



HAL
open science

Modelling the Statistics of Cyclic Activities by Trajectory Analysis on the Manifold of Positive-Semi-Definite Matrices

Ettore Maria Celozzi, Luca Ciabini, Luca Cultrera, Pietro Pala, Stefano Berretti, Mohamed Daoudi, Alberto Del Bimbo

► **To cite this version:**

Ettore Maria Celozzi, Luca Ciabini, Luca Cultrera, Pietro Pala, Stefano Berretti, et al.. Modelling the Statistics of Cyclic Activities by Trajectory Analysis on the Manifold of Positive-Semi-Definite Matrices. 15th IEEE International Conference on Automatic Face and Gesture Recognition, Nov 2020, Buenos Aires, Argentina. hal-02863164

HAL Id: hal-02863164

<https://hal.science/hal-02863164>

Submitted on 16 Jun 2020

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

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

Modelling the Statistics of Cyclic Activities by Trajectory Analysis on the Manifold of Positive-Semi-Definite Matrices

Ettore Maria Celozzi¹, Luca Ciabini¹, Luca Cultrera¹, Pietro Pala¹, Stefano Berretti¹,
Mohamed Daoudi², Alberto Del Bimbo¹

¹ Media Integration and Communication Center, Dept. of Information Engineering, University of Florence, Italy

²IMT Lille Douai, Univ. Lille, CNRS, UMR 9189 CRIStAL, F-59000 Lille, France

Abstract—In this paper, a model is presented to extract statistical summaries to characterize the repetition of a cyclic body action, for instance a gym exercise, for the purpose of checking the compliance of the observed action to a template one and highlighting the parts of the action that are not correctly executed (if any). The proposed system relies on a Riemannian metric to compute the distance between two poses in such a way that the geometry of the manifold where the pose descriptors lie is preserved; a model to detect the begin and end of each cycle; a model to temporally align the poses of different cycles so as to accurately estimate the *cross-sectional* mean and variance of poses across different cycles. The proposed model is demonstrated using gym videos taken from the Internet.

I. INTRODUCTION

In the last decades, automatic analysis of human motion has been an active research topic, with applications that have been exploited in a number of different contexts, including video surveillance, semantic annotation of videos, entertainment, human computer interaction and home care rehabilitation, to say a few. Differences in body proportion (size, height, corpulence), body stiffness and training, influence the way different people perform an action. Even one same person is not able to perform the same action twice, exactly replicating the same sequence of body poses in space and time. This variability makes the task of human motion analysis very challenging [1], [2], [3], [4].

Recent CNN architectures have demonstrated real-time and reliable extraction of the coordinates of body joints from RGB streams. These advances make it possible to use a skeleton based body representation in a much broader range of operative contexts than before, not limited by the short operative range of RGB-D sensors, typically indoor and in the range of a few meters.

Many real situations do exist where computing a statistical description is of interest to characterize the execution of a particular action by a person or group of persons. For instance, extraction of such descriptors from physical activity training of athletes can be used to confirm or not the effects of some particular training exercise on the overall physical preparation of the athlete. In a different domain, the extraction of statistical descriptors from body motion can provide useful data to check the effectiveness and progress of smart rehabilitation programs. Imagine a patient that is undergoing a rehabilitation exercise at home. A smart assistant is capable of

recognizing the patient's actions, analyzing the correctness of the exercise, providing the patient with adequate feedback to improve the execution of the exercise, and sending statistical summaries of the rehabilitation activity to a remote server for inspection by a human therapist.

Analysis of body motion can address several different tasks, such as action recognition, motion prediction, and statistical characterization of body dynamics [5]. One of the first approaches to analyze the trajectories of tracked body joints was presented in Matikainen *et al.* [6] to address the task of action recognition. Their method used trajectories of tracked feature points in a bag of words paradigm. More recently, using the human skeleton as extracted from RGB-D images, Devanne *et al.* [7] 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 normalized and parameterized curve in \mathbb{R}^N . In the same direction, Su *et al.* [8] proposed a metric that considers time-warping on a Riemannian manifold, thus allowing the registration of trajectories and the computation of statistics on the trajectories. Similar ideas have been developed by Ben Amor *et al.* [9] on the Kendall's shape space with application to action recognition using rate-invariant analysis of skeletal shape trajectories. Anirudh *et al.* [10] started from the framework of Transported Square-Root Velocity Fields (TSRVF). 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 the temporal rate map to the same point. In [11], [12], Vemulapalli *et al.* proposed a Lie group trajectory representation of the skeletal data on a product space of special Euclidean (*SE*) groups. More recently, Kacem *et al.* [13] proposed a geometric approach for modeling and classifying dynamic 2D and 3D landmark sequences based on Gramian matrices derived from the static landmarks. This results in an affine-invariant representation of the data. Since Gramian matrices are positive-semidefinite, the authors rely on the geometry of the manifold of fixed-rank positive-semidefinite matrices, and more specifically, to the metric investigated in [14]. All the approaches described above rely on *hand-crafted features* enabling representation of the action as a trajectory or point in some suitable manifold. Differently from these approaches, many neural network models have been proposed that rely on training for the extraction of

deeply learned features [15], [16], [17], [18], [19]. Although these approaches result in state-of-the-art performance [19] on public action recognition benchmarks, it is not possible to define a formal mathematical framework to compute a valid metric on the internal, learned feature representation so as to perform a statistical analysis of the learned actions.

In this paper, a system is presented to extract statistical summaries to characterize the repetition of a cyclic body action for the purpose of checking the compliance of the observed action to a template one and highlighting the parts of the action that are not correctly performed (if any). Compared to previous work that focus on the task of action recognition and classification, the proposed solution aims to define a model for measuring the compliance of a user performed movement to a template movement and identify those parts of the movement that are not correctly executed. The proposed system relies on a Riemannian metric to compute the distance between two poses in such a way that the geometry of the manifold where the pose descriptors lie is preserved; a model to detect the begin and end of each cycle; a model to temporally align the poses of different cycles so as to accurately estimate the *cross-sectional* mean and variance of poses across different cycles.

II. BODY MOTION REPRESENTATION

We assume that the observed activity corresponds to the repetition of a cyclic action and our main goal is to extract a statistics, primarily in terms of mean and deviation from this mean, to describe how the action is performed across different cycles. The availability of these statistics makes it possible to compare the execution of a same training exercise by two different persons, aiming to identify the presence of parts of the exercise that are systematically performed in a different way by the two persons.

Adopting a skeleton based representation, the motion of a human body is represented by a sequence of body poses. Therefore, within each cycle c the action is represented by a trajectory $\Gamma_c : [0, 1] \mapsto \mathbb{R}^n$ in the space of skeleton poses, $\Gamma_c(0)$ and $\Gamma_c(1)$ corresponding to the poses at the beginning and end of the cycle, respectively, n the dimensionality of the skeleton pose representation. To model this process, descriptors and similarity measures are defined that account for *i*) how to represent a generic skeleton pose so as to guarantee invariance to a number of transforms; *ii*) how to measure the dissimilarity between two poses according to the adopted skeleton pose representation; *iii*) how to measure the dissimilarity between sequences of skeleton poses, so as to account for the different speed of execution of the same exercise across different cycles and across different persons.

A. Skeleton Pose Representation and Distance Measure

The most elementary representation of a body pose corresponds to the output of the skeleton body fitting module and is in the form of a $n \times d$ matrix Γ of coordinates of skeleton joints (where n is the number of joints, and d is 3 if the x, y, z coordinate of each joint is considered). This representation, that is provided by operating on RGB-D data

(like with the Microsoft Kinect) or on RGB data (*e.g.*, the OpenPose library [20]) may prove inadequate for the task of comparing two poses, especially if we want this comparison to be invariant to some major transforms such as rotation and translation of the skeleton. In this case, it results much more convenient to represent a body pose through the *Gram matrix* of joint coordinates:

$$G = \Gamma \Gamma^T \in \mathbb{R}^{n \times n}. \quad (1)$$

In fact, there is not a one-to-one correspondence between Gram matrices and skeleton joints matrices. Rather, any two sets of mean centered joint configurations that are the same up to an arbitrary rotation yield the same Gram matrix. Thus, the Gram matrix is representative of the equivalence class of mean centered skeleton joint matrices with respect to the rotation transform. By adopting a metric defined in the space of Gram matrices, it is possible to compute a rotation invariant dissimilarity measure between two skeleton poses.

Gram matrices as expressed in (1) do not form a linear vector space, rather they span the manifold of Positive-Semi-Definite (PSD) matrices of rank d , the PSD cone $\mathcal{S}^+(d, n)$. As a consequence, using the Frobenius norm for computing the distance between two Gram matrices may prove inadequate, because it would not conform to the geometry of the manifold. The Riemannian geometry of $\mathcal{S}^+(d, n)$ was studied in [14], [21], [22], [23], [24], [25], [26]. Accordingly, a measure of the geodesic distance between two Gram matrices G_i and G_j is expressed as:

$$\delta(G_i, G_j) = \left[\text{tr}(G_i) + \text{tr}(G_j) - 2 \text{tr} \left(\left(G_i^{\frac{1}{2}} G_j G_i^{\frac{1}{2}} \right)^{\frac{1}{2}} \right) \right]^{\frac{1}{2}}. \quad (2)$$

B. Similarity between Sequences of Skeleton Poses

Since it cannot be expected that the action is performed exactly at the same speed in different cycles the poses at the same normalized time t for two different cycles c_1 and c_2 of the same action may not correspond to each other. Thus, the computation of the *cross-sectional* mean and variance of poses at the same normalized time t , $\{\Gamma_1(t), \Gamma_2(t), \dots, \Gamma_C(t)\}$ would provide a fair statistical summary, with an artificially inflated variance. Actually, to provide a more accurate measure of the cross-sectional mean and variance, the different trajectories should be first aligned, that is, registered in time. In our approach, Dynamic Time Warping [27] is used to compute the re-parametrization function $\gamma : [0, 1] \mapsto [0, 1]$ that best aligns the poses $\Gamma_i(t)$ and $\Gamma_j(t)$ of two generic cycles. Actually, this poses some limitation being DTW not commutative and thus not a true metric.

III. THE VIRTUAL TRAINING COACH

The formal models defined above for computing the distance between poses and between sequences of poses have been used to design a *Virtual Training Coach* application that helps practitioners, athletes and people following some physical rehabilitation therapy, to execute a gymnastic exercise with better technique. In fact, many people train

themselves at home by using training coach videos on the Internet. As in the gym, the best way to achieve results and avoid pain is to match a good form of the exercise. However, at home, without a trainer or an expert that monitors the movements, it can be difficult to know if the body movement is correct. The Virtual Training Coach aims to help users to improve their technique by comparing their activity with the same activity performed by a trainer so as to highlight those parts of the exercise (if any) that are not correctly executed.

The application relies on the following main steps:

- The user performs the exercise in front of a camera. The OpenPose library [20] is used to extract, frame-by-frame, the coordinates of the body skeleton joints. Eventually, the sequence of all body skeleton data is stored. The same step is operated to process the video of the same exercise performed by a training coach.
- The sequence of body skeleton data represents the repetition of an elementary action several times. Thus, it is necessary to detect this elementary action so as to extract a statistics for each pose in it. This statistics is expressed in terms of mean and variance.
- The statistics extracted from the elementary actions performed by the user is compared to the statistics of the training coach to highlight those parts of the exercise performed by the user that are not correctly executed.

More details about these steps are provided in the following.

A. Extraction of Body Skeleton

OpenPose [20] performs pose estimation using part affinity fields, which are vectors that encode the position and orientation of limbs. For the experiments described in Section IV, the BODY_25 model was used. The output consists of lists containing the ordered x - and y -coordinates of model keypoints for each video frame. For each (x, y) coordinate a confidence value is also provided.

To make the skeleton representation invariant with respect to the frame resolution, the scale of the person in the video as well as the specific biometry of the person, the coordinates of skeleton joints are normalized. In particular, joint coordinates are subject to rigid translation so as to center the hip joint to the origin of the (X, Y) coordinate system. Afterwards, joint coordinates are subject to uniform scaling so as to have unitary torso length (this is the distance between the neck, the 1-st joint, and the mid-hip, the 8-th joint).

B. Detection of Action Cycles

The proposed model targets the extraction of a statistical summary from the observation of an activity corresponding to the cyclic repetition, multiple times, of a base elementary action. Hence, the sequence of skeleton poses extracted from a video is processed in order to detect this elementary action, whose repetition corresponds to the observed activity. For this purpose, the sub-sequence of the first T skeleton poses is used as template and it is compared, using a sliding window approach, to all the sub-sequences of length T of the activity. The following *cycle profile* function is computed:

$$CP_T(t) = \mathcal{D}(S_0^{T-1}, S_t^{t+T-1}), \quad (3)$$

being S_i^j the sequence of skeleton poses extracted from the i -th to the j -th frames and $\mathcal{D}(\cdot, \cdot)$ the distance function between two skeleton sequences. In order to detect the period P of the elementary action, local minima of the cycle profile function are detected and clustered, by minimum value, using a hierarchical clustering approach. The length P of the elementary action, that is the period of the activity, is computed based on the mean gap between consecutive minima of the selected cluster.

C. Alignment and Statistical Summaries

In order to extract statistical summaries to characterize the execution of the elementary action, different cycles have to be aligned with each other. This is accomplished by using Dynamic Time Warping to register in time the sequence of poses of the generic i -th cycle to those of the $(i+1)$ -th cycle, as described in Section II-B. Eventually, each pose in the first cycle is associated with its *pose-set* that groups all the corresponding poses in the next cycles. By using DTW for the alignment, differences in terms of speed of execution of the elementary action across different cycles are coped with, enabling proper computation of the cross-sectional mean and variance of corresponding poses.

For each pose-set the mean pose and the average deviation from this pose, using (2), are computed. Mean and deviation extracted from the pose sets are used to compare the execution of an exercise by a user against the execution of the same exercise by the training coach. Again, to perform this comparison the user pose sets to the training coach pose set have to be aligned. This is accomplished by using DTW to match the sequence of mean poses extracted from the user and training coach pose sets. In this way, given a generic pose extracted from the user activity it is possible to identify the corresponding user pose set, its mean pose μ_u and average deviation σ_u . Value of this deviation provides a quantitative index of the ability of the user to replicate the same pose in that part of the exercise, across different cycles. Given a generic user pose set it is also possible to identify the corresponding trainer pose set, its mean pose μ_t and average deviation σ_t . Value of this deviation provides a quantitative index of the ability of the trainer to replicate the same pose in that part of the exercise, across different cycles.

To measure the accuracy of execution of the user poses across different cycles both the distances between user vs trainer mean poses as well as the deviations from these mean poses can be considered. In particular, an *index of precision*, inversely proportional to the average deviation, is associated with each pose set. Given two corresponding pose sets, for the user and the training coach, the ratio between their indexes of precision is an indicator of the relative accuracy of execution of the poses and can be used to highlight those parts of the exercise the user should focus on in order to improve the technique. For this purpose, an *index of fit* is computed for each user mean pose as $f_i = p_i^t/p_i^u$, being p_i^t and p_i^u the values of the index of precision for the i -th pose set of the training coach and the user, respectively. The parts of the exercise that are performed by the user with the same

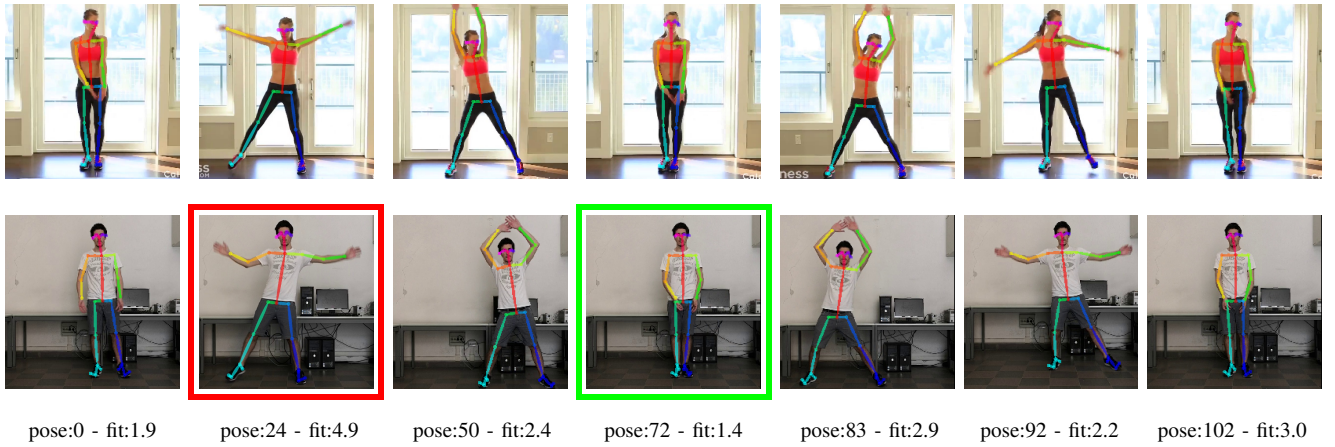


Fig. 1. Representative poses of one action performed by a training coach (top-row) and a user (bottom-row). The user poses characterized by the best and worst index of fit are highlighted with green and red colored frames, respectively.

precision as the training coach are associated with f_i values close to 1. Differently, the higher the f_i value the lower the user precision in the execution of the i -th pose.

IV. EXPERIMENTAL RESULTS

To demonstrate the potential and effectiveness of the proposed modeling framework, we downloaded some videos of fitness exercises from YouTube and used these videos as examples of exercises performed by a training coach. Then, for each one of these videos, we asked some users to replicate the exercise in the video.

Figure 1 shows (top-row) some representative poses of one action performed by a training coach and the corresponding poses (bottom-row), after alignment, of the same action replicated by a user. The corresponding values of the indexes of precision and the index of fit are shown in Figure 2.

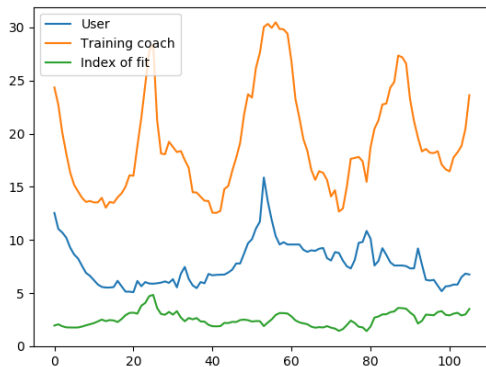


Fig. 2. Indexes of precision measuring the accuracy of execution for the trainer and the user shown in Figure 1. The index of fit is also shown.

The user poses characterized by the best and worst index of fit are highlighted with a colored frame, green for the best and red for the worst one. By the analysis of the index of fit across the action, the user can identify the parts of the exercise to focus on in order to improve the execution. For instance, a peak of the index of fit occurs in correspondence to the normalized time step 25. This corresponds to the pose

set with the worst index of fit: the training coach is able to replicate quite precisely the same pose across different cycles, whereas the user is not able to replicate these poses with the same precision. Figure 3 also compares these poses for the user and the training coach: The poses of the training coach are highly similar with each other; in contrast, a higher deviation is observed among the user poses.




Fig. 3. User (bottom-row) and trainer (top-row) poses corresponding to a bad index of fit.

V. CONCLUSION AND FUTURE WORK

In this paper, a model is described to extract statistical summaries for representing the accuracy of repetition of cyclic human activities. Computation of the statistical summaries is accomplished by modeling actions as trajectories in a Riemannian manifold and using a measure of the geodesic distance on this manifold to temporally align two trajectories. The proposed modeling framework is used to design a Virtual Training Coach application conceived to help users to improve their technique in the execution of a physical exercise. Future work will target a thorough experimental evaluation of the proposed model aiming to demonstrate its effectiveness and accuracy on a representative collection of gym videos, investigating the possible gap of accuracy that is observed comparing statistics extracted from 3D and 2D skeleton data.

REFERENCES

- [1] Avi Barliya, Lars Omlor, Martin A. Giese, Alain Berthoz, and Tamar Flash. Expression of emotion in the kinematics of locomotion. *Experimental Brain Research*, 225(2):159–176, 2013.
- [2] Frank E. Pollick, Helena M. Paterson, Armin Bruderlin, and Anthony J. Sanford. Perceiving affect from arm movement. *Cognition*, 82(2):B51–B61, 2001.
- [3] Samitha Herath, Mehrtash Harandi, and Fatih Porikli. Going deeper into action recognition: A survey. *Image and Vision Computing*, 2017.
- [4] Stefano Berretti, Mohamed Daoudi, Pavan Turaga, and Anup Basu. Representation, analysis and recognition of 3D humans: A survey. *ACM Trans. on Multimedia Computing, Communication and Applications*, 14(1s):1–35, March 2018.
- [5] Yu Kong and Yun Fu. Human action recognition and prediction: A survey, 2018.
- [6] Pyry Matikainen, Martial Hebert, and Rahul Sukthankar. Trajectons: Action recognition through the motion analysis of tracked features. In *Workshop on Video-Oriented Object and Event Classification (ICCV)*, September 2009.
- [7] Maxime Devanne, Hazem Wannous, Stefano Berretti, Pietro Pala, Mohamed Daoudi, and Alberto Del Bimbo. 3-D human action recognition by shape analysis of motion trajectories on Riemannian manifold. *IEEE Trans. on Cybernetics*, 45(7):1340–1352, 2015.
- [8] 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*, 8(1), 2014.
- [9] Boulbaba Ben Amor, Jingyong Su, and Anuj Srivastava. Action recognition using rate-invariant analysis of skeletal shape trajectories. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 38(1):1–13, 2016.
- [10] Rushil Anirudh, Pavan K. Turaga, Jingyong Su, and Anuj Srivastava. Elastic functional coding of riemannian trajectories. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 39(5):922–936, 2017.
- [11] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human action recognition by representing 3D skeletons as points in a Lie group. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 588–595, 2014.
- [12] Raviteja Vemulapalli and Rama Chellappa. Rolling rotations for recognizing human actions from 3d skeletal data. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4471–4479, 2016.
- [13] A. Kacem, M. Daoudi, B. Ben Amor, S. Berretti, and J. C. Alvarez-Paiva. A novel geometric framework on gram matrix trajectories for human behavior understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [14] Silvere Bonnabel and Rodolphe Sepulchre. Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1055–1070, 2009.
- [15] Wentao Zhu, Cuiling Lan, Junliang Xing, Wenjun Zeng, Yanghao Li, Li Shen, and Xiaohui Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *AAAI Conference on Artificial Intelligence*, 2016.
- [16] Chao Li, Qiaoyong Zhong, Di Xie, and Shiliang Pu. Skeleton-based action recognition with convolutional neural networks. In *IEEE International Conference on Multimedia & Expo Workshops*, page 597–600, 2017.
- [17] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI Conference on Artificial Intelligence*, 2018.
- [18] Bin Li, Xi Li, Zhongfei Zhang, and Fei Wu. Spatio-temporal graph routing for skeleton-based action recognition. In *AAAI Conference on Artificial Intelligence*, 2019.
- [19] Xiang Gao, Wei Hu, Jiayang Tang, Jiaying Liu, and Zongming Guo. Optimized skeleton-based action recognition via sparsified graph regression. *arXiv:1811.12013*, 2019.
- [20] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. *arXiv preprint arXiv:1812.08008*, 2018.
- [21] M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM Journal on Optimization*, 20(5):2327–2351, 2010.
- [22] Bart Vandereycken, P-A Absil, and Stefan Vandewalle. Embedded geometry of the set of symmetric positive semidefinite matrices of fixed rank. In *Statistical Signal Processing, 2009. SSP'09. IEEE/SP 15th Workshop on*, pages 389–392. IEEE, 2009.
- [23] Bart Vandereycken, P.-A. Absil, and Stefan Vandewalle. A Riemannian geometry with complete geodesics for the set of positive semidefinite matrices of fixed rank. *IMA Journal of Numerical Analysis*, 33(2):481–514, 2013.
- [24] E. Massart and P.-A. Absil. Quotient geometry with simple geodesics for the manifold of fixed-rank positive-semidefinite matrices. Preprint.
- [25] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. Image set classification by symmetric positive semi-definite matrices. In *Applications of Computer Vision (WACV), 2016 IEEE Winter Conference on*, pages 1–8. IEEE, 2016.
- [26] Gilles Meyer, Silvere Bonnabel, and Rodolphe Sepulchre. Regression on fixed-rank positive semidefinite matrices: a Riemannian approach. *Journal of Machine Learning Research*, 12(Feb):593–625, 2011.
- [27] C. S. Myers and L. R. Rabiner. A comparative study of several dynamic time-warping algorithms for connected-word recognition. *Bell System Technical Journal*, 60(7):1389–1409, 1981.