcome these drawbacks, Mei *et al.* [25] propose to use LogDet divergence to learn a local metric that can capture the relationship in high dimensional data. Che *et al.* [5] overcome the hand crafted local metric problem by

using a feed-forward network to learn local similarities. Although the proposed approaches [5,25] learn to measure the similarity between two given time series at time t , the relationship between two time steps is discarded. Moreover, finding the best alignment requires to search for all possible alignments. To address these problems, recent work focused on determining a low dimensional embedding to measure the distance between time series. To this goal, Pei *et al.* [29] and Zheng *et al.* [46] used a Siamese network which learns from pairs of inputs. While Pei *et al.* [29] trained their network by minimizing the binary cross entropy in order to predict whether the two given time series belong to the same cluster or not, Zheng *et al.* [46] propose to minimize a loss function based on the Neighbourhood Component Analysis (NCA) [32]. The main drawback of these approaches is that the siamese architecture learns the embedding by considering only the relative distances between the provided input pairs.

As for metric learning for human motion analysis, they mostly focus on directly measuring the similarity between corresponding poses along the two sequences. Lopez *et al.* [22] proposed a model based on [10] to learn a distance metric for two given human poses, while aligning the motions via Hidden Markov Models (HMM) [11]. Chen *et al.* [6] proposed a semi-supervised learning approach built on a hand-crafted geometric pose feature and aligned via DTW. By considering both the pose similarity and the pose alignment in learning, Yin *et al.* [44] proposed to learn pose embeddings with an auto-encoder trained with an alignment constraint. Notably, this approach requires an initial alignment based on DTW. The main drawback of these approaches is that their accuracy relies heavily on the accurate motion alignment provided by HMM or DTW, which is computationally expensive to obtain and prone to fail in many cases. Moreover, since the learning process considers only single poses, they lack at capturing the semantics of the entire motion.

Recent Improvements in Deep Metric Learning. Metric learning with deep networks started with Siamese architectures that minimize the contrastive loss [7,15]. Schroff *et al.* [33] suggest using a triplet loss to learn the embeddings on facial recognition and verification, showing that it performs better than contrastive loss to learn features. Since they conduct hard-negative mining, when the training set and the number of different categories increase, searching for hard-negatives become computationally inefficient. Since then, research mostly focus on carefully constructing batches and using all samples in the batch. Song *et al.* [36] proposed the lifted loss for training, so to use all samples in a batch. In [35], they further developed the idea and propose an n -pair loss that uses all negative samples in a batch. Other triplet-based approaches are [26,40]. In [31], the authors show that minimizing the loss function computed on individual pairs or triplets does not necessarily enforce the network to learn features that represent contextual relations between clusters. Magnet Loss [31] address some of these issues by learning features that compare the distributions rather than the samples. Each cluster distribution is represented by the cluster centroid obtained via k -means algorithm. A shortcoming of this approach is that computing cluster centers requires to interrupt training, this slowing down the process.

Proxy-NCA [27] tackle this issue by designing a network architecture that learns the cluster centroids in an end-to-end fashion, thus avoiding interruptions during training. Both the Magnet Loss and the Proxy-NCA use the NCA [32] loss to compare the samples. Importantly, they both represent distributions with cluster centroids which do not convey sufficient contextual information of the actual categories, and require to set a pre-defined number of clusters. In contrast, we propose to use a loss function based on MMD [14], which relies on distribution moments that do not need to explicitly determine or learn cluster centroids.

3 Metric Learning on Human Motion

The objective is to learn an embedding for human motion sequences, such that the similarity metric between two human motion sequences $X := \{x_1, x_2, \dots, x_n\}$ and $Y := \{y_1, y_2, \dots, y_m\}$ (where x_t and y_t represent the poses at time t) can be expressed directly as the squared Euclidean distance in the embedding space. Mathematically, this can be written as

$$d(f(X), f(Y)) = \|f(X) - f(Y)\|^2 \quad (1)$$

where $f(\cdot)$ is the learned embedding function that maps a varied-length motion sequence to a point in a Euclidean space, and $d(\cdot, \cdot)$ is the squared Euclidean distance. The challenge of metric learning is to find a motion embedding function f such that the distance $d(f(X), f(Y))$ should be inversely proportional to the similarity of the two sequences X and Y . In this paper, we learn f by means of a deep learning model trained with a loss function (defined in Sect. 4) which is derived from the integration of MMD with a triplet learning paradigm. In addition, its architecture (described in Sect. 5) is based on an attentive recurrent neural network.

4 Loss Function

Following the standard deep metric learning approach, we model the embedding function f by minimizing the distance $d(f(X), f(Y))$ when X and Y belong to the same category, while maximizing it otherwise. A conventional way of learning f would be to train a network with the contrastive loss [7, 15].

$$\mathcal{L}_{\text{contrastive}} = (r) \frac{1}{2} d + (1 - r) \frac{1}{2} [\max(0, \alpha_{\text{margin}} - d)]^2 \quad (2)$$

where $r \in \{1, 0\}$ indicates whether X and Y are from the same category or not, and α_{margin} defines the margin between different category samples. During training, the contrastive loss penalizes those cases where different category samples are closer than α_{margin} and when the same category samples have a distance greater than zero. This equation shows that the contrastive loss only takes into account pairwise relationships between samples, thus only partially exploiting relative relationships among categories. Conversely, triplet learning

better exploit such relationships by taking into account three samples at the same time, where the first two are from the same category while the third is from a different one. Notably, it has been shown that exploiting relative relationships among categories play a fundamental role in terms of the quality of the learned embedding [33, 45]. The triplet loss enforces embedding samples from the same category with a given margin distance with respect to samples from a different category. If we denote the three human motion samples as X , X^+ and X^- , the commonly used ranking loss [34] takes the form of

$$\mathcal{L}_{\text{triplet}} = \max(0, \|f(X) - f(X^+)\|^2 - \|f(X) - f(X^-)\|^2 + \alpha_{\text{margin}}) \quad (3)$$

where X and X^+ represent the motion samples from the same category and X^- represents the sample from a different category. In literature X , X^+ , and X^- are often referred to as anchor, positive, and negative samples, respectively [31, 33, 35, 45].

However, one of the main issue with the triplet loss is the parameterization of α_{margin} . We can overcome this problem by using the Neighbourhood Components Analysis (NCA) [32]. Thus, we can write the loss function using NCA as

$$\mathcal{L}_{\text{NCA}} = \frac{\exp(-\|f(X) - f(X^+)\|^2)}{\sum_{X^- \in C} \exp(-\|f(X) - f(X^-)\|^2)} \quad (4)$$

where C represents all categories except for that of the positive sample.

In the ideal scenario, when iterating over triplets of samples, we expect that the samples from the same category will be grouped in the same cluster in the embedding space. However, it has been shown that most of the formed triplets are not informative and visiting all possible triplet combinations is infeasible. Therefore, the model will be trained with only a few informative triplets [31, 33, 35]. An intuitive solution can be formulated by selecting those negative samples that are hard to distinguish (hard negative mining), although searching for a hard negative sample in a motion sequence dataset is computationally expensive. Another issue linked with the use of triplet loss is that, during a single update, the positive and negative samples are evaluated only in terms of their relative position in the embedding; thus, samples can end up close to other categories [35]. We address the aforementioned issue by pushing/pulling the cluster distributions instead of pushing/pulling individual samples by means of a novel loss function, dubbed MMD-NCA and described next, that is based on the distribution differences of the categories.

4.1 MMD-NCA

Assuming that given two different distributions p and q , the general formulation of MMD measures the distance between p and q while taking the differences of the mean embeddings in Hilbert spaces, written as

$$\text{MMD}[k, p, q]^2 = \|\mu_q - \mu_p\|^2 = E_{x, x'}[k(x, x')] - 2E_{x, y}[k(x, y)] + E_{y, y'}[k(y, y')] \quad (5)$$

where x and x' are drawn IID from p while y and y' are drawn IID from q , and k represents the kernel function

$$k(x, x') = \sum_{q=1}^K k_{\sigma_q}(x, x') \quad (6)$$

where k_{σ_q} is a Gaussian kernel with bandwidth parameter σ_q , while K (number of kernels) is a hyperparameter. If we replace the expected values from the given samples, we obtain

$$\text{MMD}[k, X, Y]^2 = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m k(x_i, x'_j) - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, y_j) + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k(y_i, y'_j) \quad (7)$$

where $X := \{x_1, x_2, \dots, x_m\}$ is the sample set from p and $Y := \{y_1, y_2, \dots, y_n\}$ is the sample set from q . Hence, (7) allows us to measure the distance between the distribution of two sets.

We formulate our loss function in order to force the network to decrease the distance between the distribution of the anchor samples and that of the positive samples, while increasing the distance to the distribution of the negative samples.

Therefore, we can rewrite (4) for a given number N of anchor-positive sample pairs as $\{(X_1, X_1^+), (X_2, X_2^+), \dots, (X_N, X_N^+)\}$ and $N \times M$ negative samples from the M different categories $C = \{c_1, c_2, \dots, c_M\}$ as $\{X_{c_1, 1}^-, X_{c_1, 2}^-, \dots, X_{c_1, N}^-, \dots, X_{c_M, N}^-\}$; then,

$$\mathcal{L}_{\text{MMD-NCA}} = \frac{\exp(-\text{MMD}[k, f(X), f(X^+)])}{\sum_{j=1}^M \exp(-\text{MMD}[k, f(X), f(X_{c_j}^-)])} \quad (8)$$

where X and X^+ represent motion samples from the same category, while X_{c_j} represents samples from category $c_j \in C$. Our single update contains M different negative classes randomly sampled from the training data.

Since the proposed MMD-NCA loss minimizes the overlap between different category distributions in the embedding while keeping the samples from the same distribution as close as possible, we believe it is more effective for our task than the triplet loss. We demonstrate this quantitatively and qualitatively in Sect. 7.

5 Network Architecture

Our architecture is illustrated in Fig. 2. This model has two main parts: the bi-directional long short-term memory (BiLSTM) [16] and the self-attention mechanism. The reason for using the long short-term memory (LSTM) [16] is to overcome the vanishing gradient problem of the recurrent neural networks. In [12, 13], they show that LSTM can capture long term dependencies. In the next sections, we briefly describe the layer normalization mechanism and attention mechanism that used in our architecture.

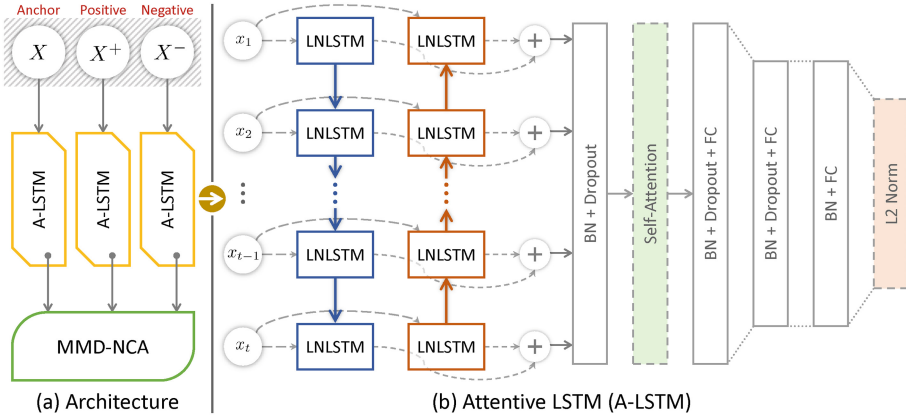


Fig. 2. (a) The proposed architecture for sequence distance learning. (b) The proposed attention-based model that uses layer normalization.

5.1 Layer Normalization

In [7, 26, 27, 36], they have shown that batch normalization plays a fundamental role on the triplet model’s accuracy. However, its straightforward application to LSTM architectures can decrease the accuracy of model [19]. Due to this, we used the layer normalized LSTM [3].

Suppose that n time steps of motion $X = (x_1, x_2, \dots, x_n)$ are given, then the layer normalized LSTM is described by

$$\mathbf{f}_t = \sigma(\mathbf{W}_{fh}\mathbf{h}_{t-1} + \mathbf{W}_{fx}x_t + \mathbf{b}_f) \tag{9}$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_{ih}\mathbf{h}_{t-1} + \mathbf{W}_{ix}x_t + \mathbf{b}_i) \tag{10}$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_{oh}\mathbf{h}_{t-1} + \mathbf{W}_{ox}x_t + \mathbf{b}_o) \tag{11}$$

$$\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_{ch}\mathbf{h}_{t-1} + \mathbf{W}_{cx}x_t + \mathbf{b}_c) \tag{12}$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t \tag{13}$$

$$\mathbf{m}_t = \frac{1}{H} \sum_j^H \mathbf{c}_t^j, \mathbf{v}_t = \sqrt{\frac{1}{H} \sum_j^H (\mathbf{c}_t^j - \mathbf{m}_t)^2} \tag{14}$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh\left(\frac{\gamma_t}{v_t} \odot (\mathbf{c}_t - \mathbf{m}_t) + \beta\right) \tag{15}$$

where c_{t-1} and h_{t-1} denotes the cell memory and cell state which comes from the previous time steps, x_t denotes the input human pose at time t . $\sigma(\cdot)$ and \odot represent the element-wise sigmoid function and multiplication respectively, and H denotes the number of hidden units in LSTM. The parameters $W_{\cdot,\cdot}$, γ and β are learned while γ and β has the same dimension of \mathbf{h}_t . Contrary to the standard LSTM, the hidden state \mathbf{h}_t is computed by normalizing the cell-memory \mathbf{c}_t .

5.2 Self-attention Mechanism

Intuitively, in a sequence of human motion, some poses are more informative than others. Therefore, we use the recently proposed self-attention mechanism [21] to assign a score for each pose in a motion sequence. Specifically, assuming that the sequence of states $S = \{h_1, h_2, \dots, h_n\}$ computed from a motion sequence X that consists of n time steps with (9) to (15), we can effectively compute the scores for each of them by

$$r = W_{s_2} \tanh(W_{s_1} S^T) \quad \text{and} \quad a_i = -\log \left(\frac{\exp(r_i)}{\sum_j \exp(r_j)} \right) \quad (16)$$

where r_i is i -th element of the r while W_{s_1} and W_{s_2} are the weight matrices in $R^{k \times l}$ and $R^{l \times 1}$, respectively. a_i is the assigned score i -th pose in the sequence of motion. Thus, the final embedding E can be computed by multiplying the scores $A = [a_1, a_2, \dots, a_n]$ and S , written as $E = AS$. Note that the final embedding size only depends on the number of hidden states in the LSTM and W_{s_2} . This allows us to encode the varying size LSTM outputs to a fixed sized output. More information about the self-attention mechanism can be found in [21].

6 Implementation Details

We use the TensorFlow framework [2] for all deep metric models that are described in this paper. Our model has three branches as shown in Fig. 2. Each branch consists of an attention based bidirectional layer normalized LSTM (LNLSTM) (see Sect. 5.1). Bidirectional LNLSTM follows a forward and backward passing of the given sequence of motion. We then denote $s_t = [s_{t,f}, s_{t,b}]$ such that $s_{t,f} = \overrightarrow{\text{LNLSTM}}(w_t, x_t)$ for $t \in [0, N]$ and $s_{t,b} = \overleftarrow{\text{LNLSTM}}(w_t, x_t)$ for $t \in [N, 0]$.

Given n time steps of a motion sequence X , we compute $S = (s_1, s_2, \dots, s_n)$ where s_t is the concatenated output of the backward and forward pass of the LNLSTM which has 128 hidden units. The bidirectional LSTM is followed by the dropout and the standard batch normalization. The output of the batch normalization layer is forwarded to the attention layer (see Sect. 5.2), which produces the fixed size of the output. The attention layer is followed by the structure: {FC(320,), dropout, BN, FC(320), BN, FC(128), BN, l_2 Norm}, where FC(m) means fully connected layer with m as the hidden units and BN means batch normalization. All the FC layers are followed by the rectified linear units except for the last FC layer. The self-attention mechanism is derived from the implementation of [21]. Here, the W_{s_1} and W_{s_2} parameters from (16) have the dimensionality of $R^{200 \times 10}$ and $R^{10 \times 1}$, respectively. We use the dropout rate of 0.5. The same dropout mask is used in all branches of the network in Fig. 2. In our model, all squared weight matrices are initialized with random orthogonal matrices while the others are initialized with uniform distribution with zero mean and 0.001 standard deviation. The parameters γ and β in (15) are initialized with zeros and ones, respectively.

Kernel Designs. The MMD-NCA loss function is implicitly associated with a family of characteristic kernels. Similar to the prior MMD papers [20, 38], we consider a mixture of K radial basis functions in (6). We fixed $K = 5$ and σ_q to be 1, 2, 4, 8, 16.

Training. Our single batch consists of randomly selected categories where each category has 25 samples. We selected 5 category as negative. Although the MMD [14] metric requires a high number of samples to understand the distribution moments, we found that 25 is sufficient for our tasks. Training each batch takes about 10s on a Titan X GPU. All the networks are trained with 5000 updates and they all converged before the end of training. During training, analogous to the curriculum learning, we start training on the samples without noise and then added Gaussian noise with zero mean and increasing standard deviation. We use stochastic gradient descent with the moment as an optimizer for all models. The momentum value is set to 0.9, and the learning rate started from 0.0001 with an exponential decay of 0.96 every 50 updates. We clip the whole gradients by their global norm to the range of -25 and 25 .

7 Experimental Results

We compare our MMD-NCA loss against the methods from DTW [42], MDDTW [25], CTW [47] and GDTW [48], as well as four state-of-the-art deep metric learning approaches: DCTW [41], triplet [33], triplet+GOR [45], and the N -Pairs deep metric loss [14]. Primarily, these methods are evaluated through action recognition task in Sect. 7.1. In order to look closely into the performance of this evaluation, we analyze the actions retrieved by the proposed method in the same section and the contribution of the self-attention mechanism from Sect. 5.2 into the algorithm in Sect. 7.3. Since one of the datasets [1] labeled the actions with their corresponding subjects, we also investigate the possibility of performing a person identification task wherein, instead of measuring the similarity of the pose, we intend to measure the similarity the actors themselves based on their movement. To have a fair comparison, we only used our attention based LSTM architecture for all methods and only changed the loss function except the DCTW [41]. Proposed loss function in DCTW [41] requires the two sequences, therefore we remove the attention layer and use only our LSTM model. Notably, all deep metric learning methods are evaluated and trained with the same data splits.

Performance Evaluation. We follow the same evaluation protocol as defined in [36, 45]. All models are evaluated for the clustering quality and false positive rate (FPR) on the same test set which consists of unseen motion categories. We compute the FPR for 90%, 80% and 70% true positive rates. In addition, we also use the Normalized Mutual Information measure (NMI) and F_1 score to measure the cluster quality where the NMI is the ratio between mutual information and sum of class and cluster labels entropies while the F_1 score is the harmonic mean of precision and recall.

Datasets and Pre-processing. We tested the models on two different datasets: (1) the CMU Graphics Lab motion capture database (CMU mocap) [1]; and, (2) the Human3.6M dataset [17]. The former [1] contains 144 different subjects where each subject performs natural motions such as *walking*, *dancing* and *jumping*. Their data is recorded with the mocap system and the poses are represented with 38 joints in 3D space. Six joints are excluded because they have no movement. We align the poses with respect to the torso and, to avoid the gimbal lock effect, the poses are expressed in the exponential map [39]. Although the original data runs at 120 Hz with different lengths of motion sequences, we down-sampled the data to 30 Hz during training and testing.

Furthermore, the Human3.6M dataset [17] consists of 15 different actions and each action was performed by seven different professional actors. The actions are mostly selected from daily activities such as *walking*, *smoking*, *engaging in a discussion*, *taking pictures* and *talking on the phone*. We process the dataset in the same way as the same as CMU mocap.

7.1 Action Recognition

In this experiment, we tested our model on both the CMU mocap [1] and the Human3.6M [17] datasets for unseen motion categories. We categorize the CMU mocap dataset into 38 different motion categories where the motion sequences which contain more than one category are excluded. Among them, we selected 19 categories for training and 19 categories for testing. For the Human3.6M [17],

Table 1. False positive rate of action recognition for CMU mocap and Human3.6M datasets.

	CMU			Human3.6M		
	FPR-90	FPR-80	FPR-70	FPR-90	FPR-80	FPR-70
DTW [42]	47.98	42.92	37.62	49.64	47.96	44.38
MDDTW [25]	44.60	39.07	34.04	49.72	45.87	44.51
CTW [47]	46.02	40.96	39.11	47.63	43.10	42.18
GDTW [48]	45.61	39.95	35.24	46.06	42.72	40.04
DCTW [41]	40.56	38.83	26.95	41.39	39.18	36.71
Triplet [33]	39.72	33.82	28.77	42.78	40.15	36.01
Triplet + GOR [45]	40.32	33.97	27.78	42.03	37.61	33.95
N-Pair [35]	40.11	32.35	26.16	40.46	39.56	36.52
MMD-NCA (<i>Ours</i>)	32.66	25.66	20.29	38.42	36.54	33.13
– without Attention	41.22	35.36	30.04	45.03	42.07	41.01
– without LN	37.27	30.21	27.95	44.25	41.69	38.09
– Linear Kernel	39.80	33.92	29.00	46.35	41.68	37.69
– Polynomial Kernel	36.80	30.35	24.98	43.60	40.03	35.62

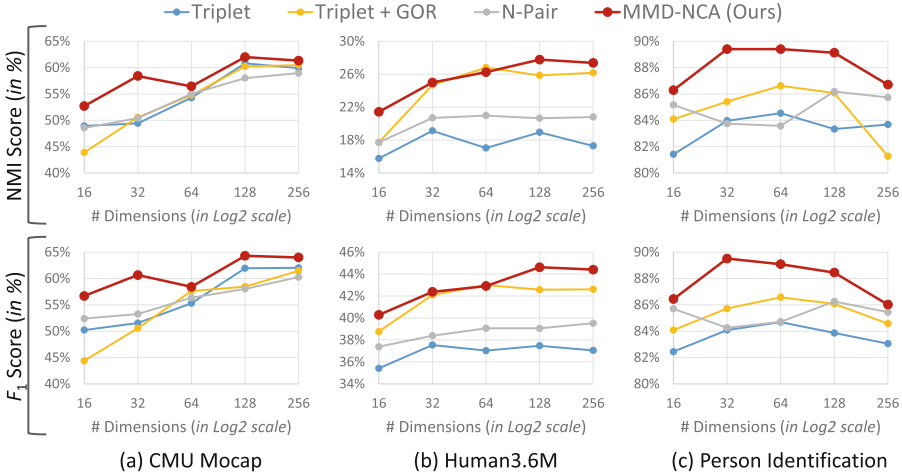


Fig. 3. NMI and F_1 score for the action recognition task using the (a) CMU Mocap and (b) Human3.6M datasets; and, (b) for person identification task.

we used all the given categories, and selected 8 categories for training and 7 categories for testing.

Although our model allows us to train with varying sizes of motion sequence, we train with a fixed size, since varying sizes slow down the training process. We divided the motion sequences into 90 consecutive frames (*i.e.* approximately 3s) and leave a gap of 30 frames. However, at test time, we divided the motion sequences only if it is longer than 5s by leaving a 1-s gap; otherwise, we keep the original motion sequence. We found this processing effective since we observe that, in sequence of motions longer than 5s, the subjects usually repeat their action. We also consider training without clipping but it was not possible with available the GPU resources.

False Positive Rate. The FPR at different percentages on CMU mocap and Human3.6M are reported in Table 1. With a true positive rate of 70%, the learning approaches [33, 35, 41, 45] including our approach achieve up to 17% improvement in FPR relative to DTW [42], MDDTW [25], CTW [47] and GDTW [48]. Moreover, our approach further improves the results up to 6% and 0.8% for CMU mocap and Human3.6m datasets, respectively, against the state-of-the-art deep learning approaches [33, 35, 41, 45].

NMI and F_1 Score. Figure 3(a) plots the NMI and F_1 score with varying size of embedding for the CMU mocap dataset. In both the NMI and F_1 metrics, our approach produces the best clusters at all the embedding sizes. Compared to other methods, the proposed approach is less sensitive to the changes of the embedding size. Moreover, Fig. 3(b) illustrates the NMI and F_1 score on



Fig. 4. Comparison of cartwheel motion query on the CMU mocap dataset between our approach and DTW [42]. The motion in the first row is query and the rest are four nearest neighbors for each method, which are sorted by the distance.

Human3.6M dataset where we observe similar performance as the CMU mocap dataset and acquire the best results.

Action Retrieval. In order to investigate further, we query a specific motion from the CMU mocap test set, and compare the closest action sequences that our approach and DTW [42] retrieve based on their respective similarity measure. In Fig. 4, we demonstrate this task as we query the challenging cartwheel motion (see first row). Our approach successfully retrieves the semantically similar motions sequences, despite the high variation on the length of sequences. On the other hand, DTW [42] fails to match the query to the dataset because the distinctive pose appears on a small portion of the sequence. This implies that the large portion, where the actor stands, dominates the similarity measure. Note that we do not have the same problem due to the self-attention mechanism from Sect. 5.2 (see Sect. 7.3 for the evaluation).

7.2 Person Identification

Since the CMU mocap dataset also includes the specific subject associated to each motion, we explore the potential application of person identification. In contrast to the action recognition and action retrieval from Sect. 7.1 where the similarity measure is calculated based on the motion category, this task tries to measure the similarity with respect the actor. In this experiment, we construct

Table 2. False positive rate of person identification for CMU mocap dataset.

	FPR-95	FPR-90	FPR-85	FPR-80	FPR-75	FPR-70
DTW [42]	46.22	43.19	38.70	32.36	27.61	22.85
MDDTW [25]	49.67	45.89	40.36	35.46	31.69	28.44
CTW [47]	45.23	40.14	35.69	29.50	25.91	20.35
GDTW [48]	44.65	40.54	35.03	28.07	24.31	19.32
DCTW [41]	32.45	20.24	18.15	15.91	13.78	10.31
Triplet [33]	22.58	18.13	11.30	9.63	8.36	6.51
Triplet + GOR [45]	28.37	16.69	10.27	8.64	7.28	4.38
N-Pair [35]	22.84	15.31	8.94	5.69	4.82	4.56
MMD-NCA (<i>Ours</i>)	19.31	10.42	8.26	5.62	3.91	2.55
– without Attention	36.10	26.15	22.48	20.94	19.21	16.78
– without LN	26.63	18.43	12.81	10.27	8.58	7.36
– Linear Kernel	35.75	30.97	25.93	15.13	11.93	10.42
– Polynomial Kernel	27.25	21.18	17.91	10.93	8.97	5.93

the training and test set in the same way as Sect. 7.1. We included the subjects which have more than three motion sequences, which resulted in 68 subjects. Among them, we selected 39 subjects for training and the rest of the 29 subjects for testing.

Table 2 shows the FPR for the person identification task for varying percentages of true positive rate with embedding size of 64. Here, all deep metric learning approaches including our work significantly improve the accuracy against the DTW, MDDTW, CTW and GDTW. Overall, our method outperforms all the approaches for all FPR with a 20% improvement against DTW [42], MDDTW [25], CTW [47] and GDTW [48], and a 2% improvement compared to the state-of-the-art deep learning approaches [33, 35, 41, 45]. Moreover, when we evaluate the NMI and the F_1 score for the clustering quality in different embedding sizes, Fig. 3(c) demonstrates that our approach obtains the state-of-the-art results with a significant margin.

7.3 Attention Visualization

The objective of the self-attention mechanism from Sect. 5.2 is to focus on the poses which are the most informative about the semantics of the motion sequence. Thus, we expect our attention mechanism to focus on the descriptive poses in the motion, which allows the model to learn more expressive embeddings. Based on the peaks of A which is composed of a_i from (16), we illustrate this behavior in Fig. 5 where the first two rows belong to the basketball sequence while the third belong to the bending sequence. Notably, all the sequences have different lengths.




Fig. 5. Attention visualization: the poses in red show where the model mostly focused its attention. Specifically, we mark as red those frames associated with each column-wise global maximum in A , together with the previous and next 2 frames. For visualization purposes, the sequences are subsampled by a factor of 4.

Despite the variations in the length of the motion, the model focuses when the actor throws the ball which is the most informative part of the motion for Fig. 5(a-b); while, for the bending motion in Fig. 5(c), it also focuses on the distinctive regions of the motion sequence. Therefore, this figure illustrates that the self-attention mechanism successfully focuses on the most informative part of the sequence. This implies that the model discards the non-informative parts in order to embed long motion sequences to a low dimensional space without losing the semantic information.

8 Ablation Study

We evaluate our architecture with different configurations to better appreciate each of our contributions separately. All models are trained with MMD-NCA loss and with an embedding of size 128. Tables 1 and 2 show the effect of the layer normalization [3], the self-attention mechanism [21] and the kernel selection in terms of FPR. We use the same architecture for linear, polynomial, and MMD-NCA and only change the kernel function in (6). Notably, the removal of the self-attention mechanism yields the biggest drop in NMI and F_1 on all the datasets. In addition, both the layer normalization and the self-attention improve the resulting FPR by 7% and 10%, respectively. In terms of kernel selection, the results show that selecting the kernel which takes into account higher moments yields better results. Comparing the two tasks, the person identification is the one that benefits from our architecture the most.

9 Conclusion

In this paper, we propose a novel loss function and network architecture to measure the similarity of two motion sequences. Experimental results on the CMU mocap [1] and Human3.6M [17] datasets show that our approach obtains state-of-the-art results. We also have shown that metric learning approaches based on deep learning can improve the results up to 20% against metrics commonly used

for similarity among human motion sequences. As future work, we plan to generalize the proposed MMD-NCA framework to time-series, as well as investigate different types of kernels.

References

1. Carnegie mellon university - CMU graphics lab - motion capture library (2010). <http://mocap.cs.cmu.edu/>. Accessed 03 Nov 2018
2. Abadi, M., et al.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). <https://www.tensorflow.org/>. Software available from [tensorflow.org](https://www.tensorflow.org/)
3. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. CoRR abs/1607.06450 (2016). <http://arxiv.org/abs/1607.06450>
4. Berndt, D.J., Clifford, J.: Using dynamic time warping to find patterns in time series. In: KDD Workshop, Seattle, WA, vol. 10, pp. 359–370 (1994)
5. Che, Z., He, X., Xu, K., Liu, Y.: DECADE: a deep metric learning model for multivariate time series (2017)
6. Chen, C., Zhuang, Y., Nie, F., Yang, Y., Wu, F., Xiao, J.: Learning a 3D human pose distance metric from geometric pose descriptor. *IEEE Trans. Vis. Comput. Graph.* **17**(11), 1676–1689 (2011)
7. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2005, vol. 1, pp. 539–546. IEEE (2005)
8. Chu, X., Yang, W., Ouyang, W., Ma, C., Yuille, A.L., Wang, X.: Multi-context attention for human pose estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017
9. Cuturi, M., Vert, J.P., Birkenes, O., Matsui, T.: A kernel for time series based on global alignments. In: 2007 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2007, vol. 2, pp. II–413. IEEE (2007)
10. Davis, J.V., Kulis, B., Jain, P., Sra, S., Dhillon, I.S.: Information-theoretic metric learning. In: Proceedings of the 24th International Conference on Machine Learning, pp. 209–216. ACM (2007)
11. Eddy, S.R.: Hidden markov models. *Curr. Opin. Struct. Biol.* **6**(3), 361–365 (1996)
12. Graves, A., Mohamed, A.R., Hinton, G.: Speech recognition with deep recurrent neural networks. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6645–6649. IEEE (2013)
13. Greff, K., Srivastava, R.K., Koutník, J., Steunebrink, B.R., Schmidhuber, J.: LSTM: a search space odyssey. *IEEE Trans. Neural Netw. Learn. Syst.* **28**(10), 2222–2232 (2017)
14. Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *J. Mach. Learn. Res.* **13**, 723–773 (2012)
15. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2, pp. 1735–1742. IEEE (2006)
16. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
17. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6M: large scale datasets and predictive methods for 3D human sensing in natural environments. *IEEE Trans. Patt. Anal. Mach. Intell.* **36**(7), 1325–1339 (2014)

18. Keogh, E.J., Pazzani, M.J.: Derivative dynamic time warping. In: Proceedings of the 2001 SIAM International Conference on Data Mining, pp. 1–11. SIAM (2001)
19. Laurent, C., Pereyra, G., Brakel, P., Zhang, Y., Bengio, Y.: Batch normalized recurrent neural networks. In: 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2657–2661. IEEE (2016)
20. Li, Y., Swersky, K., Zemel, R.: Generative moment matching networks. In: Proceedings of the 32nd International Conference on Machine Learning (ICML 2015), pp. 1718–1727 (2015)
21. Lin, Z., et al.: A structured self-attentive sentence embedding. In: Proceedings of International Conference on Learning Representations (ICLR) (2017)
22. López-Méndez, A., Gall, J., Casas, J.R., Van Gool, L.J.: Metric learning from poses for temporal clustering of human motion. In: BMVC, pp. 1–12 (2012)
23. Martinez, J., Black, M.J., Romero, J.: On human motion prediction using recurrent neural networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017
24. Mehta, D., et al.: VNect: real-time 3D human pose estimation with a single RGB camera. *ACM Trans. Graph. (TOG)* **36**(4), 44 (2017)
25. Mei, J., Liu, M., Wang, Y.F., Gao, H.: Learning a mahalanobis distance-based dynamic time warping measure for multivariate time series classification. *IEEE Trans. Cybern.* **46**(6), 1363–1374 (2016)
26. Mishchuk, A., Mishkin, D., Radenovic, F., Matas, J.: Working hard to know your neighbor’s margins: local descriptor learning loss. In: Proceedings Conference on Neural Information Processing Systems (NIPS), December 2017
27. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: The IEEE International Conference on Computer Vision (ICCV), October 2017
28. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9912, pp. 483–499. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_29
29. Pei, W., Tax, D.M., van der Maaten, L.: Modeling time series similarity with siamese recurrent networks. *CoRR abs/1603.04713* (2016)
30. Ratanamahatana, C.A., Keogh, E.: Making time-series classification more accurate using learned constraints. In: SIAM (2004)
31. Rippel, O., Paluri, M., Dollar, P., Bourdev, L.: Metric learning with adaptive density discrimination. In: International Conference on Learning Representations (2016)
32. Roweis, S., Hinton, G., Salakhutdinov, R.: Neighbourhood component analysis. *Adv. Neural Inf. Process. Syst. (NIPS)* **17**, 513–520 (2004)
33. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: a unified embedding for face recognition and clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 815–823 (2015)
34. Schultz, M., Joachims, T.: Learning a distance metric from relative comparisons. In: *Advances in Neural Information Processing Systems*, pp. 41–48 (2004)
35. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: *Advances in Neural Information Processing Systems*, pp. 1857–1865 (2016)
36. Song, H.O., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4004–4012. IEEE (2016)
37. Sun, X., Shang, J., Liang, S., Wei, Y.: Compositional human pose regression. In: *The IEEE International Conference on Computer Vision (ICCV)*, vol. 2 (2017)

38. Sutherland, D.J., et al.: Generative models and model criticism via optimized maximum mean discrepancy. In: Proceedings of the 32nd International Conference on Machine Learning (ICML 2017) (2017)
39. Taylor, G.W., Hinton, G.E., Roweis, S.T.: Modeling human motion using binary latent variables. In: Advances in Neural Information Processing Systems, pp. 1345–1352 (2007)
40. Tian, B.F.Y., Wu, F.: L2-Net: deep learning of discriminative patch descriptor in Euclidean space. In: Conference on Computer Vision and Pattern Recognition (CVPR), vol. 2 (2017)
41. Trigeorgis, G., Nicolaou, M.A., Schuller, B.W., Zafeiriou, S.: Deep canonical time warping for simultaneous alignment and representation learning of sequences. *IEEE Trans. Patt. Anal. Mach. Intell.* **5**, 1128–1138 (2018)
42. Vintsyuk, T.K.: Speech discrimination by dynamic programming. *Cybernetics* **4**(1), 52–57 (1968)
43. Yang, W., Li, S., Ouyang, W., Li, H., Wang, X.: Learning feature pyramids for human pose estimation. In: The IEEE International Conference on Computer Vision (ICCV), October 2017
44. Yin, X., Chen, Q.: Deep metric learning autoencoder for nonlinear temporal alignment of human motion. In: 2016 IEEE International Conference on Robotics and Automation (ICRA), pp. 2160–2166. IEEE (2016)
45. Zhang, X., Yu, F.X., Kumar, S., Chang, S.F.: Learning spread-out local feature descriptors. In: The IEEE International Conference on Computer Vision (ICCV), October 2017
46. Zheng, Y., Liu, Q., Chen, E., Zhao, J.L., He, L., Lv, G.: Convolutional nonlinear neighbourhood components analysis for time series classification. In: Cao, T., Lim, E.-P., Zhou, Z.-H., Ho, T.-B., Cheung, D., Motoda, H. (eds.) PAKDD 2015. LNCS (LNAI), vol. 9078, pp. 534–546. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18032-8_42
47. Zhou, F., Torre, F.: Canonical time warping for alignment of human behavior. In: Advances in Neural Information Processing Systems, pp. 2286–2294 (2009)
48. Zhou, F., De la Torre, F.: Generalized canonical time warping. *IEEE Trans. Patt. Anal. Mach. Intell.* **38**(2), 279–294 (2016)