



A Novel Geometric Framework on Gram Matrix Trajectories for Human Behavior Understanding

Anis Kacem, Mohamed Daoudi, Boulbaba Ben Amor, Stefano Berretti, Juan Carlos Alvarez-Paiva

► To cite this version:

Anis Kacem, Mohamed Daoudi, Boulbaba Ben Amor, Stefano Berretti, Juan Carlos Alvarez-Paiva. A Novel Geometric Framework on Gram Matrix Trajectories for Human Behavior Understanding. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020, 42 (1), pp.1-14. 10.1109/TPAMI.2018.2872564 . hal-01883247

HAL Id: hal-01883247

<https://hal.science/hal-01883247>

Submitted on 27 Sep 2018

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

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

A Novel Geometric Framework on Gram Matrix Trajectories for Human Behavior Understanding

Anis Kacem, Mohamed Daoudi, Boulbaba Ben Amor, Stefano Berretti, and Juan Carlos Alvarez-Paiva

Abstract—In this paper, we propose a novel space-time geometric representation of human landmark configurations and derive tools for comparison and classification. We model the temporal evolution of landmarks as parametrized trajectories on the Riemannian manifold of positive semidefinite matrices of fixed-rank. Our representation has the benefit to bring naturally a second desirable quantity when comparing shapes – the spatial covariance – in addition to the conventional affine-shape representation. We derived then geometric and computational tools for rate-invariant analysis and adaptive re-sampling of trajectories, grounding on the Riemannian geometry of the underlying manifold. Specifically, our approach involves three steps: (1) landmarks are first mapped into the Riemannian manifold of positive semidefinite matrices of fixed-rank to build time-parameterized trajectories; (2) a temporal warping is performed on the trajectories, providing a geometry-aware (dis-)similarity measure between them; (3) finally, a pairwise proximity function SVM is used to classify them, incorporating the (dis-)similarity measure into the kernel function. We show that such representation and metric achieve competitive results in applications as action recognition and emotion recognition from 3D skeletal data, and facial expression recognition from videos. Experiments have been conducted on several publicly available up-to-date benchmarks.

Index Terms—Landmark configurations, Gram matrices, Riemannian Geometry, Symmetric Positive Semidefinite Manifolds, Grassmann Manifold, Action Recognition, Emotion Recognition from Body Movements, Facial Expression Recognition.



1 INTRODUCTION

SEVERAL human-related Computer Vision problems can be approached by first detecting and tracking landmarks from visual data. A relevant example of this is given by the estimated 3D location of the joints of the skeleton in depth streams [1], and their use in action and daily activity recognition. More sophisticated solutions for automatic tracking of the skeleton do exist, as the IR body markers used in Mo-Cap systems, but they are expensive in cost and time. Another relevant example is represented by the face, for which several approaches have been proposed for fiducial points detection and tracking in video [2], [3]. These techniques generate temporal sequences of landmark configurations, which exhibit several variations due to affine and projective transformations, inaccurate detection and tracking, missing data, etc. While there have been many efforts in the analysis of temporal sequences of landmarks, the problem is far from being solved and the current solutions are facing many technical and practical problems. For instance, many general techniques for temporal sequence analysis rely on computing the Euclidean distance between two temporal sequences and do not take into account the implicit dynamics of the sequences [4]. In practice, when analyzing the temporal dynamics of landmark configurations, there are four main aspects to deal with that require us to define: (1) A shape representation invariant to undesirable transformations; (2) A temporal modeling of landmark sequences; (3) A suitable rate-invariant distance between arbitrary sequences, and (4) A solution for temporal sequence classification.

In this paper, we propose a method that effectively model the comparison and classification of temporal sequences of landmarks. In doing so, we define new solutions for the four points listed above. Considering the first issue, we propose a novel shape modeling, invariant to rigid motion, by embedding shapes represented with their corresponding Gram matrix into the Positive Semidefinite Riemannian manifold. Such representation has the advantage of bringing naturally a second desirable quantity when comparing shapes – the spatial covariance – in addition to the conventional affine-shape representation. For the second issue, that is to model the dynamics and dependency relations in both temporal and spatial domains, we represent the temporal evolution of landmarks as parametrized trajectories on the Riemannian manifold of positive semidefinite matrices of fixed-rank. For what concerns the third issue, geometric and computational tools for rate-invariant analysis and adaptive re-sampling of trajectories, grounding on the Riemannian geometry of the manifold, are proposed. Finally, a variant of SVM that takes into account the nonlinearity of this space is proposed for trajectory classification. An overview of the full approach is given in Fig. 1.

A preliminary version of this work appeared in [5], with application to facial expression recognition using 2D landmarks. In this work, we generalize the idea significantly, by considering new applications using 3D landmarks. In summary, the main contributions of this work are:

- A novel static shape representation based on Gramian matrices of centered 2D and 3D landmark configurations. A comprehensive study of the Riemannian geometry of the space of representations, termed the cone of Positive Semi-Definite $n \times n$ matrices of fixed-rank d is conducted ($d = 2$ or $d = 3$

• A. Kacem, M. Daoudi, and B. Ben Amor are with IMT Lille Douai and CRISTAL laboratory, CNRS-UMR-9189, France.
 • S. Berretti is with the University of Florence, Italy.
 • J.C. Alvarez-Paiva is with the University of Lille and Painlevé laboratory, CNRS-UMR-8524, France.

for 2D or 3D landmark configurations, respectively). Despite the large use of these matrices in several research fields, to our knowledge, this is the first application in static and dynamic shape analysis.

- In addition to the affine-invariant analysis, our representation brings a spatial covariance of the landmarks. In comparison to the preliminary version of this work appeared in [5], the proposed framework has been extended to study trajectories of 3D landmarks. The effectiveness of the Gramian representation, *i.e.*, including the spatial covariance, compared to the Grassmannian was confirmed by experiments of 3D action and emotion recognition.
- A temporal extension of the representation via parametrized trajectories in the underlying Riemannian manifold, with associated computational tools for temporal alignment and adaptive re-sampling of trajectories.
- A solution for trajectory classification based on pairwise proximity function SVM (ppfSVM), where pairwise (dis-)similarity measures between trajectories are computed using the metric of the underlying manifold.
- Extensive experiments of our framework in three applications – 3D action recognition, emotion recognition from 3D human motion, and 2D facial expression recognition – demonstrate its competitiveness with respect to the state-of-the-art.

The remaining of the paper is organized as follows: In Section 2, we discuss on the related works that use trajectories to model the temporal dynamics in different application contexts; In Section 3, we propose a method to represent static landmark configurations using the Gram matrix, and also provide a comprehensive study of the Riemannian geometry of the space of these matrices; Based on this mathematical background, in Section 4, temporal sequences of landmarks are modeled as trajectories on the Gramian manifold, and a suitable measure for comparing trajectories on the manifold is proposed. The classification of the trajectories on the manifold is described in Section 5. An extensive experimental validation of the proposed approach is reported in Section 6. Experiments account for different application contexts, including 3D human action recognition, emotion recognition from 3D body movement and 2D facial expression recognition, also comparing with state-of-the-art solutions. Finally, conclusions are drawn in Section 7.

2 RELATED WORK

In the following, we first review essential literature solutions proposed for modeling the temporal evolution of landmark sequences, then we focus on mapping static shape representations as trajectories on Riemannian manifolds. Note that, often, shape-preserving transformations are filtered out from the static representation, and the rate-invariance is proposed at the trajectory level.

A. Temporal Modeling of Landmark Sequences – In the work of Slama *et al.* [6], a temporal sequence was represented as a Linear Dynamical System (LDS). The observability matrix of the LDS was then approximated by a

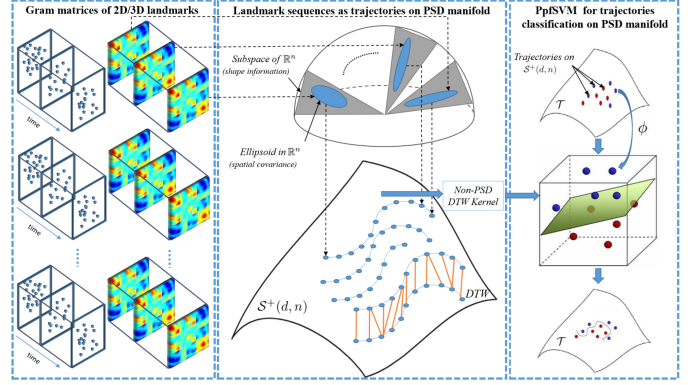


Fig. 1. Overview of the proposed approach. Given a landmark sequence, the Gram matrices are computed for each landmark configuration to build trajectories on $S^+(d, n)$. A moving shape is hence assimilated to an ellipsoid traveling along d -dimensional subspaces of \mathbb{R}^n , with d_{S^+} used to compare static ellipsoids. Dynamic Time Warping (DTW) is then used to align and compare trajectories in a rate-invariant manner. Finally, the ppfSVM is used on these trajectories for classification.

finite matrix [7]. The subspace spanned by the columns of this finite observability matrix corresponds to a point on a Grassmann manifold. Thus, the LDS is represented at each time-instant as a point on the Grassmann manifold. Each video sequence is modeled as an element of the Grassmann manifold, and action learning and recognition is cast to a classification problem on this manifold. Proximity between two spatio-temporal sequences is measured by a distance between two subspaces on the Grassmann manifold. Huang *et al.* [8] formulated the LDS as an infinite Grassmann manifold, and Venkataraman *et al.* [9] proposed a shape-theoretic framework for analysis of non-linear dynamical systems. Applications were shown to activity recognition using motion capture and RGB-D sensors, and to activity quality assessment for stroke rehabilitation. Taking a different direction, the authors of [10] and [11] proposed to map full skeletal sequences onto SPD manifolds. That is, given an arbitrary sequence, it is summarized by a covariance matrix derived from the velocities computed from neighboring frames or from the 3D landmarks themselves, respectively. In both of these works kernelized versions of covariance matrices are considered. Zhang *et al.* [12] represented temporal landmark sequences using regularized Gram matrices derived from the Hankel matrices of landmark sequences. The authors show that the Hankel matrix of a 3D landmark sequence is related to an Auto-Regressive (AR) model [13], where only the linear relationships between landmark static observations are captured. The Gram matrix of the Hankel matrix is computed to reduce the noise and is seen as a point on the positive semi-definite manifold. To analyze/compare the Gram matrices, they regularized their ranks resulting in positive definite matrices and considered metrics on the positive definite manifold.

Several solutions have experimented the application of Recurrent Neural Networks (RNNs) and Long Short Term Memory (LSTM) networks to the case of 3D landmarks represented by the joints of the human skeleton for 3D human action recognition. In fact, human actions can be interpreted as time series of body configurations, which can be represented in an effective and compact way by the

3D locations of the joints of the skeleton. In this way, each video sample can be modeled as a sequential representation of configurations. This approach was followed by Vee-riah *et al.* [14] who presented a family of differential RNNs (dRNNs) that extend LSTM by a new gating mechanism to extract the derivatives of the internal state (DoS). The DoS was fed to the LSTM gates to learn salient dynamic patterns in 3D skeleton data. Du *et al.* [15] proposed an end-to-end hierarchical RNN for skeleton based action recognition. First, the human skeleton was divided into five parts, which are then feed to five subnets. As the number of layers increases, the representations in the subnets are hierarchically fused to be the inputs of higher layers. The final representations of the skeleton sequences are fed into a single-layer perceptron, and the temporally accumulated output of the perceptron is the final decision. To ensure effective learning of the deep model, Zhu *et al.* [16] designed an in-depth dropout algorithm for the LSTM neurons in the last layer, which helps the network to learn complex motion dynamics. To further regularize the learning, a co-occurrence inducing norm was added to the networks cost function, which enforced the learning of groups of co-occurring and discriminative joints. A part aware LSTM model was proposed by Shahroudy *et al.* [17] to utilize the physical structure of the human body to improve the performance of the LSTM learning framework. Instead of keeping a long-term memory of the entire body's motion in the cell, this is split to part-based cells. In this way, the context of each body part is kept independently, and the output of the part based LSTM (P-LSTM) unit is represented as a combination of independent body part context information. Each part cell has therefore its individual input, forget, and modulation gates, but the output gate is shared among the body parts. A broader review of methods that apply deep architectures can be found in the survey by Ioannidou *et al.* [18] for the case of generic 3D data, and in the survey by Wang *et al.* [19] and Berretti *et al.* [20] for the particular case of human action recognition from 3D data.

B. Analyzing Shape Trajectories on Riemannian Manifolds – One promising idea is to formulate the motion features as trajectories. Matikainen *et al.* [21] presented a method for using the trajectories of tracked feature points in a bag of words paradigm for video action recognition. Despite of the promising results obtained, the authors did not take into account the geometric information of the trajectories. More recently, in the case of human skeleton in RGB-Depth images, Devanne *et al.* [22] proposed to formulate the action recognition task as the problem of computing a distance between trajectories generated by the joints moving during the action. An action is then interpreted as a parameterized curve on the hyper-sphere of the human skeleton. However, this approach does not take into account the relationship between the joints. In the same direction, Su *et al.* [23] proposed a metric that considers the time-warping on a Riemannian manifold, thus allowing trajectories registration and the computation of statistics on the trajectories. Su *et al.* [24] applied this framework to the problem of visual speech recognition. Similar ideas have been developed by Ben Amor *et al.* [25] on the Kendall's shape space with application to action recognition using rate-invariant analysis of skeletal shape trajectories.

Anirudh *et al.* [26] started from the framework of Trans-ported Square-Root Velocity Fields (TSRVF), which has desirable properties including a rate-invariant metric and vector space representation. Based on this framework, they proposed to learn an embedding such that each action trajectory is mapped to a single point in a low-dimensional Euclidean space, and the trajectories that differ only in temporal rates map to the same point. The TSRVF representation and accompanying statistical summaries of Riemannian trajectories are used to extend existing coding methods such as PCA, KSVD, and Label Consistent KSVD to Riemannian trajectories. In the experiments, it is shown such coding efficiently captures trajectories in action recognition, stroke rehabilitation, visual speech recognition, clustering, and diverse sequence sampling.

In [27], Vemulapalli *et al.* proposed a Lie group trajectory representation of the skeletal data on the product space of Special Euclidean (SE) groups. For each frame, the latter representation is obtained by computing the Euclidean transformation matrices encoding rotations and translations between different joint pairs. The temporal evolution of these matrices is seen as a trajectory on $SE(3) \times \dots \times SE(3)$ and mapped to the tangent space of a reference point. A one-versus-all SVM, combined with Dynamic Time Warping and Fourier Temporal Pyramid (FTP) is used for classification. One limitation of this method is that mapping trajectories to a common tangent space using the logarithm map could result in significant approximation errors. Aware of this limitation, the same authors proposed in [28] a mapping combining the usual logarithm map with a rolling map that guarantees a better flattening of trajectories on Lie groups.

3 REPRESENTATION OF STATIC LANDMARK CONFIGURATIONS

Let us consider an arbitrary sequence of landmark configurations $\{Z_0, \dots, Z_\tau\}$. Each configuration Z_i ($0 \leq i \leq \tau$) is an $n \times d$ matrix of rank d encoding the positions of n distinct landmark points in d dimensions. In our applications, we only consider the configurations of landmark points in two- or three-dimensional space (*i.e.*, $d=2$ or $d=3$) given by, respectively, $p_1 = (x_1, y_1), \dots, p_n = (x_n, y_n)$ or $p_1 = (x_1, y_1, z_1), \dots, p_n = (x_n, y_n, z_n)$. We are interested in studying such sequences or curves of landmark configurations up to Euclidean motions. In the following, we will first propose a representation for static observations, then adopt a time-parametrized representation for temporal analysis.

As a first step, we seek a shape representation that is invariant up to Euclidean transformations (rotation and translation). Arguably, the most natural choice is the matrix of pairwise distances between the landmarks of the same shape augmented by the distances between all the landmarks and their center of mass p_0 . Since we are dealing with Euclidean distances, it will turn out to be more convenient to consider the matrix of the squares of these distances. Also note that by subtracting the center of mass from the coordinates of the landmarks, these can be considered as *centered*: the center of mass is always at the origin. From now on, we will assume $p_0 = (0, 0)$ for $d = 2$ (or $p_0 = (0, 0, 0)$ for $d = 3$).

for $d = 3$). With this provision, the augmented pairwise square-distance matrix \mathcal{D} takes the form,

$$\mathcal{D} := \begin{pmatrix} 0 & \|p_1\|^2 & \cdots & \|p_n\|^2 \\ \|p_1\|^2 & 0 & \cdots & \|p_1 - p_n\|^2 \\ \vdots & \vdots & \ddots & \vdots \\ \|p_n\|^2 & \|p_n - p_1\|^2 & \cdots & 0 \end{pmatrix},$$

where $\|\cdot\|$ denotes the norm associated to the l^2 -inner product $\langle \cdot, \cdot \rangle$. A key observation is that the matrix \mathcal{D} can be easily obtained from the $n \times n$ Gram matrix $G := ZZ^T$. Indeed, the entries of G are the pairwise inner products of the points p_1, \dots, p_n ,

$$G = ZZ^T = \langle p_i, p_j \rangle, \quad 1 \leq i, j \leq n, \quad (1)$$

and the equality

$$\mathcal{D}_{ij} = \langle p_i, p_i \rangle - 2\langle p_i, p_j \rangle + \langle p_j, p_j \rangle, \quad 0 \leq i, j \leq n, \quad (2)$$

establishes a linear equivalence between the set of $n \times n$ Gram matrices and the augmented square-distance $(n+1) \times (n+1)$ matrices of distinct landmark points. On the other hand, Gram matrices of the form ZZ^T , where Z is an $n \times d$ matrix of rank d are characterized as $n \times n$ positive semidefinite matrices of rank d . For a detailed discussion of the relation between positive semidefinite matrices, Gram matrices, and square-distance matrices, we refer the reader to Section 6.2.1 of [29]. Conveniently for us, the Riemannian geometry of the space of these matrices, called the positive semidefinite cone $\mathcal{S}^+(d, n)$, was studied in [30], [31], [32], [33]. An alternative shape representation, considered in [34] and [35], associates to each configuration Z the d -dimensional subspace $\text{span}(Z)$ spanned by its columns. This representation, which exploits the well-known geometry of the Grassmann manifold $\mathcal{G}(d, n)$ of d -dimensional subspaces in \mathbb{R}^n is invariant under *all* invertible linear transformations. By fully encoding the set of all mutual distances between landmark points, the Euclidean shape representation proposed in this paper supplements the affine shape representation with the knowledge of the $d \times d$ covariance matrix for the centered landmarks that lie on the manifold of Symmetric Positive Definite (SPD) matrices. This leads to considerable improvements in the results of the conducted experiments.

3.1 Riemannian Geometry of $\mathcal{S}^+(d, n)$

Given an $n \times d$ matrix Z of rank d , its polar decomposition $Z = UR$ with $R = (Z^T Z)^{1/2}$ allows us to write the Gram matrix ZZ^T as $UR^2 U^T$. Since the columns of the matrix U are orthonormal, this decomposition defines a map

$$\begin{aligned} \Pi : \mathcal{V}_{d,n} \times \mathcal{P}_d &\rightarrow \mathcal{S}^+(d, n) \\ (U, R^2) &\mapsto UR^2 U^T, \end{aligned}$$

from the product of the Stiefel manifold $\mathcal{V}_{d,n}$ and the cone of $d \times d$ positive definite matrices \mathcal{P}_d to the manifold $\mathcal{S}^+(d, n)$ of $n \times n$ positive semidefinite matrices of rank d . The map Π defines a principal fiber bundle over $\mathcal{S}^+(d, n)$ with fibers

$$\Pi^{-1}(UR^2 U^T) = \{(UO, O^T R^2 O) : O \in \mathcal{O}(d)\},$$

where $\mathcal{O}(d)$ is the group of $d \times d$ orthogonal matrices. Bonnabel and Sepulchre [30] used this map and the geometry of the *structure space* $\mathcal{V}_{d,n} \times \mathcal{P}_d$ to introduce a Riemannian metric on $\mathcal{S}^+(d, n)$ and study its geometry.

3.2 Tangent Space and Riemannian Metric

The tangent space $T_{(U, R^2)}(\mathcal{V}_{d,n} \times \mathcal{P}_d)$ consists of pairs (M, N) , where M is a $n \times d$ matrix satisfying $M^T U + U^T M = 0$ and N is any $d \times d$ symmetric matrix. Bonnabel and Sepulchre defined a *connection* (see [36, p. 63]) on the principal bundle $\Pi : \mathcal{V}_{d,n} \times \mathcal{P}_d \rightarrow \mathcal{S}^+(d, n)$ by setting the horizontal subspace $\mathcal{H}_{(U, R^2)}$ at the point (U, R^2) to be the space of tangent vectors (M, N) such that $M^T U = 0$ and N is an arbitrary $d \times d$ symmetric matrix. They also defined an inner product on $\mathcal{H}_{(U, R^2)}$: given two tangent vectors $A = (M_1, N_1)$ and $B = (M_2, N_2)$ on $\mathcal{H}_{(U, R^2)}$, set

$$\langle (A, B) \rangle_{\mathcal{H}_{U, R^2}} = \text{tr}(M_1^T M_2) + k \text{tr}(N_1 R^{-2} N_2 R^{-2}), \quad (3)$$

where $k > 0$ is a real parameter.

It is easily checked that the action of the group of $d \times d$ orthogonal matrices on the fiber $\Pi^{-1}(UR^2 U^T)$ sends horizontals to horizontals isometrically. It follows that the inner product on $T_{UR^2 U^T} \mathcal{S}^+(d, n)$ induced from that of $\mathcal{H}_{(U, R^2)}$ via the linear isomorphism $D\Pi$ is independent of the choice of point (U, R^2) projecting onto $UR^2 U^T$. This procedure defines a Riemannian metric on $\mathcal{S}^+(d, n)$ for which the natural projection

$$\begin{aligned} \rho : \mathcal{S}^+(d, n) &\rightarrow \mathcal{G}(d, n) \\ G &\mapsto \text{range}(G), \end{aligned}$$

is a Riemannian submersion. This allows us to relate the geometry of $\mathcal{S}^+(d, n)$ with that of the Grassmannian $\mathcal{G}(d, n)$.

Recall that the geometry of the Grassmannian $\mathcal{G}(d, n)$ is easily described by using the map

$$\text{span} : \mathcal{V}_{d,n} \rightarrow \mathcal{G}(d, n),$$

that sends an $n \times d$ matrix with orthonormal columns U to their span $\text{span}(U)$. Given two subspaces $\mathcal{U}_1 = \text{span}(U_1)$ and $\mathcal{U}_2 = \text{span}(U_2) \in \mathcal{G}(d, n)$, the geodesic curve connecting them is

$$\text{span}(U(t)) = \text{span}(U_1 \cos(\Theta t) + M \sin(\Theta t)), \quad (4)$$

where Θ is a $d \times d$ diagonal matrix formed by the *principal angles* between \mathcal{U}_1 and \mathcal{U}_2 , while the matrix M is given by the formula $M = (I_n - U_1 U_1^T) U_2 F$, with F being the pseudoinverse $\text{diag}(\sin(\theta_1), \sin(\theta_2))$. The Riemannian distance between \mathcal{U}_1 and \mathcal{U}_2 is given by

$$d_{\mathcal{G}}^2(\mathcal{U}_1, \mathcal{U}_2) = \|\Theta\|_F^2. \quad (5)$$

3.3 Pseudo-Geodesics and Closeness in $\mathcal{S}^+(d, n)$

Bonnabel and Sepulchre [30] defined the *pseudo-geodesic* connecting two matrices $G_1 = U_1 R_1^2 U_1^T$ and $G_2 = U_2 R_2^2 U_2^T$ in $\mathcal{S}^+(d, n)$ as the curve

$$\mathcal{C}_{G_1 \rightarrow G_2}(t) = U(t) R^2(t) U^T(t), \quad \forall t \in [0, 1], \quad (6)$$

where $R^2(t) = R_1 \exp(t \log R_1^{-1} R_2^2 R_1^{-1}) R_1$ is a geodesic in \mathcal{P}_d connecting R_1^2 and R_2^2 , and $U(t)$ is the geodesic in $\mathcal{G}(d, n)$ given by Eq. (4). They also defined the *closeness* between G_1

and G_2 , $d_{S^+}(G_1, G_2)$, as the square of the length of this curve:

$$\begin{aligned} d_{S^+}(G_1, G_2) &= d_G^2(\mathcal{U}_1, \mathcal{U}_2) + k d_{\mathcal{P}_d}^2(R_1^2, R_2^2) \\ &= \|\Theta\|_F^2 + k \|\log R_1^{-1} R_2^2 R_1^{-1}\|_F^2, \end{aligned} \quad (7)$$

where \mathcal{U}_i ($i = 1, 2$) is the *span* of U_i and Θ is a $d \times d$ diagonal matrix formed by the principal angles between \mathcal{U}_1 and \mathcal{U}_2 . The closeness d_{S^+} consists of two independent contributions: the square of the distance $d_G(\text{span}(U_1), \text{span}(U_2))$ between the two associated subspaces, and the square of the distance $d_{\mathcal{P}_d}(R_1^2, R_2^2)$ on the positive cone \mathcal{P}_d (Fig. 2). Note that $\mathcal{C}_{G_1 \rightarrow G_2}$ is not necessarily a geodesic and therefore, the closeness d_{S^+} is not a true Riemannian distance. From the viewpoint of the landmark configurations Z_1 and Z_2 , with $G_1 = Z_1 Z_1^T$ and $G_2 = Z_2 Z_2^T$, the closeness encodes the distances measured between the affine shapes $\text{span}(Z_1)$ and $\text{span}(Z_2)$ in $\mathcal{G}(d, n)$ and between their spatial covariances in \mathcal{P}_d . Indeed, the spatial covariance of Z_i ($i = 1, 2$) is the $d \times d$ symmetric positive definite matrix

$$C = \frac{Z_i^T Z_i}{n-1} = \frac{(U_i R_i)^T (U_i R_i)}{n-1} = \frac{R_i^2}{n-1}. \quad (8)$$

The weight parameter k controls the relative weight of these two contributions. Note that for $k = 0$ the distance on $S^+(d, n)$ collapses to the distance on $\mathcal{G}(d, n)$. Nevertheless, the authors in [30] recommended choosing small values for this parameter. The experiments performed and reported in Section 6 are in general accordance with this recommendation.




Fig. 2. A pictorial representation of the positive semidefinite cone $S^+(d, n)$. Viewing matrices G_1 and G_2 as ellipsoids in \mathbb{R}^n ; their closeness consists of two contributions: d_G^2 (squared Grassmann distance) and $d_{\mathcal{P}_d}^2$ (squared Riemannian distance in \mathcal{P}_d).

4 MODELING TEMPORAL LANDMARK SEQUENCES AS TRAJECTORIES IN $S^+(d, n)$

We are able to compare static landmark configurations based on their Gramian representation G , the induced space, and closeness introduced in the previous Section. We need a natural and effective extension to study their temporal evolution. Following [25], [27], [35], we defined curves $\beta_G : I \rightarrow S^+(d, n)$ (I denotes the time domain, e.g., $[0, 1]$) to model the spatio-temporal evolution of elements on $S^+(d, n)$. Given a sequence of landmark configurations $\{Z_0, \dots, Z_\tau\}$ represented by their corresponding Gram matrices $\{G_0, \dots, G_\tau\}$ in $S^+(d, n)$, the corresponding curve is the trajectory of the point $\beta_G(t)$ on $S^+(d, n)$, when t

ranges in $[0, 1]$. These curves are obtained by connecting all successive Gramian representations of shapes G_i and G_{i+1} , $0 \leq i \leq \tau - 1$, by pseudo-geodesics in $S^+(d, n)$. Algorithm 1 summarizes the steps to build trajectories in $S^+(d, n)$ for temporal modeling of landmark sequences.

Algorithm 1: Computing trajectory $\beta_G(t)$ in $S^+(d, n)$ of a sequence of landmarks

input : A sequence of centered landmark

configurations $\{Z_0, \dots, Z_\tau\}$, where $Z_{0 \leq i \leq \tau}$ is an $(n \times d)$ matrix ($d = 2$ or $d = 3$) formed by the coordinates

$p_1 = (x_1, y_1), \dots, p_n = (x_n, y_n)$ or

$p_1 = (x_1, y_1, z_1), \dots, p_n = (x_n, y_n, z_n)$.

output: Trajectory $\beta_G(t)_{0 \leq t \leq \tau}$ and pseudo-geodesics $\mathcal{C}_{\beta_G(t) \rightarrow \beta_G(t+1)}$ in $S^+(d, n)$

```

1 /* Compute the Gram matrices of centered landmarks */
2 for i ← 0 to τ do
3    $G_i \leftarrow Z_i Z_i^T = \langle p_l, p_k \rangle, \quad 1 \leq l, k \leq n$ 
4   /* Compute the Polar decomposition1 of  $Z_i = U_i R_i$  */
5    $G_i \leftarrow U_i R_i^2 U_i^T$ 
6 end
7 /* Compute the pseudo-geodesic paths between successive
   Gram matrices */
8  $\beta_G(0) \leftarrow G_0$ 
9 for t ← 0 to τ - 1 do
10   $\mathcal{C}_{\beta_G(t) \rightarrow \beta_G(t+1)} \leftarrow \mathcal{C}_{G_t \rightarrow G_{t+1}}$  given by Eq. (6)
    connecting  $G_t$  and  $G_{t+1}$  in  $S^+(d, n)$ 
11   $\beta_G(t+1) \leftarrow G_{t+1}$ 
12 end
13 return trajectory  $\beta_G(t)_{0 \leq t \leq \tau}$  and pseudo-geodesics
     $\mathcal{C}_{\beta_G(t) \rightarrow \beta_G(t+1)}$  in  $S^+(d, n)$ 

```

4.1 Temporal Alignment and Rate-Invariant Comparison of Trajectories

A relevant issue to our classification problems is – how to compare trajectories while being invariant to rates of execution? One can formulate the problem of temporal misalignment as comparing trajectories when parameterized differently. The parameterization variability makes the distance between trajectories distorted. This issue was first highlighted by Veeraraghavan *et al.* [38] who showed that different rates of execution of the same activity can greatly decrease recognition performance if ignored. Veeraraghavan *et al.* [38] and Abdelkader *et al.* [39] used the Dynamic Time Warping (DTW) for temporal alignment before comparing trajectories of shapes of planar curves that represent silhouettes in videos. Following the above-mentioned state-of-the-art solutions, we adopt here a DTW solution to temporally align our trajectories. More formally, given m trajectories $\{\beta_G^1, \beta_G^2, \dots, \beta_G^m\}$ on $S^+(d, n)$, we are interested in finding functions γ_i such that the $\beta_G^i(\gamma_i(t))$ are matched optimally for all $t \in [0, 1]$. In other words, two curves $\beta_G^1(t)$ and $\beta_G^2(t)$ represent the same trajectory if their images are the same. This happens if, and only if, $\beta_G^2 = \beta_G^1 \circ \gamma$, where γ is a re-parameterization of the interval $[0, 1]$. The problem of

1. To compute the polar decomposition, we used the SVD based implementation proposed in [37].

temporal alignment is turned to find an optimal warping function γ^* according to,

$$\gamma^* = \arg \min_{\gamma \in \Gamma} \int_0^1 d_{S^+}(\beta_G^1(t), \beta_G^2(\gamma(t))) dt, \quad (9)$$

where Γ denotes the set of all monotonically-increasing functions $\gamma : [0, 1] \rightarrow [0, 1]$. The most commonly used method to solve such optimization problem is DTW. Note that accommodation of the DTW algorithm to the manifold-value sequences can be achieved with respect to an appropriate metric defined on the underlying manifold $S^+(d, n)$. Having the optimal re-parametrization function γ^* , one can define a (dis-)similarity measure between two trajectories allowing a rate-invariant comparison:

$$d_{DTW}(\beta_G^1, \beta_G^2) = \int_0^1 d_{S^+}(\beta_G^1(t), \beta_G^2(\gamma^*(t))) dt. \quad (10)$$

From now, we shall use $d_{DTW}(\cdot, \cdot)$ to compare trajectories in our manifold of interest $S^+(d, n)$.

4.2 Adaptive Re-Sampling of Trajectories in $S^+(d, n)$

One difficulty in video analysis is to capture the most relevant frames and focus on them. In fact, it is relevant to reduce the number of frames when no motion happened, and “introduce” new frames, otherwise. Our geometric framework provides tools to do so. In fact, interpolation between successive frames could be achieved using the pseudo-geodesics defined in Eq. (6), while their length (closeness defined in Eq. (7)) expresses the magnitude of the motion. Accordingly, we have designed an adaptive re-sampling tool that is able to increase/decrease the number of samples in a fixed time interval according to their relevance with respect to the geometry of the underlying manifold $S^+(d, n)$. Relevant samples are identified by a relatively low closeness d_{S^+} to the previous frame, while irrelevant ones correspond to a higher closeness level. Here, the down-sampling is performed by removing irrelevant shapes. In turn, the up-sampling is possible by interpolating between successive shape representations in $S^+(d, n)$, using pseudo-geodesics.

More formally, given a trajectory $\beta_G(t)_{t=0,1,\dots,\tau}$ on $S^+(d, n)$ for each sample $\beta_G(t)$, we compute the closeness to the previous sample, i.e., $d_{S^+}(\beta_G(t), \beta_G(t-1))$: if the value is below a defined threshold ζ_1 , the current sample is simply removed from the trajectory. In contrast, if the distance exceeds a second threshold ζ_2 , equally spaced shape representations from the pseudo-geodesic curve connecting $\beta_G(t)$ to $\beta_G(t-1)$ are inserted in the trajectory.

5 CLASSIFICATION OF TRAJECTORIES IN $S^+(d, n)$

Our trajectory representation reduces the problem of landmark sequence classification to that of trajectory classification in $S^+(d, n)$. That is, let us consider $\mathcal{T} = \{\beta_G : [0, 1] \rightarrow S^+(d, n)\}$, the set of time-parameterized trajectories of the underlying manifold. Let $\mathcal{L} = \{(\beta_G^1, y^1), \dots, (\beta_G^m, y^m)\}$ be the training set with class labels, where $\beta_G^i \in \mathcal{T}$ and $y^i \in \mathcal{Y}$, such that $y^i = f(\beta_G^i)$. The goal here is to find an approximation h to f such that $h : \mathcal{T} \rightarrow \mathcal{Y}$. In Euclidean spaces, any standard classifier (e.g., standard SVM) may be a natural and

appropriate choice to classify the trajectories. Unfortunately, this is no more suitable in our modeling, as the space \mathcal{T} built from $S^+(d, n)$ is non-linear. A function that divides the manifold is rather a complicated notion compared with the Euclidean space. In current literature, two main approaches have been used to handle the nonlinearity of Riemannian manifolds [27], [35], [40]. These methods map the points on the manifold to a tangent space or to Hilbert space, where traditional learning techniques can be used for classification. Mapping data to a tangent space only yields a first-order approximation of the data that can be distorted, especially in regions far from the origin of the tangent space. Moreover, iteratively mapping back and forth, i.e., Riemannian Logarithmic and Exponential maps, to the tangent spaces significantly increases the computational cost of the algorithm. Recently, some authors proposed to embed a manifold in a high dimensional Reproducing Kernel Hilbert Space (RKHS), where Euclidean geometry can be applied [40]. The Riemannian kernels enable the classifiers to operate in an extrinsic feature space without computing tangent space and log and exp maps. Many Euclidean machine learning algorithms can be directly generalized to an RKHS, which is a vector space that possesses an important structure: the inner product. Such an embedding, however, requires a kernel function defined on the manifold which, according to Mercer’s theorem, should be positive semi-definite.

5.1 Pairwise Proximity Function SVM Classifier

Inspired by a recent work of [41] for action recognition, we adopted the *pairwise proximity function SVM* (ppfSVM) [42], [43]. The ppfSVM requires the definition of a (dis-)similarity measure to compare samples. In our case, it is natural to consider the d_{DTW} defined in Eq. (10) for such a comparison. This strategy involves the construction of inputs such that each trajectory is represented by its (dis-)similarity to all the trajectories, with respect to d_{DTW} , in the dataset and then apply a conventional SVM to this transformed data [43]. The ppfSVM is related to the arbitrary kernel-SVM without restrictions on the kernel function [42].

Given m trajectories $\{\beta_G^1, \beta_G^2, \dots, \beta_G^m\}$ in \mathcal{T} , following [41], a proximity function $\mathcal{P}_{\mathcal{T}} : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}_+$ between two trajectories $\beta_G^1, \beta_G^2 \in \mathcal{T}$ is defined as,

$$\mathcal{P}_{\mathcal{T}}(\beta_G^1, \beta_G^2) = d_{DTW}(\beta_G^1, \beta_G^2). \quad (11)$$

According to [42], there are no restrictions on the function $\mathcal{P}_{\mathcal{T}}$. For an input trajectory $\beta_G \in \mathcal{T}$, the mapping $\phi(\beta_G)$ is given by,

$$\phi(\beta_G) = [\mathcal{P}_{\mathcal{T}}(\beta_G, \beta_G^1), \dots, \mathcal{P}_{\mathcal{T}}(\beta_G, \beta_G^m)]^T. \quad (12)$$

The obtained vector $\phi(\beta_G) \in \mathbb{R}^m$ is used to represent a sample trajectory $\beta_G \in \mathcal{T}$. Hence, the set of trajectories can be represented by a $m \times m$ matrix P , where $P(i, j) = \mathcal{P}_{\mathcal{T}}(\beta_G^i, \beta_G^j)$, $i, j \in \{1, \dots, m\}$. Finally, a linear SVM is applied to this data representation. Further details on ppfSVM can be found in [41], [42], [43]. In Algorithm 2, we provide a pseudo-code for the proposed trajectory classification in $S^+(d, n)$.

The proposed ppfSVM classification of trajectories on $S^+(d, n)$ aims to learn a proximity model of the data, which

Algorithm 2: Classification of trajectories in $\mathcal{S}^+(d, n)$

input : m training trajectories in $\mathcal{S}^+(d, n)$ with their corresponding labels $\{(\beta_G^1, y^1), \dots, (\beta_G^m, y^m)\}$
 One testing trajectory β_G^{test} in $\mathcal{S}^+(d, n)$
output: Predicted class y^{test} of β_G^{test}

```

1 /* Model training */
2 for  $i \leftarrow 1$  to  $m$  do
3   for  $j \leftarrow 1$  to  $m$  do
4      $P(i, j) = \mathcal{P}_{\mathcal{T}}(\beta_G^i, \beta_G^j)$  w.r.t Eq. (11)
5   end
6 end
7 Training a linear SVM on the data representation  $P$ 
8 /* Testing phase */
9  $\phi(\beta_G^{test}) = [\mathcal{P}_{\mathcal{T}}(\beta_G^{test}, \beta_G^1), \dots, \mathcal{P}_{\mathcal{T}}(\beta_G^{test}, \beta_G^m)]^T$ 
10  $y^{test} \leftarrow$  Linear SVM using the feature vector  $\phi(\beta_G^{test})$ 
11 return Predicted class  $y^{test}$ 

```

makes the computation of a pairwise distance function using the DTW (dis-)similarity measure on all the trajectories of the dataset quite necessary. For more efficiency, one can consider faster algorithms for trajectories alignment such as [44], [45].

For comparison purposes, we also evaluated a k -nearest neighbor solution, where for each test trajectory (sequence), we computed the k -nearest trajectories (sequences) from the training set using the same (dis-)similarity measure d_{DTW} defined in Eq. (10). The test sequence is then classified according to a majority voting of its neighbors, (i.e., it is assigned to the class that is most common among its k -nearest neighbors).

6 EXPERIMENTAL RESULTS

To validate the proposed framework, we conducted extensive experiments on three human behavior understanding applications. These scenarios show the potential of the proposed solution when landmarks capture different information on different data. First, we addressed the problem of activity recognition from depth sensors such as the Microsoft Kinect. In this case, 3D landmarks correspond to the joints of the body skeleton, as extracted from RGB-Depth frames. The number of joints per skeleton varies between 15 and 20, and their position is generally noisy. Next, we addressed the new emerging problem of finding relationships between body movement and emotions using 3D skeletal data. Here, landmarks correspond to physical markers placed on the body and tracked with high temporal rate and good estimation of the 3D position by a Motion Capture (MoCap) system. Finally, we evaluated our framework on the problem of facial expression recognition using landmarks of the face. In this case, 49 face landmarks are extracted in 2D with high accuracy using a state-of-the-art face landmark detector.

6.1 3D Human Action Recognition

Action recognition has been performed on 3D skeleton data as provided by a Kinect camera in different datasets. In this case, landmarks correspond to the estimated position of 3D joints of the skeleton ($d=3$). With this assumption, skeletons

are represented by $n \times n$ Gram matrices of rank 3 lying on $\mathcal{S}^+(3, n)$, and skeletal sequences are seen as trajectories on this manifold.

As discussed in Section 3, the information given by the Gram matrix of the skeleton is linearly equivalent to that of the pairwise distances between different joints. Thus, considering only some specific subparts of the skeletons can be more accurate for some actions. For instance, it is more discriminative to consider only the pairwise distances between the joints of left and right arms for actions that involve principally the motion of arms, (e.g., *wave hands*, *throw*). Accordingly, we divided the skeletons into three body parts, i.e., left/right arms, left/right legs and torso, while keeping a coarse information given by all the joints of the skeleton (we show an example of this decomposition in the supplementary material). For an efficient use of the information given by the different body parts, we propose a late fusion of four ppf-SVM classifiers that consists of: (1) training all the body part classifiers separately; (2) merging the contributions of the four body part classifiers. This is done by multiplying the probabilities $s_{i,j}$, output of the SVM for each class j , where $i \in \{1, 2, 3, 4\}$ denotes the body part. The class \mathcal{C} of each test sample is determined by

$$\mathcal{C} = \arg \max_j \prod_{i=1}^4 s_{i,j}, \quad j = 1, \dots, n_{\mathcal{C}}, \quad (13)$$

where $n_{\mathcal{C}}$ is the number of classes.

6.1.1 Datasets

We performed experiments on four publicly available datasets showing different challenges. All these datasets have been collected with a Microsoft Kinect sensor.

UT-Kinect dataset [56] – It contains 10 actions performed by 10 different subjects. Each subject performed each action twice resulting in 199 valid action sequences. The 3D locations of 20 joints are provided with the dataset.

Florence3D dataset [51] – It contains 9 actions performed two or three times by 10 different subjects. Skeleton comprises 15 joints. This is a challenging dataset due to variations in the view-point and large intra-class variations.

SYSU-3D dataset [47] – It contains 480 sequences. In this dataset, 12 different activities focusing on interactions with objects were performed by 40 persons. The 3D coordinates of 20 joints are provided in this dataset. The SYSU-3D dataset is very challenging since the motion patterns are highly similar among different activities.

SBU Interaction dataset [57] – This dataset includes 282 skeleton sequences of eight types of two-persons interacting with each other, including *approaching*, *departing*, *pushing*, *kicking*, *punching*, *exchanging objects*, *hugging*, and *shaking hands*. In most interactions, one subject is acting, while the other subject is reacting.

6.1.2 Experimental Settings and Parameters

For all the datasets, we used only the provided skeletons. The adaptive re-sampling of trajectories discussed in Section 4.2 has been not applied on these data. The motivation is that this operation tries to capture small shape deformations of the landmarks and this can amplify the noise of skeleton joints. For the SBU dataset, where two skeletons

TABLE 1

Overall accuracy (%) on the UT-Kinect, Florence3D, SBU interaction, and SYSU-3D datasets. Here, $(^D)$: depth; $(^C)$: color (or RGB); $(^G)$: geometry (or skeleton); *: Deep Learning based approach; last row: ours

Method	UT-Kinect		Florence3D		SBU Interaction		SYSU-3D	
	Protocol	Acc (%)	Protocol	Acc (%)	Protocol	Acc (%)	Protocol	Acc (%)
$(G+D)$ 3D ² CNN [46]*	LOSO	95.5	—	—	—	—	—	—
$(G+D+C)$ Dynamic features [47]	—	—	—	—	—	—	Half-Half	84.9 ± 2.29
$(G+D+C)$ LAFF [48]	—	—	—	—	—	—	Half-Half	80
$(^G)$ LARP [27]	5-fold	97.08	5-fold	90.88	—	—	—	—
$(^G)$ Gram Hankel [12]	LOOCV	100	—	—	—	—	—	—
$(^G)$ Motion trajectories [22]	LOOCV	91.5	LOSO	87.04	—	—	—	—
$(^G)$ Elastic func. coding [49]	5-fold	94.87	5-fold	89.67	—	—	—	—
$(^G)$ Mining key poses [50]	LOOCV	93.47	LOSO	92.25	—	—	—	—
$(^G)$ NBNn+parts+time [51]	—	—	LOSO	82	—	—	—	—
$(^G)$ LAFF (SKL) [48]	—	—	—	—	—	—	Half-Half	54.2
$(^G)$ Dynamic skeletons [47]	—	—	—	—	—	—	Half-Half	75.5 ± 3.08
$(^G)$ LSTM-trust gate [52]*	LOOCV	97.0	—	—	5-fold	93.3	Half-Half	76.5
$(^G)$ JL-distance LSTM [53]*	5-fold	95.96	—	—	5-fold	99.02	—	—
$(^G)$ Co-occurrence LSTM [16]*	—	—	—	—	5-fold	90.41	—	—
$(^G)$ Hierarchical RNN [15]*	—	—	—	—	5-fold	80.35	—	—
$(^G)$ SkeletonNet [54]*	—	—	—	—	5-fold	93.47	—	—
$(^G)$ STA-LSTM [55]*	—	—	—	—	5-fold	91.51	—	—
Traj. on $\mathcal{G}(3, n)$ (full body)	LOOCV	92.46	LOSO	75 ± 5.22	5-fold	76.3 ± 3.26	Half-Half	73.26 ± 2.27
Traj. on $\mathcal{G}(3, n)$ - BP Fusion	LOOCV	96.48	LOSO	76.4 ± 5.37	5-fold	83.56 ± 4.72	Half-Half	76.61 ± 2.86
Traj. on $\mathcal{S}^+(3, n)$ (full body)	LOOCV	96.48	LOSO	88.07 ± 4.8	5-fold	88.45 ± 2.88	Half-Half	76.01 ± 2.09
Traj. on $\mathcal{S}^+(3, n)$ - BP Fusion	LOOCV	98.49	LOSO	88.85 ± 4.6	5-fold	93.7 ± 1.59	Half-Half	80.22 ± 2.09

of two interacting persons are given in each frame, we considered all the joints of the two skeletons. In this case, a unique Gram matrix is computed for the two skeletons modeling the interaction between them. In this dataset, the decomposition into body parts is performed only for the acting person since the other person is reacting in a coarse manner.

As discussed in Section 3.3, our body movement representation involves a parameter k that controls the contribution of two information: the affine shape of the skeleton at time t , and its spatial covariance. The affine shape information is given by the Grassmann manifold $\mathcal{G}(3, n)$, while the spatial covariance is given by the SPD manifold \mathcal{P}_3 . We recall that for $k = 0$, the skeletons are considered as trajectories on the Grassmann manifold $\mathcal{G}(3, n)$. For each dataset, we performed a cross-validation grid search, $k \in [0, 3]$ with a step of 0.1, to find an optimal value k^* . In the case of skeleton decomposition into body parts, a different parameter k is used for computing the distance of each body part, *i.e.*, one parameter each for arms, legs, and torso, and one parameter for the whole skeleton). Each parameter k is evaluated separately by a cross-validation grid search in the classifier of the relative body part.

To allow a fair comparison, we adopted the most common experimental settings in literature. For the UT-Kinect dataset, we used the *leave-one-out cross-validation* (LOOCV) protocol [56], where one sequence is used for testing and the remaining sequences are used for training. For the Florence3D dataset, a *leave-one-subject-out* (LOSO) schema is adopted following [12], [22], [50]. For the SYSU3D dataset, we followed [47] and performed a *Half-Half* cross-subject test setting, in which half of the subjects were used for training and the remaining half were used for testing. Finally, a *5-fold* cross-validation was used for the SBU dataset. Note that the subjects considered in each split are those given by the datasets (SYSU3D and SBU). All our programs were

implemented in Matlab and run on a 2.8 GHZ CPU. We used the multi-class SVM implementation of the LibSVM library [58].

6.1.3 Results and Discussion

In Table 1, we compare our approach with existing methods dealing with skeletons and/or RGB-D data. Overall, our approach achieved competitive results compared to recent state-of-the-art approaches. We provide the confusion matrices for all the datasets in the supplementary material.

On the UT-Kinect dataset, we obtained an average accuracy of 96.48%, when considering the full skeletal shape. Using a late fusion of classifiers based on the body parts, as described in Section 6.1, the performance increased to 98.49% outperforming [22], [50], [52]. The highest average accuracy for this dataset was reported in [12] (100%), where Gram matrices were used for skeletal sequence representation, but in a completely different context. Specifically, the authors of [12] built a Gram matrix from the Hankel matrix of an Auto-Regressive (AR) model that represented the dynamics of the skeletal sequences. The used metric for the comparison of Gram matrices is also different than ours as they used metrics in the positive definite cone by regularizing their ranks, *i.e.*, making them full-rank.

On the SBU dataset, the fusion of body parts achieved the highest accuracy reaching 93.7%. We observed that all the interactions present in this dataset are well recognized, *e.g.*, *hugging* (100%), *approaching* (97.5%), *etc.*, except *pushing* (74.7%), which has been mainly confused with a very similar interaction, *i.e.*, *punching*. Here, our approach is ranked second after [53], where an average accuracy of 99.02% is reported. In that work, the authors compute a large number of joint-line distances per frame making their approach time consuming.

On the SYSU3D dataset, our approach achieved the best result compared to skeleton based approaches. We report

an average accuracy of 80.22% with a standard deviation of 2.09%, when the late fusion of body parts is used. Our approach, applied to the full skeleton, still achieved very competitive results and reached 76.01% with a standard deviation of 2.09%. Combining the skeletons with depth and color information, including the object, Hu *et al.* [47] obtained the highest performance with an average accuracy of 84.9% and a standard deviation of 2.29%.

On the Florence3D dataset, we obtained an average accuracy of 88.07%, improved by around 0.8% when involving body parts fusion. While high accuracies are reported for coarse actions, *e.g.*, *sitting down* (95%), *standing up* (100%), and *lacing* (96.2%), finer actions, *e.g.*, *reading watch* (73.9%) and *answering phone* (68.2%) are still challenging. Our results are outperformed by [27], [50], where the average accuracies are greater than 90%.

Baseline Experiments. In this paragraph, we discuss the effect of using the different steps in our framework and their computational complexity compared to baselines. Results of this evaluation are reported in Table 2. Firstly, in the top part of Table 2, we studied the computational cost of the proposed pipeline in the task of 3D action recognition and report running time statistics for the different steps of our approach on UT-Kinect dataset. Specifically, we provide the necessary execution time for: (1) an arbitrary trajectory construction in $\mathcal{S}^+(3, n)$ as described in Algorithm 1; (2) comparison of two arbitrary trajectories with the proposed version of DTW; (3) testing phase of an arbitrary trajectory classification with ppfSVM in $\mathcal{S}^+(3, n)$ as described in Algorithm 2.

TABLE 2
Baseline experiments on the UT-Kinect, SBU, SYSU3D, and Florence3D datasets

Pipeline component		Time (s)
Trajectory construction in $\mathcal{S}^+(3, n)$		0.007
Comparison of trajectories in $\mathcal{S}^+(3, n)$		0.93
Classification of a trajectory in $\mathcal{S}^+(3, n)$		147.71

Distance	UT-Kinect (%)	Time (s)
Flat distance $d_{\mathcal{F}^+}$	92.96	0.06
Distance $d_{\mathcal{P}_n}$ in \mathcal{P}_n	94.98	1.66
Closeness $d_{\mathcal{S}^+}$	96.48	0.93

Temp. alignment	UT-Kinect (%)	SBU (%)	Time (s)
No DTW	91.46	81.36 ± 2.78	0.02
DTW	96.48	88.45 ± 2.88	0.93

Classifier	UT-Kinect (%)	SBU (%)
K-NN – $\mathcal{G}(3, n)$	86.93	42.72 ± 5.68
Ppf-SVM – $\mathcal{G}(3, n)$	92.46	76.3 ± 3.26
K-NN – $\mathcal{S}^+(3, n)$	91.96	61.06 ± 2.3
Ppf-SVM – $\mathcal{S}^+(3, n)$	96.48	88.45 ± 2.88

Body parts	UT-Kinect (%)	SBU (%)
Arms only	87.94	80.96 ± 5.53
Legs only	35.68	83.36 ± 2.41
Torso only	72.36	80.58 ± 2.16
Whole body	96.48	88.45 ± 2.88
Late BP Fusion	98.49	93.7 ± 1.59

Body parts	Florence3D (%)	SYSU3D (%)
Arms only	75.72 ± 8.45	73.88 ± 2.64
Legs only	42.44 ± 7.69	37.6 ± 2.10
Torso only	54.33 ± 10.62	49.36 ± 3.94
Whole body	88.07 ± 4.8	76.01 ± 2.09
Late BP Fusion	88.85 ± 4.6	80.22 ± 2.09

Secondly, we can observe the large superiority of the Gramian representation over the Grassmann representation. For the Florence3D and SBU datasets, we report an improvement of about 12%. For UT-Kinect and SYSU3D, the performance increased by about 3%. Note that these improvements over the Grassmannian representation are due to the additional information of the spatial covariance given by the SPD manifold in the metric. The contribution of the spatial covariance is weighted with a parameter k . As discussed in Section 6.1.2, we performed a grid search cross-validation to find the optimal value k^* of this parameter. The optimal values are $k^* = 0.05$, $k^* = 0.81$, $k^* = 0.25$, and $k^* = 0.09$ for the UT-Kinect, SBU, Florence3D, and SYSU3D datasets, respectively. These results are in concordance with the recommendation of Bonnabel and Sepulchre [30] to use relative small values of k .

Then, we evaluated the proposed metric with respect to other metrics used in state of the art solutions. Specifically, given two matrices G_1 and G_2 in $\mathcal{S}^+(3, n)$, we compared our results with two other possible metrics: (1) as proposed in [12], [59], we used $d_{\mathcal{P}_n}$ that was defined in Eq. (7) to compare G_1 and G_2 by regularizing their ranks, *i.e.*, making them n full-rank, and considering them in \mathcal{P}_n (the space of n -by- n positive definite matrices), $d_{\mathcal{P}_n}(G_1, G_2) = d_{\mathcal{P}_n}(G_1 + \epsilon I_n, G_2 + \epsilon I_n)$; (2) we used the Euclidean flat distance $d_{\mathcal{F}^+}(G_1, G_2) = \|G_1 - G_2\|_F$, where $\|\cdot\|_F$ denotes the Frobenius-norm. Note that the provided execution times are relative to the comparison of two arbitrary sequences. We can observe that in Table 2, the closeness $d_{\mathcal{S}^+}$ between two elements of $\mathcal{S}^+(3, n)$ defined in Eq. (7) is more suitable compared to the distance $d_{\mathcal{P}_n}$ and the flat distance $d_{\mathcal{F}^+}$ defined in literature. This demonstrates the importance of considering the geometry of the manifold of interest. Another advantage of using $d_{\mathcal{S}^+}$ over $d_{\mathcal{P}_n}$ is the computational time as it involves n -by-3 and 3-by-3 matrices instead of n -by- n matrices.

To show the relevance of aligning the skeleton sequences in time before comparing them, we conducted the same experiments without using Dynamic Time Warping (DTW). In this case, the performance decreased by around 5% and 7% on UT-Kinect and SBU datasets, respectively. Here, the provided execution times are relative to the comparison of two arbitrary sequences on UT-Kinect dataset. Furthermore, we also compared the proposed ppfSVM classifier with a k -nearest neighbor classifier. The number of nearest neighbors k to consider for each dataset is chosen by cross-validation. Using the k -NN classifier, we obtained an average accuracy of 91.96% with $k = 5$ neighbors on UT-Kinect and 61.06% with $k = 4$ on the SBU dataset. These results are outperformed by the ppfSVM classifier.

Finally, in Table 2 we provide the obtained accuracies when considering the different body parts separately on all the datasets. Unsurprisingly, the highest accuracy is achieved by left and right arms in all the datasets compared to the torso and the legs, since the majority of the actions are acted using arms. One can note the considerable improvements realized by the late fusion compared to the whole skeleton in all the datasets, especially in the SBU and SYSU3D datasets, where we report improvements of about 5% and 4%, respectively.

6.2 Emotion Recognition from 3D Body Movement

Recently, the study of computational models for human emotion recognition has gained increasing attention not only for commercial applications (to get feedback on the effectiveness of advertising material), but also for gaming and monitoring of the emotional state of operators that act in risky contexts such as aviation. Most of these studies have focused on the analysis of facial expressions, but important clues can be derived by the analysis of the dynamics of body parts as well [60]. Using the same geometric framework that was proposed for action recognition, we evaluated our approach in the task of emotion recognition from human body movement. Here, the used landmarks are in 3D coordinate space, but with better accuracy and higher temporal resolution, with respect to the case of action recognition.

6.2.1 Dataset

Experiments have been performed on the Body Motion-Emotion dataset (P-BME), acquired at the Cognitive Neuroscience Laboratory (INSERM U960 - Ecole Normale Supérieure) in Paris [60]. It includes Motion Capture (MoCap) 3D data sequences recorded at high frame rate (120 frames per second) by an Opto-electronic Vicon V8 MoCap system wired to 24 cameras. The body movement is captured by using 43 landmarks that are positioned at joints.

To create the dataset, 8 subjects (professional actors) were instructed to walk following a predefined “U” shaped path that includes forward-walking, turn, and coming back. For each acquisition, actors moved along the path performing one emotion out of five different emotions, namely, *anger*, *fear*, *joy*, *neutral*, and *sadness*. So, each sequence is associated with one emotion label. Each actor performed at maximum five repetitions of a same emotional sequence for a total of 156 instances. Though there is some variation from subject to subject, the number of examples is well distributed across the different emotions: 29 *anger*, 31 *fear*, 33 *joy*, 28 *neutral*, 35 *sadness*.

6.2.2 Experimental Settings and Parameters

Since MoCap skeletons are in 3D coordinate space, we followed the same steps that have been proposed for action recognition, including the decomposition into body parts (we show an example of this decomposition in the supplementary material). Note that the same late fusion of body part classifiers, as mentioned in the previous Section, is adopted. A cross-validation grid search has been performed to find an optimal value for the weight parameter k .

Experiments on the P-BME dataset were performed by using a *leave-one-subject-out* cross validation protocol. With this solution, iteratively, all the emotion sequences of a subject are used for test, while all the sequences of the remaining subjects are used for training.

6.2.3 Results and Discussion

In Table 3, we provide the obtained results as well as a comparative study with baseline experiments on the P-BME dataset.

Similarly to the reported results for action recognition, the proposed fusion of body part classifiers achieved the highest performance with an average accuracy of 81.99%

TABLE 3

Comparative study of the proposed approach with baseline experiments on the P-BME dataset. First rows: state-of-the-art action and emotion recognition methods and human evaluator; second rows: baseline experiments; last row: ours

Method	Accuracy (%)
Human evaluator	74.20
COV3D [61]	71.14 ± 6.77
LARP [27]	74.8 ± 3.17
Traj. on $S^+(3, n)$ - Flat metric	57.41 ± 8.43
Traj. on $S^+(3, n)$ - No DTW	63.23 ± 8.62
Traj. on $S^+(3, n)$ - k NN	68.9 ± 7.63
Traj. on $\mathcal{G}(3, n)$	66.35 ± 6.43
Traj. on $\mathcal{G}(3, n)$ - BP Fusion	67.09 ± 6.82
Traj. on $S^+(3, n)$	78.15 ± 5.79
Traj. on $S^+(3, n)$ - BP Fusion	81.99 ± 4.36

and standard deviation of 4.36%. The best results were scored by *neutral* and *anger* (more than 80%), followed by *fear* (71%), *joy* (about 67%), with the lowest accuracy for *sadness* (about 65%). In the supplementary material, we provide the related confusion matrix as well as some additional experiments. Considering only the skeletons (without body parts) in the classification, the performance decreased to an average accuracy of 78.15%.

Recently, Daoudi *et al.* [61] proposed a method for emotion recognition from body movement based on covariance matrices and SPD manifold. They used the 3D covariance descriptor (COV3D) of skeleton joints across time to represent sequences without a special handling of the dynamics. They reported an average accuracy of 71.4%. They also performed a user based test in order to evaluate the performance of the proposed classification method in comparison with a human-based judgment. In this test, thirty-two naive individuals were asked to perform a force-choice task in which they had to choose between one of the five emotions. This resulted in an average value of about 74%. It is relevant to note that the user based test being based on RGB videos provides to the users much more information for evaluation, including the actor’s face. Notably, our method is capable to score better results based on the skeleton joints only.

We also compared our results with the Lie algebra relative pairs (LARP) method proposed by Vemulapalli *et al.* [27] for skeleton action recognition. In that work, each skeleton is mapped to a point on the product space of $SE(3) \times SE(3) \cdots \times SE(3)$, where it is modeled using transformations between joint pairs. The temporal evolution of these features is seen as a trajectory on $SE(3) \times SE(3) \times \cdots \times SE(3)$ and mapped to the tangent space of a reference point. A one-versus-all SVM combined with Dynamic Time Warping and Fourier temporal pyramid (FTP) is used for classification. Using this method, an average accuracy of 74.8% was obtained, which is about 8% lower than ours.

The highest accuracy (78.15%) is obtained for $k^* = 1.2$. For $k = 0$, the skeletons are considered as trajectories on the Grassmann manifold $\mathcal{G}(3, n)$, and the obtained accuracy is around 66%, which is 12% lower than the retained result. In order to show the importance of choosing a well defined Riemannian metric in the space of interest, we conducted the same experiments by changing the metric d_{S^+} defined in Eq. (7) with a flat metric, defined as the Frobenius norm of the difference between two Gram matrices (skeletons). For

this experiment, we report an average accuracy of 57.41% being lower of about 21% than using d_{S^+} .

Finally, as mentioned in Section 4.1, an important step in our approach is the temporal alignment. Avoiding this step and following the same protocol, we found that the performance decreased to 63.23%. We also studied the method when considering the different body parts separately and also when considering different sequence lengths. Results of these additional experiments are reported in the supplementary material.

6.3 2D Facial Expression Recognition

We evaluated our approach also in the task of facial expression recognition from 2D landmarks. In this case, the landmarks are in a 2D coordinate space, resulting in a Gram matrix of size $n \times n$ of rank 2 for each configuration of n landmarks. The facial sequences are then seen as time-parameterized trajectories on $S^+(2, n)$.

6.3.1 Datasets

We conducted experiments on four publicly available datasets – CK+, MMI, Oulu-CASIA, and AFEW datasets.

Cohn-Kanade Extended (CK+) dataset [62] – It contains 123 subjects and 593 frontal image sequences of posed expressions. Among them, 118 subjects are annotated with the seven labels – *anger* (An), *contempt* (Co), *disgust* (Di), *fear* (Fe), *happy* (Ha), *sad* (Sa) and *surprise* (Su). Note that only the two first temporal phases of the expression, *i.e.*, neutral and onset (with apex frames), are present.

MMI dataset [63] – It consists of 205 image sequences with frontal faces of 30 subjects labeled with the six basic emotion labels. In this dataset each sequence begins with a neutral facial expression, and has a posed facial expression in the middle; the sequence ends up with the neutral facial expression. The location of the peak frame is not provided as a prior information.

Oulu-CASIA dataset [64] – It includes 480 image sequences of 80 subjects, taken under normal illumination conditions. They are labeled with one of the six basic emotion labels. Each sequence begins with a neutral facial expression and ends with the apex of the expression.

AFEW dataset [65] – Collected from movies showing close-to-real-world conditions, which depict or simulate the spontaneous expressions in uncontrolled environment. The task is to classify each video clip into one of the seven expression categories (the six basic emotions plus the neutral).

6.3.2 Experimental Settings and Parameters

All our experiments were performed once facial landmarks were extracted using the method proposed in [2] on the CK+, MMI, and Oulu-CASIA datasets. On the challenging AFEW dataset, we have considered the corrections provided in² after applying the same detector. The number of landmarks is $n = 49$ for each face. In this case, we applied the adaptive re-sampling of trajectories proposed in Section 4.2 that enhances small facial deformations and disregards redundant frames. This step involves two parameters ζ_1 and ζ_2 for up-sampling and down-sampling, respectively. These

two parameters are chosen so that all the trajectories in the dataset have the same length, equal to the median length. For the parameter k , the same procedure as for action and emotion recognition from body movement is applied.

To evaluate our approach, we followed the experimental settings commonly used in recent works. Following [66], [67], [68], [69], we have performed 10-fold cross validation experiments for the CK+, MMI, and Oulu-CASIA datasets. In contrast, the AFEW dataset was divided into three sets: training, validation and test, according to the protocols defined in EmotiW'2013 [70]. Here, we only report our results on the validation set for comparison with [66], [68], [70].

6.3.3 Results and Discussion

On CK+, the average accuracy is 96.87%. Note that the accuracy of the trajectory representation on $\mathcal{G}(2, n)$, following the same pipeline is 2% lower, which confirms the contribution of the covariance embedded in our representation.

An average classification accuracy of 79.19% is reported for the MMI dataset. Note that based on geometric features only, our approach grounding on both representations on $S^+(2, n)$ and $\mathcal{G}(2, n)$ achieved competitive results with respect to the literature (see Table 4). On the Oulu-CASIA dataset, the average accuracy is 83.13%, hence 3% higher than the Grassmann trajectory representation. This is the highest accuracy reported in literature (refer to Table 5). Finally, we reported an average accuracy of 39.94% on the AFEW dataset. Despite being competitive with respect to recent literature (see Table 5), these results evidence that AFER “in-the-wild” is still challenging.

We highlight the superiority of the trajectory representation on $S^+(2, n)$ over the Grassmannian (refer to Table 4 and Table 5). This is due to the contribution of the covariance part further to the conventional affine-shape analysis over the Grassmannian. Recall that k serves to balance the contribution of the distance between covariance matrices living in \mathcal{P}_2 with respect to the Grassmann contribution $\mathcal{G}(2, n)$. The optimal performance are achieved for the following values – $k_{CK+}^* = 0.081$, $k_{MMI}^* = 0.012$, $k_{Oulu-CASIA}^* = 0.014$ and $k_{AFEW}^* = 0.001$. To show the importance of the proposed adaptive re-sampling step, we conducted the same experiments on the MMI and AFEW datasets avoiding this step. The performance decreases of about 5% on MMI and 3% on AFEW. The temporal alignment, the effectiveness of the used metric, and the used classifier were also evaluated according to the same conducted baseline experiments for action recognition (see Section 6.1.3). Results show the superiority of the retained framework over these baselines. Further details, as well as the confusion matrices, are available in the supplementary material of this paper.

Comparative Study with the State-of-the-Art. In Table 4 and Table 5, we compare our approach over the recent literature. Overall, our approach achieved competitive performance with respect to the most recent approaches. On CK+, we obtained the second highest accuracy. The ranked-first approach is DTAGN [67], in which two deep networks are trained on shape and appearance channels, then fused. Note that the geometry deep network (DTGN) achieved 92.35%, which is much lower than ours. Furthermore, our approach outperforms the ST-RBM [66] and the STM-ExpLet [68]. On the MMI dataset, our approach outperforms

2. <http://sites.google.com/site/chehrahome>

the DTAGN [67] and the STM-ExpLet [68]. However, it is behind ST-RBM [66].

TABLE 4
Overall accuracy (%) on CK+ and MMI datasets. Here, ^(A): appearance (or color); ^(G): geometry (or shape); *: Deep Learning based approach; last row: ours

Method	CK+	MMI
^(A) 3D HOG (from [67])	91.44	60.89
^(A) 3D SIFT (from [67])	-	64.39
^(A) Cov3D (from [67])	92.3	-
^(A) STM-ExpLet [68] (10-fold)	94.19	75.12
^(A) CSPL [69] (10-fold)	89.89	73.53
^(A) F-Bases [71] (LOSO)	96.02	75.12
^(A) ST-RBM [66] (10-fold)	95.66	81.63
^(A) 3DCNN-DAP [72] * (15-fold)	87.9	62.2
^(A) DTAN [67] * (10-fold)	91.44	62.45
^(A+G) DTAGN [67] * (10-fold)	97.25	70.24
^(G) DTGN [67] * (10-fold)	92.35	59.02
^(G) TMS [73] (4-fold)	85.84	-
^(G) HMM [74] (15-fold)	83.5	51.5
^(G) ITBN [74] (15-fold)	86.3	59.7
^(G) Velocity on $\mathcal{G}(n, 2)$ [35]	82.8	-
^(G) traj. on $\mathcal{G}(2, n)$ (10-fold)	94.25 ± 3.71	78.18 ± 4.87
^(G) traj. on $\mathcal{S}^+(2, n)$ (10-fold)	96.87 ± 2.46	79.19 ± 4.62

TABLE 5
Overall accuracy on Oulu-CASIA and AFEW dataset (following the EmotiW'13 protocol [70])

Method	Oulu-CASIA	AFEW
^(A) HOG 3D [75]	70.63	26.90
^(A) 3D SIFT [76]	55.83	24.87
^(A) LBP-TOP [77]	68.13	25.13
^(A) EmotiW [70]	-	27.27
^(A) STM [68]	-	29.19
^(A) STM-ExpLet [68]	74.59	31.73
^(A+G) DTAGN [67] * (10-fold)	81.46	-
^(A) ST-RBM [66]	-	46.36
^(G) traj. on $\mathcal{G}(2, n)$	80.0 ± 5.22	39.1
^(G) traj. on $\mathcal{S}^+(2, n)$	83.13 ± 3.86	39.94

On the Oulu-CASIA dataset, our approach shows a clear superiority to existing methods, in particular STM-ExpLet [68] and DTGN [67]. Elaiwat *et al.* [66] do not report any results on this dataset, however, their approach achieved the highest accuracy on AFEW. Our approach is ranked second showing a superiority to remaining approaches on AFEW.

7 CONCLUSION

In this paper, we have proposed a geometric approach for effectively modeling and classifying dynamic 2D and 3D landmark sequences for human behavior understanding. Based on Gramian matrices derived from the static landmarks, our representation consists of an affine-invariant shape representation and a spatial covariance of the landmarks. We have exploited the geometry of the space to define a closeness between static shape representations. Then, we have derived computational tools to align, re-sample and compare these trajectories giving rise to a rate-invariant analysis. Finally, landmark sequences are learned from these trajectories using a variant of SVM, called

ppfSVM, which allows us to deal with the nonlinearity of the space of representation. We evaluated our approach in three different applications, namely, 3D human action recognition, 3D emotion recognition from body movement, and 2D facial expression recognition. Extensive experiments on nine publicly available datasets showed that the proposed approach achieves competitive or better results than state-of-art solutions.

8 ACKNOWLEDGEMENTS

We thank Professor Yvonne Delevoye for fruitful discussions on emotion from 3D human dynamics. This work has been partially supported by PIA, ANR (grant ANR-11-EQPX- 0023).

REFERENCES

- [1] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook, and R. Moore, "Real-time human pose recognition in parts from single depth images," *Communications of the ACM*, vol. 56, no. 1, pp. 116–124, 2013.
- [2] A. Athana, S. Zafeiriou, S. Cheng, and M. Pantic, "Incremental face alignment in the wild," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 1859–1866.
- [3] X. Xiong and F. D. la Torre, "Supervised descent method and its applications to face alignment," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 532–539.
- [4] L. Zafeiriou, E. Antonakos, S. Zafeiriou, and M. Pantic, "Joint unsupervised deformable spatio-temporal alignment of sequences," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 3382–3390.
- [5] A. Kacem, M. Daoudi, B. Ben Amor, and J. C. Alvarez-Paiva, "A novel space-time representation on the positive semidefinite cone for facial expression recognition," in *IEEE Int. Conf. on Computer Vision (ICCV)*, October 2017.
- [6] R. Slama, H. Wannous, M. Daoudi, and A. Srivastava, "Accurate 3D action recognition using learning on the grassmann manifold," *Pattern Recognition*, vol. 48, no. 2, pp. 556–567, 2015.
- [7] P. K. Turaga, A. Veeraraghavan, A. Srivastava, and R. Chellappa, "Statistical computations on grassmann and stiefel manifolds for image and video-based recognition," *IEEE Trans. Pattern Analysis Machine Intelligence*, vol. 33, no. 11, pp. 2273–2286, 2011.
- [8] W. Huang, F. Sun, L. Cao, D. Zhao, H. Liu, and M. Harandi, "Sparse coding and dictionary learning with linear dynamical systems," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 3938–3947.
- [9] V. Venkataraman and P. K. Turaga, "Shape distributions of non-linear dynamical systems for video-based inference," *CoRR*, vol. abs/1601.07471, 2016.
- [10] J. Cavazza, A. Zunino, M. San Biagio, and V. Murino, "Kernelized covariance for action recognition," in *Pattern Recognition (ICPR)*, 2016 23rd International Conference on. IEEE, 2016, pp. 408–413.
- [11] L. Wang, J. Zhang, L. Zhou, C. Tang, and W. Li, "Beyond covariance: Feature representation with nonlinear kernel matrices," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 4570–4578.
- [12] X. Zhang, Y. Wang, M. Gou, M. Szaier, and O. Camps, "Efficient temporal sequence comparison and classification using Gram matrix embeddings on a riemannian manifold," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [13] B. Li, O. I. Camps, and M. Szaier, "Cross-view activity recognition using hankels," in *Computer Vision and Pattern Recognition (CVPR)*, 2012 IEEE Conference on. IEEE, 2012, pp. 1362–1369.
- [14] V. Veeriah, N. Zhuang, and G. J. Qi, "Differential recurrent neural networks for action recognition," in *IEEE Int. Conf. on Computer Vision (ICCV)*, Dec 2015, pp. 4041–4049.
- [15] Y. Du, W. Wang, and L. Wang, "Hierarchical recurrent neural network for skeleton based action recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 1110–1118.

- [16] W. Zhu, C. Lan, J. Xing, W. Zeng, Y. Li, L. Shen, and X. Xie, "Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks," in *AAAI Conf. on Artificial Intelligence (AAAI)*, 2016, pp. 3697–3703.
- [17] A. Shahroudy, J. Liu, T. T. Ng, and G. Wang, "Ntu rgb+d: A large scale dataset for 3d human activity analysis," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 1010–1019.
- [18] A. Ioannidou, E. Chatzilaris, S. Nikolopoulos, and I. Kompatsiaris, "Deep learning advances in computer vision with 3D data: A survey," *ACM Computer Surveys*, vol. 50, no. 2, pp. 20:1–20:38, Apr. 2017.
- [19] P. Wang, W. Li, P. O. Ogunbona, J. Wan, and S. Escalera, "RGB-D-based motion recognition with deep learning: A survey," *Int. Journal on Computer Vision*, to appear 2017.
- [20] S. Berretti, M. Daoudi, P. Turaga, and A. Basu, "Representation, analysis and recognition of 3D humans: A survey," *ACM Trans. on Multimedia Computing, Communication and Applications*, vol. 14, no. 1s, pp. 1–35, March 2018.
- [21] P. Matikainen, M. Hebert, and R. Sukthankar, "Trajectons: Action recognition through the motion analysis of tracked features," in *Workshop on Video-Oriented Object and Event Classification (ICCV)*, September 2009.
- [22] M. Devanne, H. Wannous, S. Berretti, P. Pala, M. Daoudi, and A. Del Bimbo, "3-D human action recognition by shape analysis of motion trajectories on Riemannian manifold," *IEEE Trans. on Cybernetics*, vol. 45, no. 7, pp. 1340–1352, 2015.
- [23] J. Su, S. Kurtek, E. Klassen, and A. Srivastava, "Statistical analysis of trajectories on riemannian manifolds: Bird migration, hurricane tracking and video surveillance," *Annals of Applied Statistics*, vol. 8, no. 1, 2014.
- [24] J. Su, A. Srivastava, F. D. M. de Souza, and S. Sarkar, "Rate-invariant analysis of trajectories on riemannian manifolds with application in visual speech recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [25] B. Ben Amor, J. Su, and A. Srivastava, "Action recognition using rate-invariant analysis of skeletal shape trajectories," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 38, no. 1, pp. 1–13, 2016.
- [26] R. Anirudh, P. K. Turaga, J. Su, and A. Srivastava, "Elastic functional coding of riemannian trajectories," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 39, no. 5, pp. 922–936, 2017.
- [27] R. Vemulapalli, F. Arrate, and R. Chellappa, "Human action recognition by representing 3D skeletons as points in a Lie group," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 588–595.
- [28] R. Vemulapalli and R. Chellappa, "Rolling rotations for recognizing human actions from 3d skeletal data," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 4471–4479.
- [29] M. M. Deza and M. Laurent, *Geometry of cuts and metrics*. Springer, 2009, vol. 15.
- [30] S. Bonnabel and R. Sepulchre, "Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank," *SIAM Journal on Matrix Analysis and Applications*, vol. 31, no. 3, pp. 1055–1070, 2009.
- [31] M. Faraki, M. T. Harandi, and F. Porikli, "Image set classification by symmetric positive semi-definite matrices," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2016, pp. 1–8.
- [32] G. Meyer, S. Bonnabel, and R. Sepulchre, "Regression on fixed-rank positive semidefinite matrices: a riemannian approach," *Journal of Machine Learning Research*, vol. 12, no. Feb, pp. 593–625, 2011.
- [33] B. Vandereycken, P.-A. Absil, and S. Vandewalle, "Embedded geometry of the set of symmetric positive semidefinite matrices of fixed rank," in *IEEE/SP Workshop on Statistical Signal Processing (SSP)*, 2009, pp. 389–392.
- [34] E. Begelfor and M. Werman, "Affine invariance revisited," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2006, pp. 2087–2094.
- [35] S. Taheri, P. Turaga, and R. Chellappa, "Towards view-invariant expression analysis using analytic shape manifolds," in *IEEE Int. Conf. on Automatic Face & Gesture Recognition and Workshops (FG)*, 2011, pp. 306–313.
- [36] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*. Interscience Publishers, 1963, vol. 1.
- [37] N. J. Higham, "Computing the polar decomposition with applications," *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 4, pp. 1160–1174, 1986.
- [38] A. Veeraraghavan, R. Chellappa, and A. Roy-Chowdhury, "The function space of an activity," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 2006, pp. 959–968.
- [39] M. F. Abdelkader, W. Abd-Elmageed, A. Srivastava, and R. Chellappa, "Silhouette-based gesture and action recognition via modeling trajectories on riemannian shape manifolds," *Computer Vision and Image Understanding*, vol. 115, no. 3, pp. 439–455, Mar. 2011.
- [40] S. Jayasumana, R. I. Hartley, M. Salzmann, H. Li, and M. T. Harandi, "Kernel methods on riemannian manifolds with gaussian RBF kernels," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 37, no. 12, pp. 2464–2477, 2015.
- [41] M. A. Bagheri, Q. Gao, and S. Escalera, "Support vector machines with time series distance kernels for action classification," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2016, pp. 1–7.
- [42] T. Graepel, R. Herbrich, P. Bollmann-Sdorra, and K. Obermayer, "Classification on pairwise proximity data," *Advances in Neural Information Processing Systems*, pp. 438–444, 1999.
- [43] S. Gudmundsson, T. P. Runarsson, and S. Sigurdsson, "Support vector machines and dynamic time warping for time series," in *IEEE World Congress on Computational Intelligence*, 2008, pp. 2772–2776.
- [44] S. Salvador and P. Chan, "Toward accurate dynamic time warping in linear time and space," *Intelligent Data Analysis*, vol. 11, no. 5, pp. 561–580, 2007.
- [45] M. Cuturi, "Fast global alignment kernels," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 929–936.
- [46] Z. Liu, C. Zhang, and Y. Tian, "3d-based deep convolutional neural network for action recognition with depth sequences," *Image and Vision Computing*, vol. 55, pp. 93–100, 2016.
- [47] J.-F. Hu, W.-S. Zheng, J. Lai, and J. Zhang, "Jointly learning heterogeneous features for rgb-d activity recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 5344–5352.
- [48] J.-F. Hu, W.-S. Zheng, L. Ma, G. Wang, and J. Lai, "Real-time RGB-D activity prediction by soft regression," in *European Conf. on Computer Vision (ECCV)*, 2016, pp. 280–296.
- [49] R. Anirudh, P. Turaga, J. Su, and A. Srivastava, "Elastic functional coding of riemannian trajectories," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 39, no. 5, pp. 922–936, May 2017.
- [50] C. Wang, Y. Wang, and A. L. Yuille, "Mining 3d key-pose-motifs for action recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2639–2647.
- [51] L. Seidenari, V. Varano, S. Berretti, A. Bimbo, and P. Pala, "Recognizing actions from depth cameras as weakly aligned multi-part bag-of-poses," in *IEEE Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2013, pp. 479–485.
- [52] J. Liu, A. Shahroudy, D. Xu, and G. Wang, *Spatio-Temporal LSTM with Trust Gates for 3D Human Action Recognition*. Cham: Springer Int. Publishing, 2016, pp. 816–833.
- [53] S. Zhang, X. Liu, and J. Xiao, "On geometric features for skeleton-based action recognition using multilayer lstm networks," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2017, pp. 148–157.
- [54] Q. Ke, S. An, M. Bennamoun, F. Sohel, and F. Boussaid, "Skeletonnet: Mining deep part features for 3-d action recognition," *IEEE Signal Processing Letters*, vol. 24, no. 6, pp. 731–735, 2017.
- [55] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu, "An end-to-end spatio-temporal attention model for human action recognition from skeleton data," in *AAAI*, 2017, pp. 4263–4270.
- [56] L. Xia, C. Chen, and J. Aggarwal, "View invariant human action recognition using histograms of 3d joints," in *IEEE Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2012, pp. 20–27.
- [57] K. Yun, J. Honorio, D. Chattopadhyay, T. L. Berg, and D. Samaras, "Two-person interaction detection using body-pose features and multiple instance learning," in *IEEE Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2012, pp. 28–35.
- [58] C.-C. Chang and C.-J. Lin, "Libsvm: a library for support vector machines," *ACM Trans. on Intelligent Systems and Technology*, vol. 2, no. 3, p. 27, 2011.
- [59] R. Wang, H. Guo, L. S. Davis, and Q. Dai, "Covariance discriminative learning: A natural and efficient approach to image set classification," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2496–2503.

- [60] H. Hicheur, H. Kadone, J. Grèzes, and A. Berthoz, "The combined role of motion-related cues and upper body posture for the expression of emotions during human walking," in *Modeling, Simulation and Optimization of Bipedal Walking*. Springer Berlin Heidelberg, 2013, pp. 71–85.
- [61] M. Daoudi, S. Berretti, P. Pala, Y. Delevoeye, and A. Bimbo, "Emotion recognition by body movement representation on the manifold of symmetric positive definite matrices," in *Int. Conf. on Image Analysis and Processing*, to appear 2017.
- [62] P. Lucey, J. F. Cohn, T. Kanade, J. M. Saragih, Z. Ambadar, and I. A. Matthews, "The extended Cohn-Kanade dataset (CK+): A complete dataset for action unit and emotion-specified expression," in *IEEE Conf. on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2010, pp. 94–101.
- [63] M. F. Valstar and M. Pantic, "Induced disgust, happiness and surprise: an addition to the mmi facial expression database," in *Int. Conf. on Language Resources and Evaluation, Workshop on EMOTION*, Malta, May 2010, pp. 65–70.
- [64] G. Zhao, X. Huang, M. Taini, S. Z. Li, and M. Pietikäinen, "Facial expression recognition from near-infrared videos," *Image Vision Computing*, vol. 29, no. 9, pp. 607–619, 2011.
- [65] A. Dhall, R. Goecke, S. Lucey, and T. Gedeon, "Collecting large, richly annotated facial-expression databases from movies," *IEEE MultiMedia*, vol. 19, no. 3, pp. 34–41, 2012.
- [66] S. Elaiwat, M. Bennamoun, and F. Boussaid, "A spatio-temporal rbm-based model for facial expression recognition," *Pattern Recognition*, vol. 49, pp. 152–161, 2016.
- [67] H. Jung, S. Lee, J. Yim, S. Park, and J. Kim, "Joint fine-tuning in deep neural networks for facial expression recognition," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2015, pp. 2983–2991.
- [68] M. Liu, S. Shan, R. Wang, and X. Chen, "Learning expressionlets on spatio-temporal manifold for dynamic facial expression recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 1749–1756.
- [69] L. Zhong, Q. Liu, P. Yang, B. Liu, J. Huang, and D. N. Metaxas, "Learning active facial patches for expression analysis," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2012, pp. 2562–2569.
- [70] A. Dhall, R. Goecke, J. Joshi, M. Wagner, and T. Gedeon, "Emotion recognition in the wild challenge (EmotiW) challenge and workshop summary," in *Int. Conf. on Multimodal Interaction (ICMI)*, 2013, pp. 371–372.
- [71] E. Sariyanidi, H. Gunes, and A. Cavallaro, "Learning bases of activity for facial expression recognition," *IEEE Trans. on Image Processing*, vol. PP, no. 99, pp. 1–1, 2017.
- [72] M. Liu, S. Li, S. Shan, R. Wang, and X. Chen, "Deeply learning deformable facial action parts model for dynamic expression analysis," in *Asian Conf. on Computer Vision*, 2014, pp. 143–157.
- [73] S. Jain, C. Hu, and J. K. Aggarwal, "Facial expression recognition with temporal modeling of shapes," in *IEEE Int. Conf. on Computer Vision Workshops (ICCV)*, 2011, pp. 1642–1649.
- [74] Z. Wang, S. Wang, and Q. Ji, "Capturing complex spatio-temporal relations among facial muscles for facial expression recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 3422–3429.
- [75] A. Kläser, M. Marszalek, and C. Schmid, "A spatio-temporal descriptor based on 3d-gradients," in *British Machine Vision Conf. (BMVC)*, 2008, pp. 1–10.
- [76] P. Scovanner, S. Ali, and M. Shah, "A 3-dimensional sift descriptor and its application to action recognition," in *Int. Conf. on Multimedia*, 2007, pp. 357–360.
- [77] G. Zhao and M. Pietikäinen, "Dynamic texture recognition using local binary patterns with an application to facial expressions," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 29, no. 6, pp. 915–928, 2007.



derstanding.



nals and international conferences. He is Associate Editor of *Image and Vision Computing Journal* and *IEEE Transactions on Multimedia*. He is a General co-Chair of IEEE International Conference on Automatic Face and Gesture Recognition 2019. He is Fellow of IAPR and IEEE Senior member.



human behavior analysis and assessment.



Transactions on Multimedia Computing, Communications, and Applications, and Associate Editor of the *IET Computer Vision journal*.



Anis Kacem received his engineering degree in applied sciences and technology in 2014, from the National Institute of Applied Sciences and Technology (INSAT), Tunisia. Currently, he is a Ph.D. candidate in Computer Science at Institut Mines Télécom (IMT) Lille-Douai, France, working with the 3DSAM research team of CRISTAL laboratory (CNRS UMR 9189). His research interests are mainly focused on pattern recognition and shape analysis with applications to facial expression recognition and human behavior un-

Mohamed Daoudi is a Full Professor of Computer Science at IMT Lille Douai and the Head of Image group at CRISTAL Laboratory (UMR CNRS 9189). He received his Ph.D. degree in Computer Engineering from the University of Lille 1 (France) in 1993 and Habilitation Diriger des Recherches from the University of Littoral (France) in 2000. His research interests include pattern recognition, shape analysis and computer vision. He has published over 150 papers in some of the most distinguished scientific journals

and international conferences. He is Associate Editor of *Image and Vision Computing Journal* and *IEEE Transactions on Multimedia*. He is a General co-Chair of IEEE International Conference on Automatic Face and Gesture Recognition 2019. He is Fellow of IAPR and IEEE Senior member.

Boulbaba Ben Amor is full professor of Computer Science with the Mines-Telecom Institute (IMT) Lille Douai and member of the UMR CNRS Research Center CRISTAL, in France. He is recipient of the prestigious Fulbright scholarship (2016-2017) and he is senior member of the IEEE, since 2015. He holds an Habilitation from University of Lille (France). He earned his Ph.D. in Computer Science from the Ecole Centrale de Lyon, in 2006. His research topics lie to 3D shape and their dynamics analysis for advanced

Stefano Berretti is an Associate Professor at the University of Florence. He received the Ph.D. in Computer Engineering in 2001. His research interests focus on image content modeling, indexing and retrieval, 3D computer vision for face biometrics, human emotion and behavior understanding. On these themes, he has published over 140 articles in some of the most distinguished conferences and journals in the area of pattern recognition, computer vision and multimedia. He is the Information Director of the ACM

Transactions on Multimedia Computing, Communications, and Applications, and Associate Editor of the *IET Computer Vision journal*.

Juan Carols Alvarez-Paiva is a Full Professor of Mathematics at University of Lille, France, and member of *Laboratoire Painlevé* (Lille 1/CNRS UMR 8524), since 2007. He received his Ph.D. degree in Mathematics in 1995 from Rutgers University, New Brunswick. His research interests lie in the area of differential geometry.