

MULTI-SHOT PERSON RE-IDENTIFICATION VIA RELATIONAL STEIN DIVERGENCE

Azadeh Alavi, Yan Yang, Mehrtash Harandi, Conrad Sanderson

NICTA, GPO Box 2434, Brisbane, QLD 4001, Australia
University of Queensland, School of ITEE, QLD 4072, Australia
Queensland University of Technology, Brisbane, QLD 4000, Australia

ABSTRACT

Person re-identification is particularly challenging due to significant appearance changes across separate camera views. In order to re-identify people, a representative human signature should effectively handle differences in illumination, pose and camera parameters. While general appearance-based methods are modelled in Euclidean spaces, it has been argued that some applications in image and video analysis are better modelled via non-Euclidean manifold geometry. To this end, recent approaches represent images as covariance matrices, and interpret such matrices as points on Riemannian manifolds. As direct classification on such manifolds can be difficult, in this paper we propose to represent each manifold point as a vector of similarities to class representers, via a recently introduced form of Bregman matrix divergence known as the Stein divergence. This is followed by using a discriminative mapping of similarity vectors for final classification. The use of similarity vectors is in contrast to the traditional approach of embedding manifolds into tangent spaces, which can suffer from representing the manifold structure inaccurately. Comparative evaluations on benchmark ETHZ and iLIDS datasets for the person re-identification task show that the proposed approach obtains better performance than recent techniques such as Histogram Plus Epitome, Partial Least Squares, and Symmetry-Driven Accumulation of Local Features.

Index Terms— surveillance, person re-identification, manifolds.

1. INTRODUCTION

Person re-identification is the process of matching persons across non-overlapping camera views in diverse locations. Within the context of surveillance, re-identification needs to function with a large set of candidates and be robust to pose changes, occlusions of body parts, low resolution and illumination variations. The issues can be compounded, making a person difficult to recognise even by human observers (see Fig. 1 for examples). Compared to classical biometric cues (eg. face, gait) which may not be reliable due to non-frontality, low resolution and/or low frame-rate, person re-identification approaches typically use the entire body.

While appearance based person re-identification methods are generally modelled in Euclidean spaces [8, 11, 24], it has been argued that some applications in image and video analysis are better modelled on non-Euclidean manifold geometry [28]. To this end, recent approaches represent images as covariance matrices [3], and interpret such matrices as points on Riemannian manifolds [12, 28]. A popular way of analysing manifolds is to embed them into tangent spaces, which are Euclidean spaces. This process which can be interpreted as warping the feature space [27]. Embedding manifolds is not without problems, as pairwise distances between arbitrary points on a tangent space may not represent the structure of the manifold accurately [12, 13].




Fig. 1. Examples of challenges in person re-identification, where each column contains images of the same person from two separate camera views. Challenges include pose changes, occlusions of body parts, low resolution and illumination variations.

In this paper we present a multi-shot appearance based person re-identification method on Riemannian manifolds, where embedding the manifolds into tangent spaces is not required. We adapt a recently proposed technique for analysing Riemannian manifolds, where points on the manifolds are represented through their similarity vectors [2]. The similarity vectors contain similarities to class representers. We obtain each similarity with the aid of a recently introduced form of Bregman matrix divergence known as the Stein divergence [13, 25]. The classification task on manifolds is hence converted into a task in the space of similarity vectors, which can be tackled using learning methods devised for Euclidean spaces, such as Linear Discriminant Analysis [5]. Unlike previous person re-identification methods, the proposed method does not require separate settings for new datasets.

We continue the paper as follows. In Section 2 several recent methods for person re-identification are briefly described. The proposed approach is detailed in Section 3. A comparative performance evaluation on two public datasets is given in Section 4. The main findings are summarised in Section 5.

2. PREVIOUS WORK

Given an image of an individual to be re-identified, the task of person re-identification can be categorised into two main classes. (i) Single-vs-Single (SvS), where there is only one image of each person in the gallery and one in the probe; this can be seen as a one-to-one comparison. (ii) Multiple-vs-Single (MvS), or multi-shot, where there are multiple images of each person available in gallery and one image in the probe. Below we summarise several person re-identification methods: Partial Least Squares (PLS) [24], Context based method [31], Histogram Plus Epitome (HPE) [4], and Symmetry-Driven Accumulation of Local Features (SDALF) [8].

The PLS method [24] first decomposes a given image into overlapping blocks, and extracts a rich set of features from each block. Three types of features are considered: textures, edges, and colours. The dimensionality of the feature space is then reduced by employing Partial Least Squares regression (PLSR) [30], which models relations between sets of observed variables by means of latent variables. To learn a PLSR discriminatory model for each person, one-against-all scheme is used [9]. Nearest neighbour is then employed for classification.

The Context-based method [31] enriches the description of a person by contextual visual knowledge from surrounding people. The method represents a group by considering two descriptors: **(a)** ‘center rectangular ring ratio-occurrence’ descriptor, which describes the information ratio of visual words between and within various rectangular ring regions, and **(b)** ‘block based ratio-occurrence’ descriptor, which describes local spatial information between visual words that could be stable. For group image representation only features extracted from foreground pixels are used to construct visual words.

HPE [4] considers multiple instances of each person to create a person signature. The structural element (STEL) generative model approach [16] is employed for foreground detection. The combination of a global (person level) HSV histogram and epitome regions of foreground pixels is then calculated, where an image epitome [15] is computed by collapsing the given image into a small collage of overlapped patches. The patches contain the essence of textural, shape and appearance properties of the image. Both the generic epitome (epitome mean) and local epitome (probability that a patch is in an epitome) are computed.

SDALF [8] considers multiple instances of each person. Foreground features are used to model three complementary aspects of human appearance extracted from various body parts. First, for each pedestrian image, axes of asymmetry and symmetry are found. Then, complementary aspects of the person appearance are detected on each part, and their features are extracted. To select salient parts of a given pedestrian image, the features are then weighted by exploiting perceptual principles of symmetry and asymmetry.

The above methods assume that classical Euclidean geometry is capable of providing meaningful solutions (distances and statistics) for modelling and analysing images and videos, which might not be always correct [27]. Furthermore, they require separate parameter tuning for each dataset.

3. PROPOSED APPROACH

Our goal is to automatically re-identify a given person among a large set of candidates in diverse locations over various non-overlapping camera views. The proposed method is comprised of three main stages: **(i)** feature extraction and generation of covariance descriptors, **(ii)** measurement of similarities on Riemannian manifolds via the Stein divergence, and **(iii)** creation of similarity vectors and discriminative mapping for final classification. Each of the stages is elucidated in more detail in the following subsections.

3.1. Feature Extraction and Covariance Descriptors

As per [4, 8], to reduce the effect of varying background, foreground pixels are extracted from each given image of a person via the STEL generative model approach [16]. We note that it is also possible to use more advanced approaches, such as [21].

Based on preliminary experiments, for each each foreground pixel located at (x, y) , the following feature vector is calculated:

$$\mathbf{f} = [x, y, HSV_{xy}, CIELAB_{xy}, \Lambda_{xy}, \Theta_{xy}]^T \quad (1)$$

where $HSV_{xy} = [H_{xy}, S_{xy}, \hat{V}_{xy}]$ are the colour values of the HSV channels, employing histogram equalisation for channel V , $CIELAB_{xy} = [L_{xy}, a_{xy}, b_{xy}]$ are the values of CIELAB colour space [1], while $\Lambda_{xy} = [\lambda_{xy}^R, \lambda_{xy}^G, \lambda_{xy}^B]$ and $\Theta_{xy} = [\theta_{xy}^R, \theta_{xy}^G, \theta_{xy}^B]$ indicate gradient magnitudes and orientations for each channel in RGB colour space. We note that we have selected this relatively straightforward set of features as a starting point, and that it is certainly possible to use other features. However, a thorough evaluation of possible features is beyond the scope of this paper.

Given a set $F = \{\mathbf{f}_i\}_{i=1}^N$ of extracted features, with its mean represented by $\boldsymbol{\mu}$, each image is represented as a covariance matrix:

$$\mathbf{C} = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{f}_i - \boldsymbol{\mu})(\mathbf{f}_i - \boldsymbol{\mu})^T \quad (2)$$

Representing an image with a covariance matrix has several advantages [3]: **(i)** it is a low-dimensional (compact) representation that is independent of image size, **(ii)** the impact of noisy samples is reduced via the averaging during covariance computation, and **(iii)** it is a straightforward method of fusing correlated features.

3.2. Riemannian Manifolds and Stein Divergence

Covariance matrices belong to the group of symmetric positive definite (SPD) matrices, which can be interpreted as points on Riemannian manifolds. As such, the underlying distance and similarity functions might not be accurately defined in Euclidean spaces [23].

Efficiently handling Riemannian manifolds is non-trivial, due largely to two main challenges [26]: **(i)** as manifold curvature needs to be taken into account, defining divergence or distance functions on SPD matrices is not straightforward; **(ii)** high computational requirements, even for basic operations such as distances. For example, the Riemannian structure induced by considering the Affine Invariant Riemannian Metric (AIRM) has been shown somewhat useful for analysing SPD matrices [14, 20]. For $\mathbf{A}, \mathbf{B} \in \mathcal{S}_{++}^d$, where \mathcal{S}_{++}^d is the space of positive definite matrices of size $d \times d$, AIRM is defined as:

$$\delta_R(\mathbf{A}, \mathbf{B}) := \left\| \log \left(\mathbf{B}^{-\frac{1}{2}} \mathbf{A} \mathbf{B}^{-\frac{1}{2}} \right) \right\|_F \quad (3)$$

where $\log(\cdot)$ is the principal matrix logarithm [25]. However, AIRM is computationally demanding as it essentially needs eigen-decomposition of \mathbf{A} and \mathbf{B} . Furthermore, the resulting structure has negative curvature which prevents the use of conventional learning algorithms for classification purposes.

To simplify the handling of Riemannian manifolds, they are often first embedded into higher dimensional Euclidean spaces, such as tangent spaces [18, 19, 22, 29]. However, only distances between points to the tangent pole are equal to true geodesic distances, meaning that distances between arbitrary points on tangent spaces may not represent the manifold accurately.

As an alternative to measuring distances on tangent spaces, in this work we use the recently introduced Stein divergence, which is a version of the Bregman matrix divergence for SPD matrices [25]. To measure dissimilarity between two SPD matrices \mathbf{A} and \mathbf{B} , the Bregman divergence is defined as [17]:

$$D_\phi(\mathbf{A}, \mathbf{B}) \triangleq \phi(\mathbf{A}) - \phi(\mathbf{B}) - \langle \nabla_\phi(\mathbf{B}), \mathbf{A} - \mathbf{B} \rangle \quad (4)$$

where $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}^T \mathbf{B})$ and $\phi : \mathcal{S}_{++}^d \rightarrow \mathbb{R}$ is a real-valued, strictly convex and differentiable function. The divergence in (4) is asymmetric which is often undesirable. The Jensen-Shannon symmetrisation of Bregman divergence is defined as [17]:

$$D_\phi^{JS}(\mathbf{A}, \mathbf{B}) \triangleq \frac{1}{2} D_\phi\left(\mathbf{A}, \frac{\mathbf{A} + \mathbf{B}}{2}\right) + \frac{1}{2} D_\phi\left(\mathbf{B}, \frac{\mathbf{A} + \mathbf{B}}{2}\right) \quad (5)$$

By selecting ϕ in (5) to be $-\log(\det(\mathbf{A}))$, which is the barrier function of semi-definite cone [25], we obtain the symmetric Stein divergence, also known as the Jensen Bregman Log-Det divergence [6]:

$$J_\phi(\mathbf{A}, \mathbf{B}) \triangleq \log\left(\det\left(\frac{\mathbf{A} + \mathbf{B}}{2}\right)\right) - \frac{1}{2} \log(\det(\mathbf{AB})) \quad (6)$$

The symmetric Stein divergence is invariant under congruence transformations and inversion [6]. It is computationally less expensive than AIRM, and is related to AIRM in several aspects which establish a bound between the divergence and AIRM [6].

3.3. Similarity Vectors and Discriminative Mapping

For each query point (an SPD matrix) to be classified, a similarity to each training class is obtained, forming a similarity vector. We obtain each similarity with the aid of the Stein divergence described in the preceding section. The classification task on manifolds is hence converted into a task in the space of similarity vectors, which can be tackled using learning methods devised for Euclidean spaces.

Given a training set of points on a Riemannian manifold, $\mathbb{X} = \{(\mathbf{X}_1, y_1), (\mathbf{X}_2, y_2), \dots, (\mathbf{X}_n, y_n)\}$, where $y_i \in \{1, 2, \dots, m\}$ is a class label, and m is the number of classes, we define the similarity between matrix \mathbf{X}_i and class l as:

$$s_{i,l} = \frac{1}{N_l} \sum_{j \neq i} J_\phi(\mathbf{X}_i, \mathbf{X}_j) \delta(y_j - l) \quad (7)$$

where $\delta(\cdot)$ is the discrete Dirac function and

$$N_l = \begin{cases} n_l - 1 & \text{if } y_i = l \\ n_l & \text{otherwise} \end{cases} \quad (8)$$

where n_l is the number of training matrices in class l . Using Eqn. (7), the similarity between \mathbf{X}_i and all classes is obtained, where $i \in \{1, 2, \dots, n\}$. Each matrix \mathbf{X}_i is hence represented by a similarity vector:

$$\mathbf{s}_i = [s_{i,1}, s_{i,2}, \dots, s_{i,m}]^T \quad (9)$$

Classification on Riemannian manifolds can now be reinterpreted as a learning task in \mathbb{R}^m . Given the similarity vectors of training data, $\mathbb{S} = \{(\mathbf{s}_1, y_1), (\mathbf{s}_2, y_2), \dots, (\mathbf{s}_n, y_n)\}$, we seek a way to label a query matrix \mathbf{X}_q , represented by a similarity vector $\mathbf{s}_q = [s_{q,1}, s_{q,2}, \dots, s_{q,m}]^T$. As a starting point, we have chosen linear discriminant analysis [5], where we find a mapping \mathbf{W}^* that minimises the intra-class distances while simultaneously maximising inter-class distances:

$$\mathbf{W}^* = \underset{\mathbf{W}}{\operatorname{argmax}} \operatorname{trace} \left\{ \left[\mathbf{W} \mathbf{S}_W \mathbf{W}^T \right]^{-1} \left[\mathbf{W} \mathbf{S}_B \mathbf{W}^T \right] \right\} \quad (10)$$

where \mathbf{S}_B and \mathbf{S}_W are the between class and within class scatter matrices [5]. The query similarity vector \mathbf{s}_q can then be mapped into the new space via:

$$\mathbf{x}_q = \mathbf{W}^{*T} \mathbf{s}_q \quad (11)$$

We can now use a straightforward nearest neighbour classifier [5] to assign a class label to \mathbf{x}_q . We shall refer to this approach as Relational Divergence Classification (RDC).

4. EXPERIMENTS AND DISCUSSION

In this section we evaluate the proposed RDC approach by providing comparisons against several methods on two person re-identification datasets: iLIDS [31] and ETHZ [7, 24]. The VIPeR dataset [10] was not used as it only has one image from each person in the gallery, and is hence not suitable for testing MvS approaches. Each dataset covers various aspects and challenges of the person re-identification task. The results are shown in terms of the Cumulative Matching Characteristic (CMC) curves, where each CMC curve represents the expectation of finding the correct match in the top n matches.

In order to show the improvement caused by using similarity vectors in conjunction with linear discriminant analysis, we also evaluate the performance of directly using the Stein divergence in conjunction with a nearest neighbour classifier (ie. direct classification on manifolds, without creating similarity vectors). We refer to this approach as the *direct Stein* method.

4.1. iLIDS Dataset

The iLIDS dataset is a publicly available video dataset capturing real scenarios at an airport arrival hall under a multi-camera CCTV network. From these videos a dataset of 479 images of 119 pedestrians was extracted and the images were normalised to 128×64 pixels (height \times width) [31]. The extracted images were chosen from non-overlapping cameras, and are subject to illumination changes and occlusions [31].

We randomly selected N images for each person to build the gallery set, while the remaining images form the probe set. The whole procedure is repeated 10 times in order to estimate an average CMC curve. We compared the performance of the proposed RDC approach against the direct Stein method, as well as the algorithms described in Section 2 (SDALF and Context based) for a commonly used setting of $N = 3$. The results, shown in Fig. 2, indicate that the proposed method generally outperforms the other techniques. The results also show that the use of similarity vectors in conjunction with linear discriminant analysis is preferable to directly using the Stein divergence.




Fig. 2. Performance on the iLIDS dataset [31] for $N=3$, using the proposed RDC method, the direct Stein method, SDALF [8], context based method [31]. HPE results for $N=3$ were not provided in [4].

4.2. ETHZ Dataset

The ETHZ dataset [7, 24] was captured from a moving camera, with the images of pedestrians containing occlusions and wide variations in appearance. Sequence 1 contains 83 pedestrians (4857 images), Sequence 2 contains 35 pedestrians (1936 images), and Sequence 3 contains 28 pedestrians (1762 images).

We downsampled all the images to 64×32 (height \times width). For each subject, the training set consisted of N randomly selected images, with the rest used for the test set. The random selection of the training and testing data was repeated 10 times.

Results were obtained for the commonly used setting of $N = 10$ and are shown in Fig. 3. On sequences 1 and 2, the proposed RDC method considerably outperforms PLS, SDALF, HPE and the direct Stein method. On sequence 3, RDC obtains performance on par with SDALF.

Note that the random selection used by the RDC approach to create the gallery is more challenging and more realistic than the data selection strategy employed by SDALF and HPE on the same dataset [4, 8]. SDALF and HPE both apply clustering beforehand on the original frames, and then select randomly one frame for each cluster to build their gallery set. In this way they can ensure that their gallery set includes the keyframes to use for the multi-shot signature calculation. In contrast, we haven't applied any clustering for the proposed RDC method in order to be closer to real life scenarios.

5. CONCLUSION

We have proposed a novel appearance based person re-identification method comprised of: (i) representing each image as a compact covariance matrix constructed from feature vectors extracted from foreground pixels, (ii) treating covariance matrices as points on Riemannian manifolds, (iii) representing each manifold point as a vector of similarities to class representers with the aid of the recently introduced Stein divergence, and (iv) using a discriminative mapping of similarity vectors for final classification. The use of similarity vectors is in contrast to the traditional approach of analysing manifolds via embedding them into tangent spaces. The latter might result in inaccurate modelling, as the structure of the manifolds is only partially taken into account [12, 13].

Person re-identification experiments on the iLIDS [31] and ETHZ [7, 24] datasets show that the proposed approach outperforms several recent methods, such as Histogram Plus Epitome [4], Partial Least Squares [24], and Symmetry-Driven Accumulation of Local Features [8].

6. ACKNOWLEDGEMENTS

NICTA is funded by the Australian Government as represented by the *Department of Broadband, Communications and the Digital Economy*, as well as the Australian Research Council through the *ICT Centre of Excellence* program.




Fig. 3. Performance on the ETHZ dataset [24] for $N = 10$, using Sequences 1 to 3 (top to bottom). Results are shown for the proposed RDC method, direct Stein method, HPE [4], PLS [24] and SDALF [8].

7. REFERENCES

- [1] T. Acharya and A. K. Ray. *Image Processing: Principles and Applications*. 2005.
- [2] A. Alavi, M. T. Harandi, and C. Sanderson. Relational divergence based classification on Riemannian manifolds. In *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 111–116, 2013.
- [3] C. Anoop, M. Vassilios, and P. Nikolaos. Dirichlet process mixture models on symmetric positive definite matrices for appearance clustering in video surveillance applications. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 3417–3424, 2011.
- [4] L. Bazzani, M. Cristani, A. Perina, M. Farenzena, and V. Murino. Multiple-shot person re-identification by HPE signature. In *Int. Conf. Pattern Recognition (ICPR)*, pages 1413–1416, 2010.
- [5] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [6] A. Cherian, S. Sra, A. Banerjee, and N. Papanikolopoulos. Efficient similarity search for covariance matrices via the Jensen-Bregman LogDet divergence. In *Int. Conf. Computer Vision (ICCV)*, pages 2399–2406, 2011.
- [7] A. Ess, B. Leibe, and L. Van Gool. Depth and appearance for mobile scene analysis. *Int. Conf. Computer Vision (ICCV)*, pages 1–8, 2007.
- [8] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. *IEEE Conf. Computer Vision and Pattern Recognition*, pages 2360–2367, 2010.
- [9] P. Geladi and B. Kowalski. Partial least-squares regression: a tutorial. *Analytica Chimica Acta*, 185:1–17, 1986.
- [10] D. Gray, S. Brennan, and H. Tao. Evaluating appearance models for recognition, reacquisition, and tracking. In *Proc. IEEE International Workshop on Performance Evaluation for Tracking and Surveillance (PETS)*, volume 3, page 5, 2007.
- [11] D. Gray and H. Tao. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *Computer Vision—ECCV 2008, Lecture Notes in Computer Science*, volume 5302, pages 262–275, 2008.
- [12] M. Harandi, C. Sanderson, A. Wiliem, and B. Lovell. Kernel analysis over Riemannian manifolds for visual recognition of actions, pedestrians and textures. *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 433–439, 2012.
- [13] M. T. Harandi, C. Sanderson, R. Hartley, and B. C. Lovell. Sparse coding and dictionary learning for symmetric positive definite matrices: A kernel approach. In *European Conference on Computer Vision (ECCV), Lecture Notes in Computer Science (LNCS)*, volume 7573, pages 216–229, 2012.
- [14] T. Hou and H. Qin. Efficient computation of scale-space features for deformable shape correspondences. *European Conference in Computer Vision (ECCV)*, pages 384–397, 2010.
- [15] N. Jovic, B. J. Frey, and A. Kannan. Epitomic analysis of appearance and shape. *IEEE International Conference on Computer Vision*, 1:34–41, 2003.
- [16] N. Jovic, A. Perina, M. Cristani, V. Murino, and B. Frey. STEL component analysis: Modeling spatial correlations in image class structure. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2044–2051, 2009.
- [17] B. Kulis, M. Sustik, and I. Dhillon. Low-rank kernel learning with Bregman matrix divergences. *The Journal of Machine Learning Research*, 10:341–376, 2009.
- [18] Y. Lui. Tangent bundles on special manifolds for action recognition. *IEEE Trans. Circuits and Systems for Video Technology*, 22(6):930–942, 2011.
- [19] F. Porikli, O. Tuzel, and P. Meer. Covariance tracking using model update based on Lie algebra. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 728–735, 2006.
- [20] D. Raviv, A. Bronstein, M. Bronstein, R. Kimmel, and N. Sochen. Affine-invariant geodesic geometry of deformable 3d shapes. *Computers & Graphics*, 35(3):692–697, 2011.
- [21] V. Reddy, C. Sanderson, and B. C. Lovell. Improved foreground detection via block-based classifier cascade with probabilistic decision integration. *IEEE Transactions on Circuits and Systems for Video Technology*, 23(1):83–93, 2013.
- [22] A. Sanin, C. Sanderson, M. T. Harandi, and B. C. Lovell. K-tangent spaces on Riemannian manifolds for improved pedestrian detection. In *IEEE International Conference on Image Processing (ICIP)*, pages 473–476, 2012.
- [23] A. Sanin, C. Sanderson, M. T. Harandi, and B. C. Lovell. Spatio-temporal covariance descriptors for action and gesture recognition. In *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 103–110, 2013.
- [24] W. Schwartz and L. Davis. Learning discriminative appearance-based models using partial least squares. In *Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI)*, pages 322–329, 2009.
- [25] S. Sra. Positive definite matrices and the symmetric Stein divergence. *Preprint: [arXiv:1110.1773]*, 2012.
- [26] S. Sra and A. Cherian. Generalized dictionary learning for symmetric positive definite matrices with application to nearest neighbor retrieval. *Machine Learning and Knowledge Discovery in Databases*, 6913:318–332, 2011.
- [27] P. Turaga, A. Veeraraghavan, and R. Chellappa. Statistical analysis on Stiefel and Grassmann manifolds with applications in computer vision. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008.
- [28] P. 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 and Machine Intelligence*, 33(11):2273–2286, 2011.
- [29] A. Veeraraghavan, A. Roy-Chowdhury, and R. Chellappa. Matching shape sequences in video with applications in human movement analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27(12):1896–1909, 2005.
- [30] H. Wold, S. Kotz, and N. Johnson. Partial least squares. *Encyclopedia of Statistical Sciences*, 6:581–591, 1985.
- [31] W. Zheng, S. Gong, and T. Xiang. Associating groups of people. In *British Machine Vision Conference*, volume 1, pages 1–11, 2009.