

Tensor Representations via Kernel Linearization for Action Recognition from 3D Skeletons

Piotr Koniusz^{1,3}(✉), Anoop Cherian^{2,3}, and Fatih Porikli^{1,2,3}

¹ NICTA/Data61/CSIRO, Canberra, Australia

{piotr.koniusz,fatih.porikli}@data61.csiro.au

² Australian Centre for Robotic Vision, Canberra, Australia

anoop.cherian@data61.csiro.au

³ Australian National University, Canberra, Australia

{piotr.koniusz,anoop.cherian,fatih.porikli}@anu.edu.au

Abstract. In this paper, we explore tensor representations that can compactly capture higher-order relationships between skeleton joints for 3D action recognition. We first define RBF kernels on 3D joint sequences, which are then linearized to form kernel descriptors. The higher-order outer-products of these kernel descriptors form our tensor representations. We present two different kernels for action recognition, namely (i) a *sequence compatibility kernel* that captures the spatio-temporal compatibility of joints in one sequence against those in the other, and (ii) a *dynamics compatibility kernel* that explicitly models the action dynamics of a sequence. Tensors formed from these kernels are then used to train an SVM. We present experiments on several benchmark datasets and demonstrate state of the art results, substantiating the effectiveness of our representations.

Keywords: Kernel descriptors · Skeleton action recognition · Higher-order tensors

1 Introduction

Human action recognition is a central problem in computer vision with potential impact in surveillance, human-robot interaction, elderly assistance systems, and gaming, to name a few. While there have been significant advancements in this area over the past few years, action recognition in unconstrained settings still remains a challenge. There have been research to simplify the problem from using RGB cameras to more sophisticated sensors such as Microsoft Kinect that can localize human body-parts and produce moving 3D skeletons [1]; these skeletons are then used for recognition. Unfortunately, these skeletons are often noisy due to the difficulty in localizing body-parts, self-occlusions, and sensor range

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46493-0_3](https://doi.org/10.1007/978-3-319-46493-0_3)) contains supplementary material, which is available to authorized users.

errors; thus necessitating higher-order reasoning on these 3D skeletons for action recognition.

There have been several approaches suggested in the recent past to improve recognition performance of actions from such noisy skeletons. These approaches can be mainly divided into two perspectives, namely (i) generative models that assume the skeleton points are produced by a latent dynamic model [2] corrupted by noise and (ii) discriminative approaches that generate compact representations of sequences on which classifiers are trained [3]. Due to the huge configuration space of 3D actions and the unavailability of sufficient training data, discriminative approaches have been the trend in the recent years for this problem. In this line of research, the main idea has been to compactly represent the spatio-temporal evolution of 3D skeletons, and later train classifiers on these representations to recognize the actions. Fortunately, there is a definitive structure to motions of 3D joints relative to each other due to the connectivity and length constraints of body-parts. Such constraints have been used to model actions; examples include Lie Algebra [4], positive definite matrices [5,6], using a torus manifold [7], Hanklet representations [8], among several others. While modeling actions with explicit manifold assumptions can be useful, it is computationally expensive.

In this paper, we present a novel methodology for action representation from 3D skeleton points that avoids any manifold assumptions on the data representation, instead captures the higher-order statistics of how the body-joints relate to each other in a given action sequence. To this end, our scheme combines positive definite kernels and higher-order tensors, with the goal to obtain rich and compact representations. Our scheme benefits from using non-linear kernels such as radial basis functions (RBF) and it can also capture higher-order data statistics and the complexity of action dynamics.

We present two such kernel-tensor representations for the task. Our first representation *sequence compatibility kernel* (SCK), captures the spatio-temporal compatibility of body-joints between two sequences. To this end, we present an RBF kernel formulation that jointly captures the spatial and temporal similarity of each body-pose (normalized with respect to the hip position) in a sequence against those in another. We show that tensors generated from third-order outer-products of the linearizations of these kernels can be a simple yet powerful representation capturing higher-order co-occurrence statistics of body-parts and yield high classification confidences.

Our second representation, termed *dynamics compatibility kernel* (DCK) aims at representing spatio-temporal dynamics of each sequence explicitly. We present a novel RBF kernel formulation that captures the similarity between a pair of body-poses in a given sequence explicitly, and then compare it against such body-pose pairs in other sequences. As it might appear, such spatio-temporal modeling could be expensive due to the volumetric nature of space and time. However, we show that using an appropriate kernel model can shrink the time-related variable in a small constant size representation after kernel linearization. With this approach, we can model both spatial and temporal variations in the form of co-occurrences which could otherwise have been prohibitive.

We further show through experiments that the above two representations in fact capture complementary statistics regarding the actions, and combining them leads to significant benefits. We present experiments on three standard datasets for the task, namely (i) UTKinect-Actions [9], (ii) Florence3D-Actions [10], and (iii) MSR-Action3D [11] datasets and demonstrate state-of-the-art accuracy.

To summarize, the main contributions of this paper are (i) introduction of sequence and the dynamics compatibility kernels for capturing spatio-temporal evolution of body-joints for 3D skeleton based action sequences, (ii) derivations of linearization of these kernels, and (iii) their tensor reformulations. We review the related literature next.

2 Related Work

The problem of skeleton based action recognition has received significant attention over the past decades. Interested readers may refer to useful surveys [3] on the topic. In the sequel, we will review some of the more recent related approaches to the problem.

In this paper, we focus on action recognition datasets that represent a human body as an articulated set of connected body-joints that evolve in time [12]. A temporal evolution of the human skeleton is very informative for action recognition as shown by Johansson in his seminal experiment involving the moving lights display [13]. At the simplest level, the human body can be represented as a set of 3D points corresponding to body-joints such as elbow, wrist, knee, ankle, *etc.* Action dynamics has been modeled using the motion of such 3D points in [14, 15], using joint orientations with respect to a reference axis [16] and even relative body-joint positions [17, 18]. In contrast, we focus on representing these 3D body-joints by kernels whose linearization results in higher-order tensors capturing complex statistics. Noteworthy are also parts-based approaches that additionally consider the connected body segments [4, 19–21].

Our work also differs from previous works in the way it handles the temporal domain. 3D joint locations are modeled as temporal hierarchy of coefficients in [14]. Pairwise relative positions of joints were modeled in [17] and combined with a hierarchy of Fourier coefficients to capture temporal evolution of actions. Moreover, this approach uses multiple kernel learning to select discriminative joint combinations. In [18], the relative joint positions and their temporal displacements are modeled with respect to the initial frame. In [4], the displacements and angles between the body parts are represented as a collection of matrices belonging to the special Euclidean group $SE(3)$. Temporal domain is handled by the discrete time warping and Fourier temporal pyramid matching on a sequence of such matrices. In contrast, we model temporal domain with a single RBF kernel providing invariance to local temporal shifts and avoid expensive techniques such as time warping and multiple-kernel learning.

Our scheme also differs from prior works such as kernel descriptors [22] that aggregate orientations of gradients for recognition. Their approach exploits sums over the product of at most two RBF kernels handling two cues *e.g.*, gradient orientations and spatial locations, which are later linearized by Kernel PCA

and Nyström techniques. Similarly, convolutional kernel networks [23] consider stacked layers of a variant of kernel descriptors [22]. Kernel trick was utilized for action recognition in kernelized covariances [24] which are obtained in Nyström-like process. A time series kernel [25] between auto-correlation matrices is proposed to capture spatio-temporal auto-correlations. In contrast, our scheme allows sums over several multiplicative and additive RBF kernels, thus, it allows handling multiple input cues to build a complex representation. We show how to capture higher-order statistics by linearizing a polynomial kernel and avoid evaluating costly kernels directly in contrast to kernel trick.

Third-order tensors have been found to be useful for several other vision tasks. For example, in [26], spatio-temporal third-order tensors on videos is proposed for action analysis, non-negative tensor factorization is used for image denoising in [27], tensor textures are proposed for texture rendering in [28], and higher order tensors are used for face recognition in [29]. A survey of multi-linear algebraic methods for tensor subspace learning and applications is available in [30]. These applications use a single tensor, while our goal is to use the tensors as data descriptors similar to [31–34] for image recognition tasks. However, in contrast to these similar methods, we explore the possibility of using third-order representations for 3D action recognition, which poses a different set of challenges.

3 Preliminaries

In this section, we review our notations and the necessary background on shift-invariant kernels and their linearizations, which will be useful for deriving kernels on 3D skeletons for action recognition.

3.1 Tensor Notations

Let $\mathcal{V} \in \mathbb{R}^{d_1 \times d_2 \times d_3}$ denote a third-order tensor. Using Matlab style notation, we refer to the p -th slice of this tensor as $\mathcal{V}_{:, :, p}$, which is a $d_1 \times d_2$ matrix. For a matrix $\mathbf{V} \in \mathbb{R}^{d_1 \times d_2}$ and a vector $\mathbf{v} \in \mathbb{R}^{d_3}$, the notation $\mathcal{V} = \mathbf{V} \uparrow \otimes \mathbf{v}$ produces a tensor $\mathcal{V} \in \mathbb{R}^{d_1 \times d_2 \times d_3}$ where the p -th slice of \mathcal{V} is given by $\mathbf{V} v_p$, v_p being the p -th dimension of \mathbf{v} . Symmetric third-order tensors of rank one are formed by the outer product of a vector $\mathbf{v} \in \mathbb{R}^d$ in modes two and three. That is, a rank-one $\mathcal{V} \in \mathbb{R}^{d \times d \times d}$ is obtained from \mathbf{v} as $\mathcal{V} = (\uparrow \otimes_3 \mathbf{v} \triangleq (\mathbf{v} \mathbf{v}^T) \uparrow \otimes \mathbf{v})$. Concatenation of n tensors in mode k is denoted as $[\mathcal{V}_i]_{i \in \mathcal{I}_n}^{\oplus k}$, where \mathcal{I}_n is an index sequence $1, 2, \dots, n$. The Frobenius norm of tensor is given by $\|\mathcal{V}\|_F = \sqrt{\sum_{i,j,k} \mathcal{V}_{ijk}^2}$, where \mathcal{V}_{ijk} represents the ijk -th element of \mathcal{V} . Similarly, the inner-product between two tensors \mathcal{X} and \mathcal{Y} is given by $\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{ijk} \mathcal{X}_{ijk} \mathcal{Y}_{ijk}$.

3.2 Kernel Linearization

Let $G_\sigma(\mathbf{u} - \bar{\mathbf{u}}) = \exp(-\|\mathbf{u} - \bar{\mathbf{u}}\|_2^2 / 2\sigma^2)$ denote a standard Gaussian RBF kernel centered at $\bar{\mathbf{u}}$ and having a bandwidth σ . Kernel linearization refers to rewriting

this G_σ as an inner-product of two infinite-dimensional feature maps. To obtain these maps, we use a fast approximation method based on probability product kernels [35]. Specifically, we employ the inner product of d' -dimensional isotropic Gaussians given $u, u' \in \mathbb{R}^{d'}$. The resulting approximation can be written as:

$$G_\sigma(\mathbf{u} - \bar{\mathbf{u}}) = \left(\frac{2}{\pi\sigma^2} \right)^{\frac{d'}{2}} \int_{\zeta \in \mathbb{R}^{d'}} G_{\sigma/\sqrt{2}}(\mathbf{u} - \zeta) G_{\sigma/\sqrt{2}}(\bar{\mathbf{u}} - \zeta) d\zeta. \quad (1)$$

Equation (1) is then approximated by replacing the integral with the sum over Z pivots ζ_1, \dots, ζ_Z , thus writing a feature map ϕ as:

$$\begin{aligned} \phi(\mathbf{u}) &= \left[G_{\sigma/\sqrt{2}}(\mathbf{u} - \zeta_1), \dots, G_{\sigma/\sqrt{2}}(\mathbf{u} - \zeta_Z) \right]^T, \\ \text{and } G_\sigma(\mathbf{u} - \bar{\mathbf{u}}) &\approx \langle \sqrt{c}\phi(\mathbf{u}), \sqrt{c}\phi(\bar{\mathbf{u}}) \rangle, \end{aligned} \quad (2)$$

where c represents a constant. We refer to (3) as the linearization of the RBF kernel.

4 Proposed Approach

In this section, we first formulate the problem of action recognition from 3D skeleton sequences, which precedes an exposition of our two kernel formulations for describing the actions, followed by their tensor reformulations through kernel linearization.

4.1 Problem Formulation

Suppose we are given a set of 3D human pose skeleton sequences, each pose consisting of J body-keypoints. Further, to simplify our notations, we assume each sequence consists of N skeletons, one per frame¹. Mathematically, we can define such a pose sequence Π as:

$$\Pi = \{ \mathbf{x}_{is} \in \mathbb{R}^3, i \in \mathcal{I}_J, s \in \mathcal{I}_N \}. \quad (4)$$

Further, let each such sequence Π be associated with one of K action class labels $\ell \in \mathcal{I}_K$. Our goal is to use the skeleton sequence Π and generate an action descriptor for this sequence that can be used in a classifier for recognizing the action class. In the following, we will present two such action descriptors, namely (i) sequence compatibility kernel and (ii) dynamics compatibility kernel, which are formulated using the ideas of kernel linearization and tensor algebra. We present both these kernel formulations next.

¹ We assume that all sequences have N frames for simplification of presentation. Our formulations are equally applicable to sequences of arbitrary lengths *e.g.*, M and N . Therefore, we apply in practice $G_{\sigma_3}(\frac{s}{M} - \frac{t}{N})$ in Eq. (5).

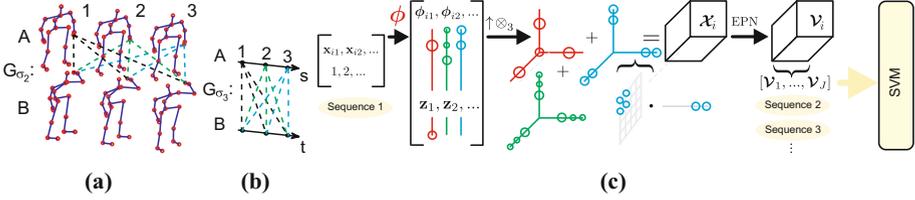


Fig. 1. Figures (a) and (b) show how SCK works – kernel G_{σ_2} compares exhaustively *e.g.* hand-related joint i for every frame in sequence A with every frame in sequence B . Kernel G_{σ_3} compares exhaustively the frame indexes. Figure (c) shows this burden is avoided by linearization – third-order statistics on feature maps $\phi(\mathbf{x}_{i,s})$ and $\mathbf{z}(s)$ for joint i are captured in tensor \mathcal{X}_i and whitened by EPN to obtain \mathcal{V}_i which are concatenated over $i=1, \dots, J$ to represent a sequence.

4.2 Sequence Compatibility Kernel

As alluded to earlier, the main idea of this kernel is to measure the compatibility between two action sequences in terms of the similarity between their skeletons and their temporal order. To this end, we assume each skeleton is centralized with respect to one of the body-joints (say, hip). Suppose we are given two such sequences Π_A and Π_B , each with J joints, and N frames. Further, let $\mathbf{x}_{i,s} \in \mathbb{R}^3$ and $\mathbf{y}_{j,t} \in \mathbb{R}^3$ correspond to the body-joint coordinates of Π_A and Π_B , respectively. We define our *sequence compatibility kernel* (SCK) between Π_A and Π_B as (See footnote 1):

$$K_S(\Pi_A, \Pi_B) = \frac{1}{A} \sum_{(i,s) \in \mathcal{J}} \sum_{(j,t) \in \mathcal{J}} G_{\sigma_1}(i-j) \left(\beta_1 G_{\sigma_2}(\mathbf{x}_{i,s} - \mathbf{y}_{j,t}) + \beta_2 G_{\sigma_3} \left(\frac{s-t}{N} \right) \right)^r, \quad (5)$$

where A is a normalization constant and $\mathcal{J} = \mathcal{I}_J \times \mathcal{I}_N$. As is clear, this kernel involves three different compatibility subkernels, namely (i) G_{σ_1} , that captures the compatibility between joint-types i and j , (ii) G_{σ_2} , capturing the compatibility between joint locations \mathbf{x} and \mathbf{y} , and (iii) G_{σ_3} , measuring the temporal alignment of two poses in the sequences. We also introduce weighting factors $\beta_1, \beta_2 \geq 0$ that adjusts the importance of the body-joint compatibility against the temporal alignment, where $\beta_1 + \beta_2 = 1$. Figures 1a and b illustrate how this kernel works. It might come as a surprise, why we need the kernel G_{σ_1} . Note that our skeletons may be noisy and there is a possibility that some of the keypoints are detected incorrectly (for example, elbows and wrists). Thus, this kernel allows incorporating some degree of uncertainty to the alignment of such joints. To simplify our formulations, in this paper, we will assume that such errors are absent from our skeletons, and thus $G_{\sigma_1}(i-j) = \delta(i-j)$. Further, the standard deviations σ_2 and σ_3 control the joint-coordinate selectivity and temporal shift-invariance respectively. That is, for $\sigma_3 \rightarrow 0$, two sequences will have to match perfectly in the temporal sense. For $\sigma_3 \rightarrow \infty$, the algorithm is invariant to any permutations of the frames. As will be clear in the sequel, the parameter r determines the order statistics of the kernel (we use $r = 3$).

Next, we present linearization of our kernel using the method proposed in Sect. 3.2 and Eq. (3) so that kernel $G_{\sigma_2}(\mathbf{x}-\mathbf{y}) \approx \phi(\mathbf{x})^T \phi(\mathbf{y})$ (see footnote²) while $G_{\sigma_3}(\frac{s-t}{N}) \approx \mathbf{z}(s/N)^T \mathbf{z}(t/N)$. With these approximations and simplification to G_{σ_1} we described above, we can rewrite our sequence compatibility kernel as:

$$K_S(\Pi_A, \Pi_B) = \frac{1}{\Lambda} \sum_{i \in \mathcal{I}_J} \sum_{s \in \mathcal{I}_N} \sum_{t \in \mathcal{I}_N} \left(\begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{x}_{is}), \text{ (see footnote 2)} \\ \sqrt{\beta_2} \mathbf{z}(s/N) \end{bmatrix}^T \cdot \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{y}_{it}) \\ \sqrt{\beta_2} \mathbf{z}(t/N) \end{bmatrix} \right)^r \quad (6)$$

$$= \frac{1}{\Lambda} \sum_{i \in \mathcal{I}_J} \sum_{s \in \mathcal{I}_N} \sum_{t \in \mathcal{I}_N} \left\langle \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{x}_{is}) \\ \sqrt{\beta_2} \mathbf{z}(s/N) \end{bmatrix}, \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{y}_{it}) \\ \sqrt{\beta_2} \mathbf{z}(t/N) \end{bmatrix} \right\rangle \quad (7)$$

$$= \sum_{i \in \mathcal{I}_J} \left\langle \frac{1}{\sqrt{\Lambda}} \sum_{s \in \mathcal{I}_N} \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{x}_{is}) \\ \sqrt{\beta_2} \mathbf{z}(s/N) \end{bmatrix}, \frac{1}{\sqrt{\Lambda}} \sum_{t \in \mathcal{I}_N} \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{y}_{it}) \\ \sqrt{\beta_2} \mathbf{z}(t/N) \end{bmatrix} \right\rangle. \quad (8)$$

As is clear, (8) expresses $K_S(\Pi_A, \Pi_B)$ as a sum of inner-products on third-order tensors ($r = 3$). This is illustrated by Fig. 1c. While, using the dot-product as the inner-product is a possibility, there are much richer alternatives for tensors of order $r >= 2$ that can exploit their structure or manipulate higher-order statistics inherent in them, thus leading to better representations. An example of such a commonly encountered property is the so-called *burstiness* [36], which is the property that a given feature appears more often in a sequence than a statistically independent model would predict. A robust sequence representation should be invariant to the length of actions *e.g.*, a prolonged *hand waving* represents the same action as a short *hand wave*. The same is true for short versus repeated *head nodding*. Eigenvalue Power Normalization (EPN) [32] is known to suppress burstiness. It acts on higher-order statistics illustrated in Fig. 1c. Incorporating EPN, we generalize (8) as:

$$K_S^*(\Pi_A, \Pi_B) = \sum_{i \in \mathcal{I}_J} \left\langle \mathcal{G} \left(\frac{1}{\sqrt{\Lambda}} \sum_{s \in \mathcal{I}_N} \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{x}_{is}) \\ \sqrt{\beta_2} \mathbf{z}(s/N) \end{bmatrix} \right), \mathcal{G} \left(\frac{1}{\sqrt{\Lambda}} \sum_{t \in \mathcal{I}_N} \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{y}_{it}) \\ \sqrt{\beta_2} \mathbf{z}(t/N) \end{bmatrix} \right) \right\rangle, \quad (9)$$

where the operator \mathcal{G} performs EPN by applying power normalization to the spectrum of the third-order tensor (by taking the higher-order SVD). Note that in general $K_S^*(\Pi_A, \Pi_B) \not\approx K_S(\Pi_A, \Pi_B)$ as \mathcal{G} is intended to manipulate the spectrum of \mathcal{X} . The final representation, for instance for a sequence Π_A , takes the following form:

$$\mathbf{v}_i = \mathcal{G}(\mathcal{X}_i), \text{ where } \mathcal{X}_i = \frac{1}{\sqrt{\Lambda}} \sum_{s \in \mathcal{I}_N} \uparrow \otimes_r \begin{bmatrix} \sqrt{\beta_1} \phi(\mathbf{x}_{is}) \\ \sqrt{\beta_2} \mathbf{z}(s/N) \end{bmatrix}. \quad (10)$$

² In practice, we use $G'_{\sigma_2}(\mathbf{x}-\mathbf{y}) = G_{\sigma_2}(x^{(x)}-y^{(x)}) + G_{\sigma_2}(x^{(y)}-y^{(y)}) + G_{\sigma_2}(x^{(z)}-y^{(z)})$ so the kernel $G'_{\sigma_2}(\mathbf{x}-\mathbf{y}) \approx [\phi(x^{(x)}); \phi(x^{(y)}); \phi(x^{(z)})]^T [\phi(y^{(x)}); \phi(y^{(y)}); \phi(y^{(z)})]$ but for simplicity we write $G_{\sigma_2}(\mathbf{x}-\mathbf{y}) \approx \phi(\mathbf{x})^T \phi(\mathbf{y})$. Note that (x) , (y) , (z) are the spatial xyz-components of joints.

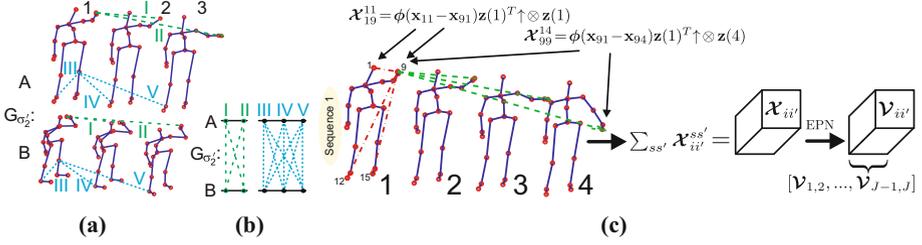


Fig. 2. Figure (a) shows that kernel G_{σ_2} in DCK captures spatio-temporal dynamics by measuring displacement vectors from any given body-joint to remaining joints spatially- and temporally-wise (*i.e.* see dashed lines). Figure (b) shows that comparisons performed by G_{σ_2} for any selected two joints are performed all-against-all temporally-wise which is computationally expensive. Figure (c) shows the encoding steps in the proposed linearization which overcome this burden.

We can further replace the summation over the body-joint indexes in (9) by concatenating \mathcal{V}_i in (10) along the fourth tensor mode, thus defining $\mathcal{V} = [\mathcal{V}_i]_{i \in \mathcal{I}_J}^{\oplus 4}$. Suppose \mathcal{V}_A and \mathcal{V}_B are the corresponding fourth order tensors for Π_A and Π_B respectively. Then, we obtain:

$$K_S^*(\Pi_A, \Pi_B) = \langle \mathcal{V}_A, \mathcal{V}_B \rangle. \quad (11)$$

Note that the tensors \mathcal{X} have the following properties: (i) super-symmetry $\mathcal{X}_{i,j,k} = \mathcal{X}_{\pi(i,j,k)}$ for indexes i, j, k and their permutation given by π , $\forall \pi$, and (ii) positive semi-definiteness of every slice, that is, $\mathcal{X}_{::,s} \in \mathcal{S}_+^d$, for $s \in \mathcal{I}_d$. Therefore, we need to use only the upper-simplex of the tensor which consists of $\binom{d+r-1}{r}$ coefficients (which is the total size of our final representation) rather than d^r , where d is the side-dimension of \mathcal{X} *i.e.*, $d = 3Z_2 + Z_3$ (see footnote 2), and Z_2 and Z_3 are the numbers of pivots used in the approximation of G_{σ_2} (see footnote 2) and G_{σ_3} respectively. As we want to preserve the above listed properties in tensors \mathcal{V} , we employ slice-wise EPN which is induced by the Power-Euclidean distance and involves rising matrices to a power γ . Finally, we re-stack these slices along the third mode as:

$$\mathcal{G}(\mathcal{X}) = [\mathcal{X}_{::,s}^\gamma]_{s \in \mathcal{I}_d}^{\oplus 3}, \text{ for } 0 < \gamma \leq 1. \quad (12)$$

This $\mathcal{G}(\mathcal{X})$ forms our tensor representation for the action sequence.

4.3 Dynamics Compatibility Kernel

The SCK kernel that we described above captures the inter-sequence alignment, while the intra-sequence spatio-temporal dynamics is lost. In order to capture these temporal dynamics, we propose a novel dynamics compatibility kernel (DCK). To this end, we use the absolute coordinates of the joints in our kernel.

Using the notations from the earlier section, for two action sequences Π_A and Π_B , we define this kernel as:

$$K_D(\Pi_A, \Pi_B) = \frac{1}{\Lambda} \sum_{\substack{(i,s) \in \mathcal{I}, \\ (i',s') \in \mathcal{I}, \\ i \neq i', s \neq s'}} \sum_{\substack{(j,t) \in \mathcal{J}, \\ (j',t') \in \mathcal{J}, \\ j \neq j', t \neq t'}} G'_{\sigma'_1}(i-j, i'-j') G_{\sigma'_2}((\mathbf{x}_{is} - \mathbf{x}_{i's'}) - (\mathbf{y}_{jt} - \mathbf{y}_{j't'})) \cdot G'_{\sigma'_3}\left(\frac{s-t}{N}, \frac{s'-t'}{N}\right) G'_{\sigma'_4}(s-s', t-t'), \quad (13)$$

where $G'_\sigma(\boldsymbol{\alpha}, \boldsymbol{\beta}) = G_\sigma(\boldsymbol{\alpha})G_\sigma(\boldsymbol{\beta})$. In comparison to the SCK kernel in (5), the DCK kernel uses the intra-sequence joint differences, thus capturing the dynamics. This dynamics is then compared to those in the other sequences. Figures 2a-c, depict schematically how this kernel captures co-occurrences. As in SCK, the first kernel $G'_{\sigma'_1}$ is used to capture sensor uncertainty in body-keypoint detection, and is assumed to be a delta function in this paper. The second kernel $G_{\sigma'_2}$ models the spatio-temporal co-occurrences of the body-joints. Temporal alignment kernels expressed as $G_{\sigma'_3}$ encode the temporal start and end-points from (s, s') and (t, t') . Finally, $G'_{\sigma'_4}$ limits contributions of dynamics between temporal points if they are distant from each other, *i.e.* if $s' \gg s$ or $t' \gg t$ and σ'_4 is small. Furthermore, similar to SCK, the standard deviations σ'_2 and σ'_3 control the selectivity over spatio-temporal dynamics of body-joints and their temporal shift-invariance for the start and end points, respectively. As discussed for SCK, the practical extensions described by the footnotes 1 and 2 apply to DCK as well.

As in the previous section, we employ linearization to this kernel. Following the derivations described above, it can be shown that the linearized kernel has the following form (see [37] or supplementary material for details):

$$K_D(\Pi_A, \Pi_B) = \sum_{\substack{i \in \mathcal{I}_j, \\ i' \in \mathcal{I}_j, \\ i' \neq i}} \left\langle \frac{1}{\sqrt{\Lambda}} \sum_{\substack{s \in \mathcal{I}_N, \\ s' \in \mathcal{I}_N, \\ s' \neq s}} G_{\sigma'_4}(s-s') \left(\boldsymbol{\phi}(\mathbf{x}_{is} - \mathbf{x}_{i's'}) \cdot \mathbf{z}\left(\frac{s}{N}\right)^T \right) \uparrow \otimes \mathbf{z}\left(\frac{s'}{N}\right), \right. \\ \left. \frac{1}{\sqrt{\Lambda}} \sum_{\substack{t \in \mathcal{I}_N, \\ t' \in \mathcal{I}_N, \\ t' \neq t}} G_{\sigma'_4}(t-t') \left(\boldsymbol{\phi}(\mathbf{y}_{it} - \mathbf{y}_{i't'}) \cdot \mathbf{z}\left(\frac{t}{N}\right)^T \right) \uparrow \otimes \mathbf{z}\left(\frac{t'}{N}\right) \right\rangle. \quad (14)$$

Equation (14) expresses $K_D(\Pi_A, \Pi_B)$ as a sum over inner-products on third-order non-symmetric tensors of third-order (c.f. Sect. 4.2 where the proposed kernel results in an inner-product between super-symmetric tensors). However, we can decompose each of these tensors with a variant of EPN which involves Higher Order Singular Value Decomposition (HOSVD) into factors stored in the so-called core tensor and equalize the contributions of these factors. Intuitively, this would prevent bursts in the statistically captured spatio-temporal co-occurrence dynamics of actions. For example, consider that a long *hand wave* versus a short one yield different temporal statistics, that is, the prolonged action

results in bursts. However, the representation for action recognition should be invariant to such cases. As in the previous section, we introduce a non-linear operator \mathcal{G} into Eq. (14) which will handle this. Our final representation, for example, for sequence Π_A can be expressed as:

$$\mathcal{V}_{ii'} = \mathcal{G}(\mathcal{X}_{ii'}), \text{ and } \mathcal{X}_{ii'} = \frac{1}{\sqrt{\Lambda}} \sum_{\substack{s \in \mathcal{I}_N, \\ s' \in \mathcal{I}_N: \\ s' \neq s}} G_{\sigma'_4}(s - s') \left(\phi(\mathbf{x}_{is} - \mathbf{x}_{i's'}) \cdot \mathbf{z} \left(\frac{s}{N} \right)^T \right) \uparrow \otimes \mathbf{z} \left(\frac{s'}{N} \right), \quad (15)$$

where the summation over the pairs of body-joint indexes in (14) becomes equivalent to the concatenation of $\mathcal{V}_{ii'}$ from (15) along the fourth mode such that we obtain tensor representations $[\mathcal{V}_{ii'}]_{i>i': i, i' \in \mathcal{I}_J}^{\oplus 4}$ for sequence Π_A and $[\bar{\mathcal{V}}_{ii'}]_{i>i': i, i' \in \mathcal{I}_J}^{\oplus 4}$ for sequence Π_B . The dot-product can be now applied between these representations for comparing them. For the operator \mathcal{G} , we choose HOSVD-based tensor whitening as proposed in [32]. However, they work with the super-symmetric tensors, such as the one we proposed in Sect. 4.2. We work with a general non-symmetric case in (15) and use the following operator \mathcal{G} :

$$(\mathcal{E}; \mathbf{A}_1, \dots, \mathbf{A}_r) = \text{HOSVD}(\mathcal{X}) \quad (16)$$

$$\hat{\mathcal{E}} = \text{Sgn}(\mathcal{E}) |\mathcal{E}|^\gamma \quad (17)$$

$$\hat{\mathcal{V}} = ((\hat{\mathcal{E}} \otimes_1 \mathbf{A}_1) \dots) \otimes_r \mathbf{A}_r \quad (18)$$

$$\mathcal{G}(\mathcal{X}) = \text{Sgn}(\hat{\mathcal{V}}) |\hat{\mathcal{V}}|^{\gamma^*} \quad (19)$$

In the above equations, we distinguish the core tensor \mathcal{E} and its power normalized variants $\hat{\mathcal{E}}$ with factors that are being evened out by rising to the power $0 < \gamma \leq 1$, eigenvalue matrices $\mathbf{A}_1, \dots, \mathbf{A}_r$ and operation \otimes_r which represents a so-called tensor-product in mode r . We refer the reader to paper [32] for the detailed description of the above steps.

5 Computational Complexity

Non-linearized SCK with kernel SVM has complexity $\mathcal{O}(JN^2T^\rho)$ given J body joints, N frames per sequence, T sequences, and $2 < \rho < 3$ which concerns complexity of kernel SVM. Linearized SCK with linear SVM takes $\mathcal{O}(JNTZ_*^r)$ for a total of Z_* pivots and tensor order $r = 3$. Note that $N^2T^\rho \gg NTZ_*^r$. For $N = 50$ and $Z_* = 20$, this is $3.5 \times$ (or $32 \times$) faster than the exact kernel for $T = 557$ (or $T = 5000$) used in our experiments. Non-linearized DCK with kernel SVM has complexity $\mathcal{O}(J^2N^4T^\rho)$ while linearized DCK takes $\mathcal{O}(J^2N^2TZ^3)$ for Z pivots per kernel, *e.g.* $Z = Z_2 = Z_3$ given $G_{\sigma'_2}$ and $G_{\sigma'_3}$. As $N^4T^\rho \gg N^2TZ^3$, the linearization is $11000 \times$ faster than the exact kernel, for say $Z = 5$. Note that EPN incurs negligible cost (see [37] for details).

6 Experiments

In this section, we present experiments using our models on three benchmark 3D skeleton based action recognition datasets, namely (i) the UTKinect-Action [9], (ii) Florence3D-Action [10], and (iii) MSR-Action3D [11]. We also present experiments evaluating the influence of the choice of various hyper-parameters, such as the number of pivots Z used for linearizing the body-joint and temporal kernels, the impact of Eigenvalue Power Normalization, and factor equalization.

6.1 Datasets

UTKinect-Action [9] dataset consists of 10 actions performed twice by 10 different subjects, and has 199 action sequences. The dataset provides 3D coordinate annotations of 20 body-joints for every frame. The dataset was captured with a stationary Kinect sensor and contains significant viewpoint and intra-class variations.

Florence3D-Action [10] dataset consists of 9 actions performed two to three times by 10 different subjects. It comprises 215 action sequences. 3D coordinate annotations of 15 body-joints are provided for every frame. This dataset was also captured with a Kinect sensor and contains significant intra-class variations *i.e.*, the same action may be articulated with the left or right hand. Moreover, some actions such as *drinking*, *performing a phone call*, etc., can be visually ambiguous.

MSR-Action3D [11] dataset is comprised from 20 actions performed two to three times by 10 different subjects. Overall, it consists of 557 action sequences. 3D coordinates of 20 body-joints are provided. This dataset was captured using a Kinect-like depth sensor. It exhibits strong inter-class similarity.

In all experiments we follow the standard protocols for these datasets. We use the cross-subject test setting, in which half of the subjects are used for training and the remaining half for testing. Similarly, we divide the training set into two halves for purpose of training-validation. Additionally, we use two protocols for MSR-Action3D according to approaches [11,17], where the latter protocol uses three subsets grouping related actions together.

6.2 Experimental Setup

For the sequence compatibility kernel, we first normalized all body-keypoints with respect to the hip joints across frames, as indicated in Sect. 4.2. Moreover, lengths of all body-parts are normalized with respect to a reference skeleton. This setup follows the pre-processing suggested in [4]. For our dynamics compatibility kernel, we use unnormalized body-joints and assume that the displacements of body-joint coordinates across frames capture their temporal evolution implicitly.

Sequence Compatibility Kernel. In this section, we first present experiments evaluating the influence of parameters σ_2 and σ_3 of kernels G_{σ_2} and G_{σ_3} which

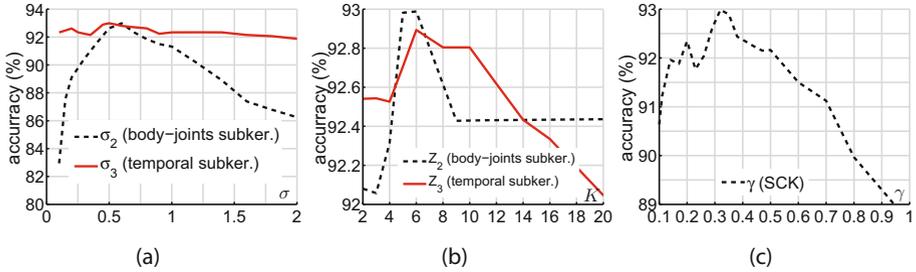


Fig. 3. Figure (a) illustrates the classification accuracy on Florence3d-Action for the sequence compatibility kernel when varying radii σ_2 (body-joints subkernel) and σ_3 (temporal subkernel). Figure (b) evaluates behavior of SCK w.r.t. the number of pivots Z_2 and Z_3 . Figure (c) demonstrates effectiveness of our slice-wise Eigenvalue Power Normalization in tackling burstiness by varying parameter γ .

control the degree of selectivity for the 3D body-joints and temporal shift invariance, respectively. See Sect. 4.2 for a full definition of these parameters.

Furthermore, recall that the kernels G_{σ_2} and G_{σ_3} are approximated via linearizations according to Eqs. (1) and (3). The quality of these approximations and the size of our final tensor representations depend on the number of pivots Z_2 and Z_3 chosen. In our experiments, the pivots ζ are spaced uniformly within interval $[-1; 1]$ and $[0; 1]$ for kernels G_{σ_2} and G_{σ_3} respectively.

Figures 3a and b present the results of this experiment on the Florence3D-Action dataset – these are the results presented on the test set as we have also observed exactly the same trends on the validation set.

Figure 3a shows that the body-joint compatibility subkernel G_{σ_2} requires a choice of σ_2 which is not too strict as the specific body-joints (*e.g.*, elbow) would be expected to repeat across sequences in the exactly same position. On the one hand, very small σ_2 leads to poor generalization. On the other hand, very large σ_2 allows big displacements of the corresponding body-joints between sequences which results in poor discriminative power of this kernel. Furthermore, Fig. 3a demonstrates that the range of σ_3 for the temporal subkernel for which we obtain very good performance is large, however, as σ_3 becomes very small or very large, extreme temporal selectivity or full temporal invariance, respectively, result in a loss of performance. For instance, $\sigma_3 = 4$ results in 91 % accuracy only.

In Fig. 3b, we show the performance of our SCK kernel with respect to the number of pivots used for linearization. For the body-joint compatibility subkernel G_{σ_2} , we see that $Z_2 = 5$ pivots are sufficient to obtain good performance of 92.98 % accuracy. We have observed that this is consistent with the results on the validation set. Using more pivots, say $Z_2 = 20$, deteriorates the results slightly, suggesting overfitting. We make similar observations for the temporal subkernel G_{σ_3} which demonstrates good performance for as few as $Z_3 = 2$ pivots. Such a small number of pivots suggests that linearizing 1D variables and generating higher-order co-occurrences, as described in Sect. 4.2, is a simple, robust, and effective linearization strategy.

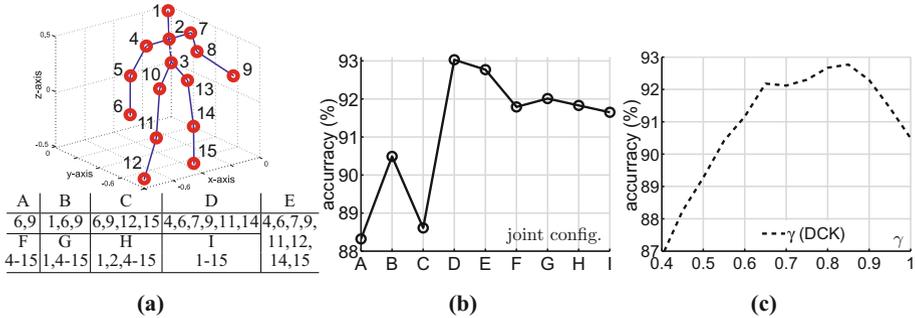


Fig. 4. Figure (a) enumerates the body-joints in the Florence3D-Action dataset. The table below lists subsets A-I of the body-joints used to build representations evaluated in Fig. (b), which demonstrates the performance of our dynamics compatibility kernel w.r.t. these subsets. Figure (c) demonstrates effectiveness of equalizing the factors in non-symmetric tensor representation by HOSVD Eigenvalue Power Normalization by varying γ .

Further, Fig. 3c demonstrates the effectiveness of our slice-wise Eigenvalue Power Normalization (EPN) described in Eq. (12). When $\gamma = 1$, the EPN functionality is absent. This results in a drop of performance from 92.98% to 88.7% accuracy. This demonstrates that statistically unpredictable bursts of actions described by the body-joints, such as long versus short *hand waving*, are indeed undesirable. It is clear that in such cases, EPN is very effective, as in practice it considers correlated bursts, *e.g.* co-occurring *hand wave* and associated with it elbow and neck motion. For more details behind this concept, see [32]. For our further experiments, we choose $\sigma_2 = 0.6$, $\sigma_3 = 0.5$, $Z_2 = 5$, $Z_3 = 6$, and $\gamma = 0.36$, as dictated by cross-validation.

Dynamics Compatibility Kernel. In this section, we evaluate the influence of choosing parameters for the DCK kernel. Our experiments are based on the Florence3D-Action dataset. We present the scores on the test set as the results on the validation set match these closely. As this kernel considers all spatio-temporal co-occurrences of body-joints, we first evaluate the impact of the joint subsets we select for generating this representation as not all body-joints need to be used for describing actions.

Figure 4a enumerates the body-joints that describe every 3D human skeleton on the Florence3D-Action dataset whilst the table underneath lists the proposed body-joint subsets A-I which we use for computations of DCK. In Fig. 4b, we plot the performance of our DCK kernel for each subset. The plot shows that using two body-joints associated with the hands from Configuration-A in the DCK kernel construction, we attain 88.32% accuracy which highlights the informativeness of temporal dynamics. For Configuration-D, which includes six body-joints such as the knees, elbows and hands, our performance reaches 93.03%. This suggests that some not-selected for this configuration body-joints may be noisy and therefore detrimental to classification.

As configuration Configuration-E includes eight body-joints such as the feet, knees, elbows and hands, we choose it for our further experiments as it represents a reasonable trade-off between performance and size of representations. This configuration scores 92.77 % accuracy. We see that if we utilize all the body-joints according to Configuration-I, performance of 91.65 % accuracy is still somewhat lower compared to 93.03 % accuracy for Configuration-D highlighting again the issue of noisy body-joints.

In Fig. 4c, we show the performance of our DCK kernel when HOSVD factors underlying our non-symmetric tensors are equalized by varying the EPN parameter γ . For $\gamma = 1$, HOSVD EPN is absent which leads to 90.49 % accuracy only. For the optimal value of $\gamma = 0.85$, the accuracy rises to 92.77 %. This again demonstrates the presence of the burstiness effect in temporal representations.

Comparison to the State of the Art. In this section, we compare the performance of our representations against the best performing methods on the three datasets. Along with comparing SCK and DCK, we will also explore the complementarity of these representations in capturing the action dynamics by combining them.

On the Florence3D-Action dataset, we present our best results in Table 1a. Note that the model parameters for the evaluation was selected by cross-validation. Linearizing a sequence compatibility kernel using these parameters resulted in a tensor representation of size 26,565 dimensions³, and producing an accuracy of 92.98 % accuracy. As for the dynamics compatibility kernel (DCK), our model selected Configuration-E (described in Fig. 4a) resulting in a representation of dimensionality 16,920 and achieved a performance of 92 %. However, somewhat better results were attained by Configuration-D, namely 92.27 % accuracy for size of 9,450. Combining both SCK representation with DCK in Configuration-E results in an accuracy of 95.23 %. This constitutes a 4.5 % improvement over the state of the art on this dataset as listed in Table 1a and demonstrates the complementary nature of SCK and DCK. To the best of our knowledge, this is the highest performance attained on this dataset.

Action recognition results on the UTKinect-Action dataset are presented in Table 1b. For our experiments on this dataset, we kept all the parameters the same as those we used on the Florence3D dataset (described above). On this

Table 1. Evaluations of SCK and DCK and comparisons to the state-of-the-art results on (a) the Florence3D-Action and (b) UTKinect-Action dataset.

	SCK	DCK		SCK+DCK		SCK	DCK	SCK+DCK	
accuracy	92.98 %	93.03 %	92.77 %	95.23 %	accuracy	96.08 %	97.5 %	98.2 %	
size	26,565	9,450	16,920	43,485	size	40,480	16,920	57,400	
Bag-of-Poses	82.00 % [10]		$SE(3)$	90.88 % [4]	3D joints. hist.	90.92 % [9]		$SE(3)$	97.08 % [4]
	(a)					(b)			

³ Note that this is the length of a vector per sequence after unfolding our tensor representation and removing duplicate coefficients from the symmetries in the tensor.

Table 2. Results on SCK and DCK and comparisons to the state of the art on MSR-Action3D.

	SCK	DCK	SCK+DCK	accuracy, protocol [17]	accuracy, protocol [11]
acc., prot. [17]	90.72 %	86.30 %	91.45 %	Actionlets 88.20 % [17]	R. Forests 90.90 % [38]
acc., prot. [11]	93.52 %	91.71 %	93.96 %	$SE(3)$ 89.48 % [4]	$SE(3)$ 92.46 % [4]
size	40,480	16,920	57,400	Kin. desc. 91.07 % [39]	

dataset, both SCK and DCK representations yield 96.08 % and 97.5 % accuracy, respectively. Combining SCK and DCK yields 98.2 % accuracy outperforming marginally a more complex approach described in [4] which uses Lie group algebra on $SE(3)$ matrix descriptors and requires practical extensions such as discrete time warping and Fourier temporal pyramids for attaining this performance, which we avoid completely.

In Table 2, we present our results on the MSR-Action3D dataset. Again, we kept all the model parameters the same as those used on the Florence3D dataset. Conforming to prior literature, we use two evaluation protocols on this dataset, namely (i) the protocol described in actionlets [17], for which the authors utilize the entire dataset with its 20 classes during the training and evaluation, and (ii) approach of [11], for which the authors divide the data into three subsets and report the average in classification accuracy over these subsets. The SCK representation results in the state-of-the-art accuracy of 90.72 % and 93.52 % for the two evaluation protocols, respectively. Combining SCK with DCK outperforms other approaches listed in the table and yields 91.45 % and 93.96 % accuracy for the two protocols, respectively.

Processing Time. For SCK and DCK, processing a single sequence with unoptimized MATLAB code on a single core i5 takes 0.2s and 1.2s, respectively. Training on full MSR Action3D with the SCK and DCK takes about 13 min. In comparison, extracting $SE(3)$ features [4] takes 5.3s per sequence, processing on the full MSR Action3D dataset takes ~ 50 min. and with post-processing (time warping, Fourier pyramids, etc.) it goes to about 72 min. Therefore, SCK and DCK is about 5.4 \times faster.

7 Conclusions

We have presented two kernel-based tensor representations for action recognition from 3D skeletons, namely the sequence compatibility kernel (SCK) and dynamics compatibility kernel (DCK). SCK captures the higher-order correlations between 3D coordinates of the body-joints and their temporal variations, and factors out the need for expensive operations such as Fourier temporal pyramid matching or dynamic time warping, commonly used for generating sequence-level action representations. Further, our DCK kernel captures the action dynamics by modeling the spatio-temporal co-occurrences of the body-joints. This tensor representation also factors out the temporal variable, whose length depends

on each sequence. Our experiments substantiate the effectiveness of our representations, demonstrating state-of-the-art performance on three challenging action recognition datasets.

Acknowledgements. AC is funded by the Australian Research Council Centre of Excellence for Robotic Vision (project number CE140100016).

References

1. 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. *Commun. ACM* **56**, 116–124 (2013)
2. Turaga, P., Chellappa, R.: Locally time-invariant models of human activities using trajectories on the grassmannian. In: *CVPR* (2009)
3. Presti, L.L., La Cascia, M.: 3D skeleton-based human action classification: a survey. *Pattern Recogn.* **53**, 130–147 (2015)
4. Vemulapalli, R., Arrate, F., Chellappa, R.: Human action recognition by representing 3D skeletons as points in a Lie Group. In: *CVPR*, pp. 588–595 (2014)
5. Harandi, M., Salzmann, M., Porikli, F.: Bregman divergences for infinite dimensional covariance matrices. In: *CVPR* (2014)
6. Hussein, M.E., Torki, M., Gowayyed, M.A., El-Saban, M.: Human action recognition using a temporal hierarchy of covariance descriptors on 3D joint locations. In: *IJCAI* (2013)
7. Elgammal, A., Lee, C.S.: Tracking people on a torus. *PAMI* **31**, 520–538 (2009)
8. Li, B., Camps, O.I., Szaier, M.: Cross-view activity recognition using hanklets. In: *CVPR* (2012)
9. Xia, L., Chen, C.C., Aggarwal, J.K.: View invariant human action recognition using histograms of 3D joints. In: *CVPR Workshops*, pp. 20–27(2012)
10. Seidenari, L., Varano, V., Berretti, S., Bimbo, A.D., Pala, P.: Recognizing actions from depth cameras as weakly aligned multi-part bag-of-poses. In: *CVPR Workshop*, June 2013
11. Li, W., Zhang, Z., Liu, Z.: Action recognition based on a bag of 3D points. In: *CVPR Workshop*, pp. 9–14 (2010)
12. Zatsiorsky, V.M.: *Kinematic of Human Motion*. Human Kinetics Publishers, Champaign (1997)
13. Johansson, G.: Visual perception of biological motion and a model for its analysis. *Percept. Psychophysics* **14**(2), 201–211 (1973)
14. Hussein, M.E., Torki, M., Gowayyed, M., El-Saban, M.: Human action recognition using a temporal hierarchy of covariance descriptors on 3D joint locations. In: *IJCAI* 2466–2472 (2013)
15. Lv, F., Nevatia, R.: Recognition and segmentation of 3-D human action using HMM and multi-class AdaBoost. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *ECCV 2006*. LNCS, vol. 3954, pp. 359–372. Springer, Heidelberg (2006). doi:[10.1007/11744085_28](https://doi.org/10.1007/11744085_28)
16. Parameswaran, V., Chellappa, R.: View invariance for human action recognition. *IJCV* **66**(1), 83–101 (2006)
17. Wu, Y., Liu, Z., Wu, Y., Yuan, J.: Mining actionlet ensemble for action recognition with depth cameras. In: *CVPR*, pp. 1290–1297 (2012)

18. Yang, X., Tian, Y.: Effective 3D action recognition using eigenjoints. *J. Vis. Comun. Image Represent.* **25**(1), 2–11 (2014)
19. Yacoob, Y., Black, M.J.: Parameterized modeling and recognition of activities. In: *ICCV*, pp. 120–128 (1998)
20. Ohn-Bar, E., Trivedi, M.M.: Joint angles similarities and HOG² for action recognition. In: *CVPR Workshop* (2013)
21. Offi, F., Chaudhry, R., Kurillo, G., Vidal, R., Bajcsy, R.: Sequence of the most informative joints (SMIJ). *J. Vis. Comun. Image Represent.* **25**(1), 24–38 (2014)
22. Bo, L., Lai, K., Ren, X., Fox, D.: Object recognition with hierarchical kernel descriptors. In: *CVPR* (2011)
23. Mairal, J., Koniusz, P., Harchaoui, Z., Schmid, C.: Convolutional kernel networks. In: *NIPS* (2014)
24. Cavazza, J., Zunino, A., Biagio, M.S., Vittorio, M.: Kernelized covariance for action recognition. *CoRR abs/1604.06582* (2016)
25. Gaidon, A., Harchoui, Z., Schmid, C.: A time series kernel for action recognition. *BMVC* **63**(1-63), 11 (2011)
26. Kim, T.K., Wong, K.Y.K., Cipolla, R.: Tensor canonical correlation analysis for action classification. In: *CVPR* (2007)
27. Shashua, A., Hazan, T.: Non-negative tensor factorization with applications to statistics and computer vision. In: *ICML* (2005)
28. Vasilescu, M.A., Terzopoulos, D.: Tensortextures: multilinear image-based rendering. *ACM Trans. Graph.* **23**(3), 336–342 (2004)
29. Vasilescu, M.A.O., Terzopoulos, D.: Multilinear analysis of image ensembles: tensorfaces. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002*. LNCS, vol. 2350, pp. 447–460. Springer, Heidelberg (2002). doi:[10.1007/3-540-47969-4-30](https://doi.org/10.1007/3-540-47969-4-30)
30. Lu, H., Plataniotis, K.N., Venetsanopoulos, A.N.: A survey of multilinear subspace learning for tensor data. *Pattern Recogn.* **44**(7), 1540–1551 (2011)
31. Koniusz, P., Yan, F., Gosselin, P., Mikolajczyk, K.: Higher-order occurrence pooling on mid- and low-level features: visual concept detection. Technical report (2013)
32. Koniusz, P., Yan, F., Gosselin, P., Mikolajczyk, K.: Higher-order occurrence pooling for bags-of-words: visual concept detection. *PAMI* (2016)
33. Koniusz, P., Cherian, A.: Sparse coding for third-order super-symmetric tensor descriptors with application to texture recognition. In: *CVPR* (2016)
34. Zhao, X., Wang, S., Li, S., Li, J.: A comprehensive study on third order statistical features for image splicing detection. In: *Digital Forensics and Watermarking*, pp. 243–256 (2012)
35. Jebara, T., Kondor, R., Howard, A.: Probability product kernels. *JMLR* **5**, 819–844 (2004)
36. Jégou, H., Douze, M., Schmid, C.: On the burstiness of visual elements. In: *CVPR*, pp. 1169–1176(2009)
37. Koniusz, P., Cherian, A., Porikli, F.: Tensor representations via kernel linearization for action recognition from 3D skeletons (extended version). *CoRR abs/1604.00239* (2016)
38. Zhu, Y., Chen, W., Guo, G.: Fusing spatiotemporal features and joints for 3D action recognition. In: *CVPR Workshop*, pp. 486–491(2013)
39. Zanfir, M., Leordeanu, M., Sminchisescu, C.: The moving pose: An efficient 3D kinematics descriptor for low-latency action recognition and detection. In: *ICCV*, pp. 2752–2759 (2013)