

Comparative Study of the Behavior of Feature Reduction Methods in Person Re-identification Task

Bahram Lavi¹, Mehdi Fatan Serj² and Domenec Puig Valls²

¹Dept. of Electrical and Electronic Engineering, University of Cagliari, Piazza d'Armi, 09123, Cagliari, Italy

²Dept. of Computer Engineering and Maths School of Engineering, Universitat Rovira i Virgili, Tarragona, Spain

Keywords: Person Re-Identification, Video Surveillance Systems, Dimensional Reduction Methods.

Abstract: One of the goals of person re-identification systems is to support video-surveillance operators and forensic investigators to find an individual of interest in videos acquired by a network of non-overlapping cameras. This is attained by sorting images of previously observed individuals for decreasing values of their similarity with a given probe individual. Existing appearance descriptors, together with their similarity measures, are mostly aimed at improving ranking quality. Many of these descriptors generate a high feature vector represented as an image signature. To tackle person re-identification in real-world scenario the processing time will be crucial, so an individual of interest within a network camera should be found out swiftly. We therefore study some feature reduction methods to achieve a significant trade-off between processing time and ranking quality. Although, observing some redundancies on the generated patterns of a given descriptor are not deniable, we suggest to employ a feature reduction method before use of it in real-world scenarios. In particular, we have tested three reduction methods: PCA, KPCA, and Isomap. We then evaluate our study on two benchmark data sets (VIPeR, and i-LIDS), by using two state-of-the-art descriptors on person re-identification task. The results presented in this paper, after applying the feature reduction step, are very promising in terms of recognition rate.

1 INTRODUCTION

Person re-identification is a computer vision task consisting of recognizing an individual who had previously been observed over a network of cameras with non-overlapping fields of view (Bedagkar-Gala and Shah, 2014). One of its applications consists of supporting video surveillance operators and forensic investigators in retrieving all the videos showing an individual of interest, given an image of him/her as a query (*aka probe*). In this application scenario, the goal of a person re-identification system is returning to the user the frames or videos of all the individuals recorded by the camera network (*aka template gallery*) sorted for decreasing similarity to the probe, so that the user can find the occurrences (if any) of the individual of interest, ideally in the first positions. This task is challenging due to several issues typical of video surveillance footage, like low resolution, unconstrained pose, illumination changes, and occlusions, which do not allow to exploit strong biometrics like face. Clothing appearance is therefore one of the most widely used cues. Other cues like gait and anthropo-

metric measures have also been investigated.

Many of the existing similarity measures (either hand-crafted or learnt from data) are indeed rather complex, and require a relatively high processing time, e.g., (Farenzena et al., 2010; Ma et al., 2014; Liao et al., 2015). On the other hand, in real-world applications the template gallery can be very large, and even if the processing time for a single matching score is low (e.g., the Euclidean distance between fixed-length feature vectors (Ma et al., 2014)), evaluating the matching scores for all the templates can be time-consuming.

During the past few years, many descriptors have been proposed in the literature based on clothing appearance. The existing descriptors are typically constructed based on either colour information, texture information, or combination of both. Despite the differences among them, the final descriptor is typically generated in high-dimension feature-size. In some cases, the operator intends to attain the result much faster, because of existing too many templates to be checked by him/her. We clearly discussed this issue in (Lavi et al., 2016) by proposing multi-stage sy-

stem to attain a significant trade-off between processing time with the expense of low recognition rate. To construct such method, one should consider on couple of simplified versions of a *given* descriptor to obtain a faster version. As a simple guideline, that we suggested, was to employ a feature reduction method on this purpose. However, apart from existence of many feature reduction methods, at this paper, we aim to study the behavior of some of them in person re-identification problem.

In this paper, three of the most popular feature reduction methods in machine learning is presented, and a comparison among all of them is constructed for person re-identification task. These three methods are principal component analysis (PCA), kernel principal component analysis (KPCA), and Isomap. PCA (Hotelling, 1933) is a well-known method in terms of compressing data pattern which consists of calculating the Eigenvectors of the covariance matrix of the original feature space, and describe the variation of a set of variables in terms of a reduced set of uncorrelated linear space of such variables with maximum variance (*aka* principal components (PCs)). KPCA (Schölkopf et al., 1999) is the nonlinear version of PCA in which the original feature space is mapped to a higher-dimensional feature space using a kernel function, and then PCA is calculated. Isomap (*aka* isometric feature mapping) is popular in terms of computing quasi-isometric from a high-dimensional feature space to a low-dimensional feature space. Isomap is highly efficient and applicable to a wide range of data points and dimensions (Balasubramanian and Schwartz, 2002). The isometric feature space can be supposed as a kernel function and so this method can also be known as a type of KPCA technique.

This paper is structured as follows. We first summarize related work in Sect. 2. In Sect. 3, we define the problem by explaining a standard person re-identification system. Then, in Sect. 4, we explain the employed feature reduction methods in our study. Finally, in Sect. 5, we present our experimental results and discuss on behaviour the feature reduction methods by using two well-known descriptors in person re-identification task on two benchmark data sets.

2 BACKGROUND

In this section we first describe some standard person re-identification techniques. We then summarize person re-identification techniques aimed at reducing processing time in the computation of matching scores.

2.1 Overview on Person Re-Identification Methods

The recent state-of-the-art around descriptor for person re-identification task. These descriptors can be categorised by considering their cue on generating image signature; colour information (Cheng et al., 2011; Dikmen et al., 2010), and many of existing works used as combination of different colour and texture information to help in attaining a better performance (Farenzena et al., 2010; Prosser et al., 2010; Ma et al., 2014; Zheng et al., 2013). For person re-identification task, both processing time and recognition accuracy play an important role to tackle this problem in online applications. In this respect, some existing approaches are time consuming due to computation of matching scores for a target image against the templates in the gallery set. This may be caused of large number of feature elements generated as an image signature. Farenzena et al. proposed (Farenzena et al., 2010) symmetry driven accumulation of local features (SDALF) descriptor which consists of subdividing body into four parts: left and right, torso and legs. Three kinds of features are extracted from each part: maximally stable color regions (MSCR), i.e., elliptical regions (blobs) exhibiting distinct color patterns (their number depends on the specific image); a weighted HSV color histogram (wHSV); and recurrent high-structured patches (RHSP) that characterize texture. A specific similarity measure is defined for each feature; the matching score is computed as their linear combination. Despite from the small feature size vectors generated by a descriptor, the whole processes for such three descriptors are very time consuming. The final descriptor in terms of the feature-size is variant for different individuals. In (Satta et al., 2012), a dissimilarity-based approach proposed to design descriptors made up of bags of local features, possibly extracted from different body parts. It consists in finding a set of M representative local features (called prototypes) from all individuals of the template gallery, and in representing each template and probe image as a vector of M dissimilarity values between the corresponding bag of local features and the templates. This allows the matching score to be computed as a distance between feature vectors, rather than using a more complex similarity measure between bags of local features.

2.2 Reducing Processing Time in Person Re-Identification

To our knowledge, the issue of processing time has been explicitly addressed so far in the context of per-

son re-identification only in (Dutra et al., 2013; Khedher and El Yacoubi, 2015). Only in (Dutra et al., 2013) the proposed solution is a multi-stage system: the first stage selects a subset of templates using a descriptor which is built upon a bag-of-words feature representation and an indexing scheme based on inverted lists, and requires a low processing time for computing matching scores; the second stage ranks only the selected templates using a different, more complex descriptor based on mean Riemann covariance. In (Dutra et al., 2013) only two stages are considered, and only a subset of templates is ranked by the whole system, possibly losing the correct identity. The method of (Khedher and El Yacoubi, 2015) reduces processing time in the specific multi-shot setting (when several images per individual are available), and for specific descriptors based on local feature matching, e.g., interest points. It first filters out irrelevant interest points, then it builds a sparse representation of the remaining ones.

3 PROBLEM DEFINITION

The existing person re-identification descriptors can be subdivided into two main categories: fixed-size feature vectors (e.g., (Ma et al., 2014; Liao et al., 2015)), and descriptor with variable size (e.g., (Farenzena et al., 2010)). Note that, at this paper, we study only the descriptors with fixed-size feature vectors. For fixed-size descriptors, an unsupervised feature reduction technique like PCA can be used. The suitability of PCA to person re-identification tasks is witnessed to its use in the pre-processing step of gBi-Cov (Ma et al., 2014). However, we aim to study also some other feature reduction methods regarding to our motivation of our work.

In standard person re-identification systems, \mathbf{X}_T and \mathbf{X}_P are the generated features by a descriptor for a template and probe image, respectively, $m(\cdot, \cdot)$ the similarity measure between two descriptors, and $G = \{\mathbf{X}_{T_1}, \dots, \mathbf{X}_{T_n}\}$ the template gallery. For a given \mathbf{X}_P , a standard re-identification system computes the matching scores $m(\mathbf{X}_P, \mathbf{X}_{T_i})$, $i = 1, \dots, n$, and returns the list of template images ranked for decreasing values of the score. Ranking accuracy is typically evaluated using the cumulative matching characteristic (CMC) curve, i.e., the probability (recognition rate) that the correct identity is within the first ranks. Hereinafter we consider only the generated *fixed-size* feature vector (e.g. \mathbf{X}) by a specific descriptor. Fig. 1 presents the whole scheme of our strategy; aiming to employing a feature reduction method on person re-identification task

Apparently, some redundancies of patterns can be occurred within a feature vector, which are intuitively effected on processing time on real-time applications. It is worth to remind the readers that the issue of processing time in person re-identification can be categorized from two point of views: the processing time of constructing descriptor (*aka descriptor generation*); which can be done off-line for the gallery set, and the processing time of computing matching score between pair of descriptors (*aka descriptors matching*); which has to be done on-line for investigating an individual of interest (i.e. probe image) as the procedure of the real-time application. At this thesis work, whereas a feature reduction method needs a training phase to project the proper patterns into the low-dimensional space as same as a step of the re-identification system which needs to construct the descriptor for each individual. We consider, instead, the issue of the matching processing time of a single probe image and a template image. We therefore study on the feature vectors generated by the descriptors, and investigate an empirical procedure to attain a significant trade-off between processing time and ranking quality in person re-identification task.

Moreover, having redundant and irrelevant patterns from the feature vectors might be caused in overfitting problem. Removing these irrelevant pattern from the feature space before tacking them in real-world application scenarios is know as prepare data step in machine learning processes. To sum up, there might be three key advantages of feature reduction methods:

1. decrease the risk of overfitting; which allows the algorithm to make a decision in less redundant data.
2. improve the recognition accuracy; which avoid the algorithm by occurrence of misleading those irrelevant data.
3. decrease the processing time; which leads the method to be faster.

We therefore study on the feature vector generated by the descriptors, and investigate an empirical procedure to attain a significant trade-off between processing time and ranking quality in person re-identification task.

4 FEATURE REDUCTION METHODS

At this paper, we discuss on three feature reduction methods and their behaviours in person re-identification task.

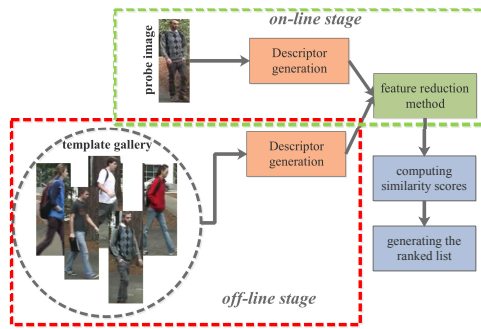


Figure 1: Application of a feature reduction method in person re-identification.

4.1 Principal Component Analysis (PCA)

PCA (Diamantaras and Kung, 1996) is pretty a well-known method for linear dimensional reduction. This method leads to identify important patterns within a data, and express the data in low-dimensional patterns by keeping the nature of the data at the same time (*aka* compressing data). PCA, typically, employs Singular Value Decomposition (SVD) of the features to project a feature vector into a lower dimensional space. SVD can be more fundamental in the concept of feature reduction method, since not only it provides direct approach to compute the principle components (PCs), but also simultaneously helpful to obtain row and column spaces (Zhang et al., 2007). At the following, we go through some brief introductions of the mathematical point of views of SVD.

Let $X = \{x_1, x_2, \dots, x_n\}$ be the given feature vector of size n to be compressed. While X is denoted as $X \in R^N$, where R represents the real numbers within the corresponding feature vector. The SVD of X is defined as

$$X = USV^H \quad (1)$$

where $U \in F^{N \times N}$ and $V \in F^{N \times N}$ are unitary matrices, and $S \in F^{N \times N}$ as a diagonal matrix, $S = \text{diag}(\alpha_1, \dots, \alpha_r, 0, \dots, 0)$. The singular values are in the decreasing order, $\alpha_1 \geq \dots \geq \alpha_r \geq 0$. Accordingly, in many applications, it can be useful to approximate X with low-rank matrix

$$X = \begin{bmatrix} U_r & U_{n-r} \end{bmatrix} \begin{bmatrix} S_r & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_r^H \\ V_{n-r}^H \end{bmatrix} \quad (2)$$

Then, we have

$$\tilde{X} = U_r S_r V_r^H \quad (3)$$

4.2 Kernel Principal Component Analysis (KPCA)

Kernel-PCA (KPCA) is an improved theory of traditional linear PCA in a high-dimensional space which is constructed by employing a kernel function. In other words, KPCA is a non-linear dimensional reduction for the features to project a lower dimensional space using the kernel method, and then compute PCA on the high-dimensional feature space. However, KPCA can be applied based on a specific kernel of k , this leads the application to choose a proper kernel function (Schölkopf et al., 1999). Given a data set X of input samples $\{x_1, x_2, \dots, x_N\}$, a kernel is defined as follows

$$k : X \times X \rightarrow \mathbb{R} \quad (4)$$

$$(x_i, x_j) \mapsto k(x_i, x_j),$$

where the kernel $k(\cdot, \cdot)$ gives a scalar that describes the similarity of the samples x_i and x_j . In this work, the Gaussian kernel (RBF) is employed as follow

$$k(x_i, x_j) = \exp\left(-\frac{\|x_i - x_j\|^2}{2\sigma^2}\right) \quad (5)$$

Using the obtained kernel, the originally linear operations of PCA are performed (as explained in 4.1) to reduce the dimensionality of the kernel feature space. The Gaussian kernel (RBF) is chosen in this studying because the linear kernel function gives the same performance as applying the normal PCA in the original feature space.

4.3 Isomap

Isomap is a non-linear dimensional reduction through isometric feature mapping which consists of calculating quasi-isometric to obtain low-dimensional embedding of a set of high-dimensional data points. The algorithm is based on estimating of geometry features of data distribution, and then mapping data to a new space. Isomap is quite straightforward technique that resolves feature reduction problem by computing geodesic distances between each data point. Geodesic distance basically computes between two points over the manifold. In order to compute the distances between data points $x_i, i = 1, 2, \dots, N$, a neighborhood graph of G is constructed in which each data point x_i is connected with its k nearest neighbors $x_{ij}, j = 1, 2, \dots, k$, within a data set, and forming an estimation of the geodesic distance between two points by taking into account the shortest path between them. One can compute the shortest-path between all the data points by employing Dijkstra's (Dijkstra, 1959) or

Floyds algorithms (Fisher, 1936), and form into a pairwise geodesic distance matrix. The computational complexity of this method, only when the algorithm measures the shortest path on the neighborhood graph and creates a graph distances matrix, is the most time consuming and it is performed in $O(N^3)$ operations.

4.4 Error Estimation

The performance of a re-identification system is typically measured using the CMC curve, defined as the probability that the correct identity is within the first rnk ranks, for $\text{rnk} = 1, \dots, n$. By definition, the CMC curve increases with rnk , and equals 1 for $\text{rnk} = n$. Whereas, this work aims to reduce the feature space of the original descriptor, we employ reconstruction error to estimate the variances the projected feature space. However, this could be simply identified from the behaviour of corresponding CMC curve, but for sake of comparison of different reduction methods, the reconstruction error is computed between the original feature space and the projected feature space. In order to estimate the reconstruction error of the projected feature vector, We employ Frobenius norm. To this aim, by recalling the projected feature space (\bar{X}) and the original feature vector (X), the reconstruction error is estimated as

$$E = \frac{\|X - \bar{X}\|_F^2}{\|X\|} \quad (6)$$

5 EXPERIMENTAL EVALUATION

We carried out our comparison with two well-known descriptors in person re-identification problem: gBiCov and LOMO, on two benchmark data sets: VIPeR and i-LIDS. At the following, we give some explanation of the used descriptors as well as the data sets.

5.1 Descriptors and Data Sets

We used the gBiCov, and LOMO descriptors, and obtained reduced versions of each of them by using PCA, KPCA, and Isomap feature reduction methods. **gBiCov** (Ma et al., 2014) is based on biologically-inspired features (BIF) obtained by Gabor filters with different scales over the HSV color channels. The resulting images are subdivided into overlapping regions of 16×16 pixels; each region is represented by a covariance descriptor that encodes shape, location and color information. BIF and covariance descriptors are concatenated, and PCA is used to reduce its dimension. The final generate image signature con-

tains ≈ 6000 elements. **LOMO** (Liao et al., 2015) extracts an $8 \times 8 \times 8$ -bins HSV histogram and two scales of the Scale Invariant Local Ternary Pattern histogram (characterizing texture) from overlapping windows of 10×10 pixels; it then retains one only histogram from all windows at the same horizontal location, obtained as the maximum value among all the corresponding bins. These histograms are concatenated with the ones computed on a down-sampled image. The final generated image signature contains ≈ 27000 elements.

We evaluated our experimental results on two benchmark data sets: VIPeR, and i-LIDS. **VIPeR** (Gray and Tao, 2008) is a challenging data set for person re-identification; it is made up of two images of 632 individuals from two camera views, with pose and illumination changes, cropped and scaled to 128×48 pixels. **i-LIDS** contains 476 images of 119 pedestrians taken at an airport hall from non-overlapping cameras, with pose and lightning variations and strong occlusions.

5.2 Experimental Setup

One image for each person was randomly selected to build the template gallery; the other images formed the probe gallery. As in (Farenzena et al., 2010), for each data set we repeated our experiments on ten different subsets of individuals, using one image of each individual as template and one as probe, and reported the average CMC curve over the ten runs. We used an Intel Core i5 2.6 GHz CPU. We applied the above-mentioned feature reduction methods on two well-known descriptors on VIPeR data set in person re-identification task. The original feature generated by a *given* descriptor (X), are reduced for different sizes in $r = \{2, 5, 20, 50, 80, 100, 130, 150, 200, 300, 500, 800, 1000, 1200, 2000, 2500, 3500\}$. For KPCA, we have chosen the Gaussian kernel because of its good performance. For Isomap, we set the number of neighborhoods to $k = 5$.

5.3 Experimental Results

Figures 2 and 3 present the corresponding CMC curves obtained by using different descriptors as well as different feature reduction methods on VIPeR, and i-LIDS data sets, respectively. PCA as a standard technique, which the new techniques are still unable to outperform it. KPCA also has a very similar behaviour in terms of the recognition accuracy in person re-identification. Also, in both techniques, the recognition accuracy outperformed the original CMC curve

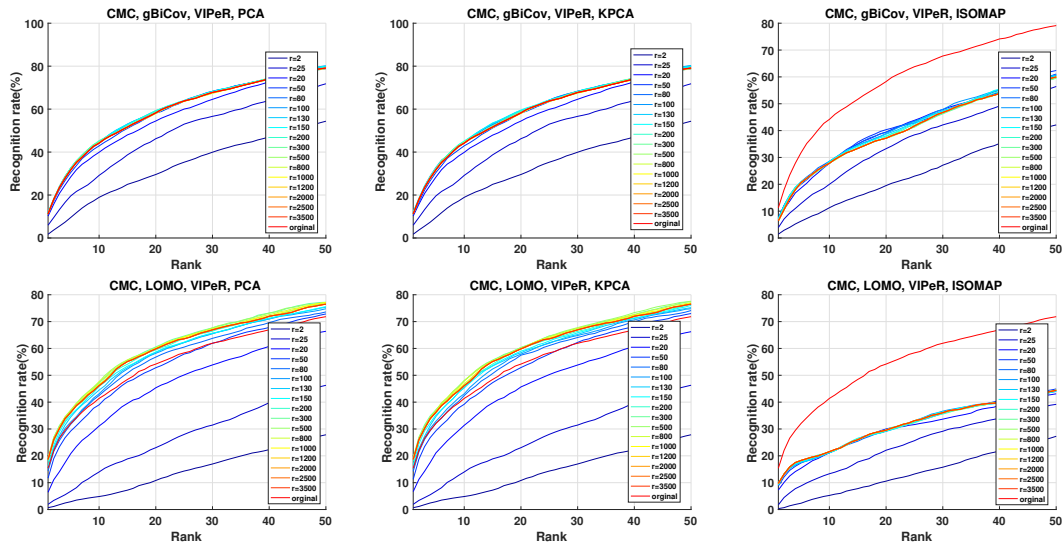


Figure 2: CMC curves obtained by gBiCov and LOMO descriptors on VIPeR data set in which the feature reduction methods have been employed. Figure is best viewed in color and under zoom.

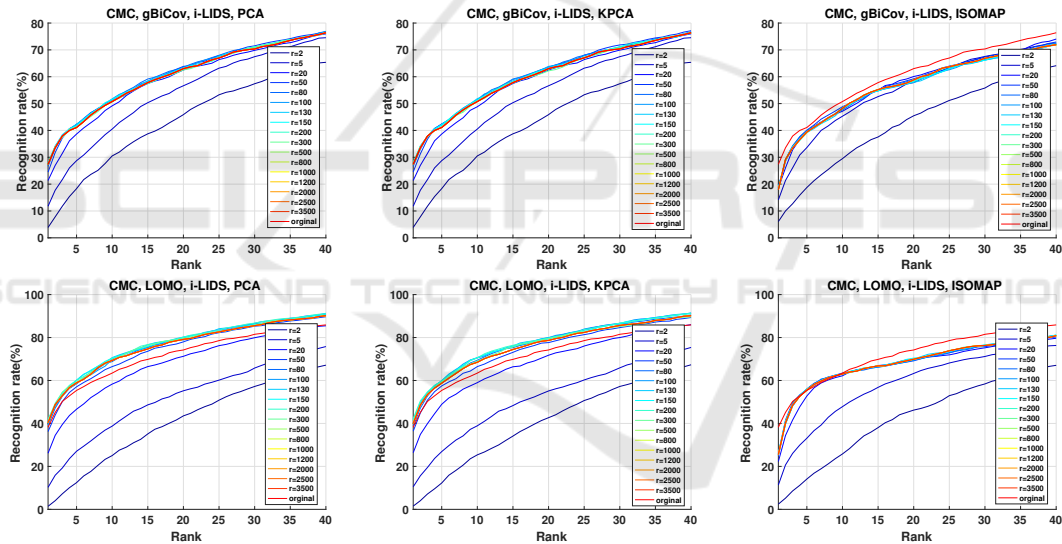


Figure 3: CMC curves obtained by gBiCov and LOMO descriptors on i-LIDS data set in which the feature reduction methods have been employed. Figure is best viewed in color and under zoom.

on LOMO descriptor when PCA and KPCA are applied for the original feature vector, and in the same way, with slightly better performance by using gBiCov on VIPeR and i-LIDS data sets. In contrast, KPCA is very time consuming because of the computational complexity when the feature vector is relatively larger in comparison to the other methods. The average processing time for computing one matching score, evaluated on VIPeR and i-LIDS, are reported in Fig. 4. Similar processing times were observed in all data sets, due to the use of the same image size. Moreover, figures 5 and 6 demonstrate the estimated errors among different feature reduction methods for

different values of r . Apparently from the presented figures, the error estimation using LOMO and gBicov descriptors with PCA and KPCA lead to be zero at the certain value of $r > 1000$ on VIPeR, and $r > 500$ on i-LIDS data set with different descriptors. However, Isomap has better behaviour only on i-LIDS data set with respect to its performance on VIPeR data set, and this is also obvious from the estimated error in which the error leads to be zero for $r > 500$.

To sum up, if one needs to achieve a significant trade-off between processing time and recognition rate, PCA could be a good solution in this regard. Despite from the processing time, one could also em-

ploy KPCA to gain higher recognition rate with expense of the higher processing time. This is due to KPCA is more time consuming than traditional PCA because of its mapping kernel function, as well as optimization of the kernel parameters.

6 CONCLUSIONS AND FUTURE WORK

At this paper we compared the performances of the most popular fundamental dimensional reduction methods, PCA, KPCA, and Isomap, on person re-identification data sets. The comparison is done through the experiments conducted by using two descriptors on two benchmark data sets. The experimental results evidenced that generated features by these descriptors might be not well-optimum. PCA and KPCA outperformed the original CMC curve on LOMO and gBiCov descriptors on VIPeR and i-LIDS data sets. This was apparent also from their error estimation of projected feature space using two descriptors on two data sets. Both these reduction methods achieve better performances rather than Isomap method in person re-identification task. The reason relies in the fact that PCA and KPCA can explore higher order information of the original inputs than Isomap. It is worth to point out that, at this work, PCA was better than others in terms of the computational cost, while KPCA was more time consuming with respect to the other two reduction methods. It therefore can be stated that PCA achieved promising performance for handling of optimization of raw data and projection of it to low-dimensional feature space. We studied this only for the descriptors with fixed-size feature vector. Finally, we point out that the optimization of the dimensional reduction methods analyzed in this paper is computationally and numerically practical in real-time applications. As the future work, we aim at carefully study the behavior of these feature reduction methods by concerning on some analytical terms, and visualize the projected data on the actual feature space to get better prospective on those behaviours.




Figure 4: Average processing time t_M (in sec.) for computing a matching score for a single probe image and one template, for each of the two descriptors with different feature sizes r .

ACKNOWLEDGMENT

We gratefully acknowledge Prof. Giuseppe Rodriguez of Dipartimento di Matematica e Informatica at University of Cagliari, for his initial analysis on the data and experimental evaluation of this paper.

REFERENCES

- Balasubramanian, M. and Schwartz, E. L. (2002). The isomap algorithm and topological stability. *Science*, 295(5552):7–7.
- Bedagkar-Gala, A. and Shah, S. K. (2014). A survey of approaches and trends in person re-identification. *Image and Vision Computing*, 32(4):270–286.
- Cheng, D. S., Cristani, M., Stoppa, M., Bazzani, L., and Murino, V. (2011). Custom pictorial structures for re-identification. In *BMVC*, page 6.
- Diamantaras, K. I. and Kung, S. Y. (1996). *Principal component neural networks: theory and applications*. John Wiley & Sons, Inc.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271.
- Dikmen, M., Akbas, E., Huang, T. S., and Ahuja, N. (2010). Pedestrian recognition with a learned metric. In *Computer Vision-ACCV 2010*, pages 501–512. Springer.
- Dutra, C. R., Schwartz, W. R., Souza, T., Alves, R., and Oliveira, L. (2013). Re-identifying people based on indexing structure and manifold appearance modeling. In *Graphics, Patterns and Images (SIBGRAPI), 2013 26th SIBGRAPI-Conference on*, pages 218–225. IEEE.
- Farenzena, M., Bazzani, L., Perina, A., Murino, V., and Cristani, M. (2010). Person re-identification by symmetry-driven accumulation of local features. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2360–2367. IEEE.
- Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of human genetics*, 7(2):179–188.
- Gray, D. and Tao, H. (2008). Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *Computer Vision-ECCV 2008*, pages 262–275. Springer.
- Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417.
- Khedher, M. I. and El Yacoubi, M. A. (2015). Two-stage filtering scheme for sparse representation based interest point matching for person re-identification. In *Advanced Concepts for Intelligent Vision Systems*, pages 345–356. Springer.
- Lavi, B., Fumera, G., and Roli, F. (2016). A multi-stage approach for fast person re-identification. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and*

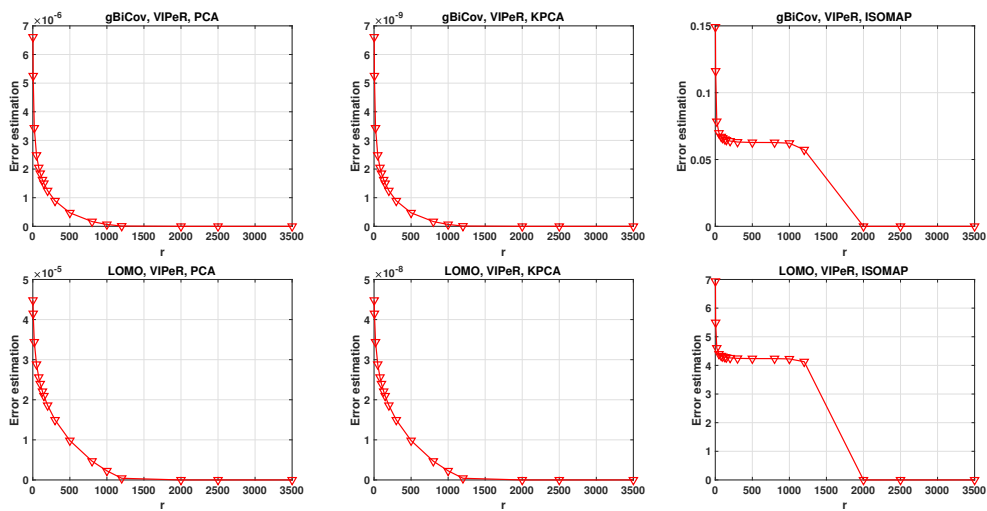


Figure 5: Error estimation by using gBiCov and LOMO descriptors on VIPeR data set for different feature reduction methods.

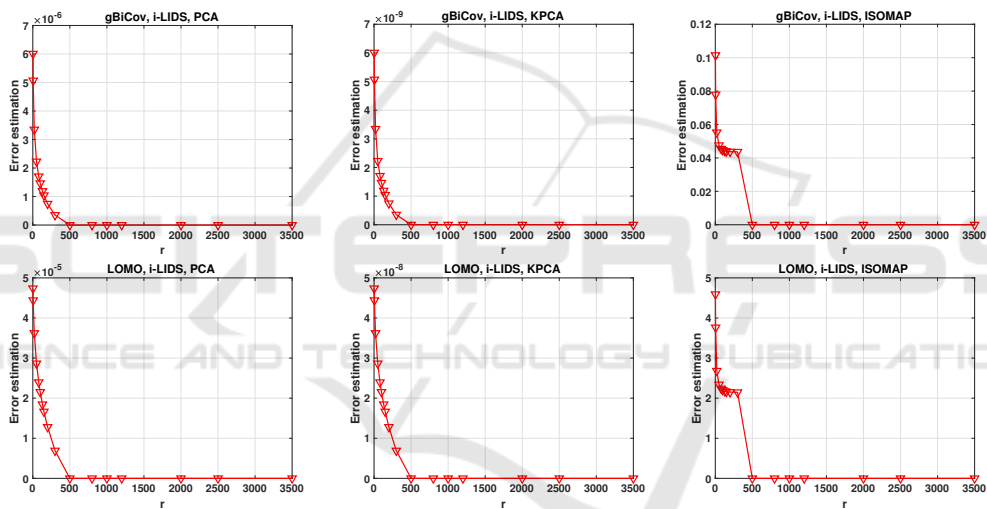


Figure 6: Error estimation by using gBiCov and LOMO descriptors on i-LIDS data set for different feature reduction methods.

Syntactic Pattern Recognition (SSPR), pages 63–73. Springer.

Liao, S., Hu, Y., Zhu, X., and Li, S. Z. (2015). Person re-identification by local maximal occurrence representation and metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2197–2206.

Ma, B., Su, Y., and Jurie, F. (2014). Covariance descriptor based on bio-inspired features for person re-identification and face verification. *Image and Vision Computing*, 32(6):379–390.

Prosser, B., Zheng, W.-S., Gong, S., Xiang, T., and Mary, Q. (2010). Person re-identification by support vector ranking. In *BMVC*, page 6.

Satta, R., Fumera, G., and Roli, F. (2012). Fast person re-identification based on dissimilarity representations. *Pattern Recognition Letters*, 33(14):1838–1848.

Schölkopf, B., Burges, C. J., and Smola, A. J. (1999). *Advances in kernel methods: support vector learning*. MIT press.

Advances in kernel methods: support vector learning. MIT press.

Zhang, L., Marron, J., Shen, H., and Zhu, Z. (2007). Singular value decomposition and its visualization. *Journal of Computational and Graphical Statistics*, 16(4):833–854.

Zheng, W.-S., Gong, S., and Xiang, T. (2013). Reidentification by relative distance comparison. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, pages 653–668.