

Dissimilarity Representations for Low-Resolution Face Recognition

Mairelys Hernández-Durán¹(✉), Veronika Cheplygina^{2,3},
and Yenisel Plasencia-Calaña^{1,3}

¹ Advanced Technologies Application Center, La Habana, Cuba
{mhduran,yplasencia}@cenatav.co.cu

² Biomedical Imaging Group Rotterdam, Erasmus Medical Center,
Rotterdam, The Netherlands

³ Pattern Recognition Laboratory, Delft University of Technology,
Delft, The Netherlands
v.cheplygina@tudelft.nl

Abstract. Low-resolution face recognition is a very difficult problem. In this setup, the training database or gallery contains high-resolution images, but the image to be recognized is of low resolution. Thus we are dealing with a resolution mismatch problem for training and test images. Standard face recognition methods fail in this setting, which suggests that current feature representation approaches are not adequate to cope with this problem. Therefore, we propose the use of dissimilarity representations based on different strategies, which differ in how images with different resolutions are compared, to solve the resolution mismatch problem. Experiments on four standard face datasets demonstrate that a strategy based on first down-scaling and afterwards up-scaling training images while up-scaling test images outperforms all the other approaches.

Keywords: Dissimilarity space · Low-resolution face recognition · Super-resolution · Prototype selection

1 Introduction

Face recognition has been studied for decades due to its wide range of applications. Although face recognition has achieved high recognition accuracy under controlled environments, in low-resolution face recognition (LR FR) systems the results are still unsatisfactory. Nowadays, there is a growing interest in real applications such as video protection and surveillance in which subjects are far away from the camera. In such scenarios, the face image sizes tend to be small and the images do not have a good definition of facial features. Moreover, discriminatory features present in the facial images used for distinguishing one person from another are lost due to the decrease in resolution, resulting in unsatisfactory performance. As a result, low-resolution (LR) images affect the performance of traditional face recognition systems. LR FR aims at recognizing face images with LR and variations such as pose and illumination. In LR FR the gallery contains

high resolution images while the test images are of low resolution, causing the so-called dimensional mismatch [1, 2].

Current approaches mainly include feature vector representations to allow a good discrimination between different faces for addressing LR FR. Methods such as the nearest neighbour (1-NN) and the bicubic interpolation are the simplest ways to increase resolution for an input LR image [3].

In [4] the authors propose a 1-NN approach for producing super-resolution images from ordinary images and videos. Sparse representation [5] and metric learning [6], are some of the feature methods for LR FR with the advantages of low computational complexity and lower requirement of training samples, making them more suitable for real applications. However, it is difficult to find a good feature representation in LR FR because most of the effective features used in high-resolution face recognition such as texture and color may fail in LR case. As a consequence, most of the successful approaches cannot be efficiently applied to LR case [3].

A representation based on dissimilarities between objects [7, 8] is an alternative to the feature-based representation. A dissimilarity-based representation is advantageous in situations where it is difficult to define sufficiently discriminative features, but it is easier to define dissimilarities. More specifically, the dissimilarity space (DS) approach is very attractive due to its efficiency and easy possibility to map new test objects compared to the Pseudo-Euclidean space representation [7].

Based on the success of previous works [7], we used the dissimilarity representation approach to tackle our problem. Intuitively, the proximity information is more important for discriminating between the classes than the composition and features of each object independently [9]. Particularly, we believe that a dissimilarity space representation can be suitable for LR FR because in the context of comparisons with the prototype objects we can compensate the noise introduced by the low resolution as well as the lack of information in such low resolution images. By using the differences with the prototype images for creating the representations we may be able to emphasize relevant information for discrimination among the classes, which, otherwise, by only analyzing the image, may be difficult to express in a feature representation. Furthermore, a dissimilarity representation has been used for other difficult problems as well such as: small sample size situations [10] or problems where the results of the 1-NN on features are still unsatisfactory [8, 11].

In this work, we present an alternative to feature-based representations for LR FR based on the DS representation. We compare the proposed dissimilarity representation with feature representations for LR FR and also for very low-resolution face recognition. Three different strategies are tested based on original or up-scaled test images, and original or down-scaled training images to address the mismatch problem between training and test images. The comparisons show that the dissimilarity space representation outperforms the feature representation and that the low-high strategy, where the training images are down-scaled and then up-scaled while the test images are up-scaled, is the best

way to cope with the mismatch problem. In particular, the linear discriminant classifier (LDC) in the dissimilarity space is very promising.

The paper is organized as follows. Section 2 presents the related work on LR FR and the dissimilarity representation. Section 3 presents our proposed reduced dissimilarity space to cope with classification of LR and very low-resolution images. Experiments and discussion are presented in Sect. 4, and concluding remarks are provided in Sect. 5.

2 Related Work

The purpose of LR FR is to recognize faces from small size or poor quality images (e.g. face inside a 32×20 pixels image) which can also present challenging facial variations such as pose, illumination, and expression. The LR of the test images causes a dimensional mismatch when having to deal with high resolution training images. Three main research lines have been considered to cope with the problem: interpolation [12, 13], down-scaling [14] and unified feature space [15]. The first approach has limitations associated to the scale factor and it is more suitable for synthesizing generic objects or scenes instead of faces. The second approach allows to match in the LR domain by down-sampling the training set, but it represents a reduction of the information useful for the recognition process. In the third approach, although it seems feasible to cope with the mismatch problem, it is not easy to find an optimal inter-resolution space.

Several methods have been used for recognizing faces from LR images. Super resolution (SR) is one of the most frequently employed techniques for dealing with this problem. SR methods recover the lost information during the image formation process by including a-priori information about the image. SR methods produce a reconstructed high-resolution image from a low-resolution one by making assumptions about the image structure or content. The first SR techniques based on reconstruction represent an intuitive approach to improve a face image, but are aimed mostly at a visual improvement, and are not designed from a pattern recognition point of view.

Recently, Zou and Yuen [14] proposed the very low recognition problem, where the resolution of the face images to be recognized is lower than 16×12 pixels. Hennings et al. included facial features as prior information into an SR method named Simultaneous Super-Resolution and Recognition (S2R2) [2] to improve the results. They showed that when faces are of very low-resolution, the approach of matching in the low-resolution domain is better than applying SR. Li et al. [15] proposed the coupled locality preserving mappings method to include robust features in a unified feature space for increasing the discriminability in the recognition process. Nevertheless, finding a resolution-robust feature representation is still far from being a solved problem.

An alternative solution is a dissimilarity representation between objects based on the general idea proposed in [7], in which dissimilarities are considered as the connection between perception and higher-level knowledge, thus being an important factor in the process of human recognition and categorization. The

dissimilarity representation is also able to deal with several problems related to the feature vector representation. A feature-based description may be difficult to find or can be inefficient for the learning task. Furthermore, the dimensionality of the feature vector is usually larger than the number of images, commonly known as the curse of dimensionality. Another advantageous property of this representation is the possibility to learn from small sample sizes [10].

The dissimilarity-based approach has successfully been used for multiple tasks such as person re-identification [16] and object classification [17]. In [16], Satta et al. convert a given appearance-based re-identification method into a dissimilarity-based one and show a reduction in both the processing time and the memory requirements. In [18], Orozco et al. use a dissimilarity-based method for face recognition which was derived by applying the eigenface transformation and, afterwards, the Euclidean distance between the eigenface representations.

Our present work differs from these works in several aspects. The application considered in this paper is very different from previous applications as we have to transform the images first to cope with the resolution mismatch problem, i.e., we propose different strategies to be able to compare test images with training images. We also propose the use of a reduced dissimilarity space by using prototype selection, including an analysis of its benefits at test time. We show experimentally that one of our proposals is very promising, and that a small dimensionality of the DS is sufficient to achieve a good discrimination among the classes.

3 Proposed Approach: Reduced Dissimilarity Space

3.1 Dissimilarity Space and Prototype Selection

Dissimilarity representations have been studied in a number of problems [18–20], however their application for LR FR has not been studied so far. We believe that this type of relational representation can cope with the poor discriminability of standard feature representations when using LR images. Let X be the space of objects, let $R = \{r_1, r_2, \dots, r_k\}$ be the set of prototypes such that $R \in X$, and let $d : X \times X \rightarrow \mathbb{R}^+$ be a suitable dissimilarity measure for the problem. For a training set $T = \{x_1, x_2, \dots, x_l\}$ such that $T \in X$, a mapping $\phi_R^d : X \rightarrow \mathbb{R}^k$ defines the embedding of training and test objects in the DS by the dissimilarities with the prototypes:

$$\phi_R^d(x_i) = [d(x_i, r_1) \ d(x_i, r_2) \ \dots \ d(x_i, r_k)]. \quad (1)$$

In a problem where training, prototype, and test images have the same resolution it is straightforward to apply the approach. However, in our setup, test images are of LR, so we need to decide how to deal with the resolution mismatch problem. We compare three different strategies to cope with the resolution mismatch between training, prototype, and test images:

- Low-resolution test images, down-scaled training images (low) and down-scaled prototypes

- up-scaled low-resolution test images, down-scaled and then up-scaled training images (low-high), and high-resolution prototypes
- up-scaled low-resolution test images, high-resolution training images (high), and high-resolution prototypes

The same training set can be used as the set of prototypes. However, for training sets of moderate to large size, a selection of the best set of prototypes is needed to find a trade-off between classification accuracy and computational efficiency. This can be achieved by selecting a reduced set of prototypes which has similar performance to using the whole set.

To select the reduced set of prototypes we need a search strategy with a suitable criterion. Different approaches have been previously studied for this purpose (see [8, 19]). Recently, a genetic algorithm (GA) was proposed in [21], which showed to be very fast and accurate in selecting a good set of prototypes. It proposes a number of improvements to the simple GA such as the use of indexes for codifying the prototypes instead of binary chromosomes, and an early stopping criterion which was shown to be adequate for this type of problem. In addition, only scalable criteria are considered for the fitness function to evaluate each solution (set of prototypes), therefore the method is fast and scalable. We will use the supervised prototype selection strategy from [21] to find an adequate set of prototypes for a given or desired cardinality of the DS.

The GA can also be used for feature selection by using a slightly different selection criterion. The criterion for selecting prototypes is based on maximizing matching labels between the prototypes and their nearest neighbours. Therefore, for selecting features, it is replaced by a criterion minimizing the nearest neighbour error in the training set for a feature set of a given cardinality.

3.2 Considerations at Test Time

We want to remark the advantages of a reduced dissimilarity space (RDS) by prototype selection in comparison with a RDS by feature extraction as well as the advantages over a reduced feature space (RFS) by feature selection or by feature extraction.

Suppose we have these spaces with the same dimensionality. The problem of a feature space with selected features is that we lose the information contained in the discarded features, especially in problems where the majority of the features are informative. Even if only the selected features are informative, due to the nature of the representation (such as a histogram), all features might need to be extracted before discarding the non-informative ones. In contrast, once the prototypes were selected to create a RDS, for a new test object we only need to measure the dissimilarities with the selected prototypes. Besides, a small set of prototypes is often enough to represent the data properly which is not the case for handcrafted feature representations [22].

Feature extraction methods, both in a feature space or a DS, present even stronger disadvantages in terms of computing time at test time. These methods

always require the computation of the full set of features (or alternatively dissimilarities with the large set of prototypes) before applying the transformation to a reduced space, which is performed by expensive floating-point multiplications of the test object representation with elements from a mapping or projection matrix. These costs are not adequate for deployment in real-world scenarios [22].

4 Experiments and Discussion

This section presents the experimental comparison, results and discussion of different feature-based and dissimilarity based strategies for the classification of LR images where the gallery is composed by high resolution images.

4.1 Databases Description

Four different standard face datasets were used for the experiments. In each case, the test images were obtained by down-scaling the original images using a bicubic interpolation. All images were geometrically normalized by the center of the eyes to a LR size of 10×12 pixels or 24×30 pixels during experiments. A bicubic interpolation was also applied in the up-scaling process to obtain high resolution images of 64×80 pixels.

Olivetti Research Database (ORL) [23]. The ORL database contains 400 grayscale images of 40 individuals, 10 images per person. Some images are taken with a certain time difference. They present variations in facial expression (including opening and closing the eyes), illumination changes, different details on the face (with and without glasses) and a slight difference in scale. Figure 1 shows examples of variations on this database.



Fig. 1. Some examples of ORL database

Yale Database [24]. The Yale database contains images with variations in lighting condition (left-light, center-light, right-light), facial expression (normal, happy, sad, sleepy, surprised, and wink), and with/without glasses. Figure 2 shows example images with some variations for the individuals. During the experiments we used a subset of the database, which consists of 200 images belonging to 10 subjects with different variations. Some subsets were removed because they have strong differences in lighting conditions and addressing this problem is not the purpose in this work.




Fig. 2. Some examples of Yale database

Essex Database [25]. The database contains single light source images with racial diversity, and variations with glasses, beards, and so forth. The images are captured from a fixed distance with different orientation and different facial expression. The database consists of images of 153 individuals (20 images each). Each image has plain green background with no head scale but with very minor variation in head turn, tilt and slant. Some example images are shown in Fig. 3.

During the experiments we used a subset of the database which consists of 720 images in total belonging to 20 different subjects having 36 images per person with different variations. Some subsets were removed to focus on the low-resolution problem.



Fig. 3. Some examples of Essex database

Labeled Faces in the Wild (LFW) [26]. It contains 13233 labelled faces of 5749 people. For 1680 people two or more faces are available. The data is challenging, as the faces are detected in images “in the wild”, taken from Yahoo! News. The faces present some variations including changes in scale, pose, background, hairstyle, clothing, expression, image resolution, focus, and others. During the experiments we used a subset of the database consisting of 3 832 images belonging to 178 classes, by selecting the classes with 8 or more images. Some example images are shown in Fig. 4.

The characteristics of the datasets are summarized in Table 1.

4.2 Experimental Setup

We randomly divided the datasets into two sets for training and testing of equal size five times, ensuring that each class is equally represented in each set. The classifiers as well as the prototype selectors are trained using the training set and classification errors are computed for the test set. The average error values are reported.



Fig. 4. Some examples of LFW database

Table 1. Characteristics of the datasets used for the experiments

Datasets	# Classes	# Obj in total
ORL	40	400
Yale	10	200
Essex	36	720
LFW	178	3832

We consider two different representation spaces: a feature space (feat) and a dissimilarity space (DS). Furthermore, we consider two different classifiers: the linear discriminant classifier (LDC), which assumes equal covariance matrices for the classes, and the 1-NN.

In order to obtain the feature representation, we compute local binary patterns on local blocks of the geometrically normalized images. Histograms were computed on each block and concatenated. Chi square distances are used for the 1-NN classifiers as well as for creating the DS. Note that, in our case, the dissimilarity measure was computed on top of a feature representation, therefore we suffer from the cost of first computing the feature representation. However, a dissimilarity representation can also be computed by directly matching the images if we have a good dissimilarity measure for this purpose.

As it would be convenient to compute the dissimilarity measure by matching the images directly, we reviewed the literature to find good (dis)similarity measures for this purpose. However, we found that such measures are not as heavily used for face recognition as feature-based measures. This happens because several conditions affect facial images such as differences in pose, illumination, expression, and other capturing conditions, which directly affect image matching measures such as correlation. Unfortunately, despite several attempts to create good illumination and pose normalization methods to improve the original images so they can be used for direct matching, it is easier to use features that intrinsically deal with these problems such as the local binary patterns histograms that we used as base for computing the dissimilarities. The definition of such a measure that is able to deal with the mentioned problems directly is still an open issue.

In general, our motivations behind the use of dissimilarities on top of features for the experiments are: first, we can perform a fair comparison between the feature representation and the dissimilarity representation since it was computed on top of the same feature representation, second, the Chi square distance

measure on top of the local binary patterns histograms have shown very good performances in previous works for face recognition [27]. Therefore, it is a good starting point for our research.

Different DSs are created for each of the strategies and classifiers. However, as a baseline, the results of the 1-NN and LDC in the feature space are shown only for the best performing resolution strategy in the DS.

As parameters for the GA for prototype selection we used very similar parameters to [21]:

- 40 chromosomes for the population
- 30 generations reached or 10 generations without change in the fitness value as stopping criteria
- Reproduction probability equal to 0.5
- Mutation probability equal to 0.02

For the feature representation the same GA was used for feature selection to compare the feature space and DS space with the same dimensionality. The criterion used for feature selection is an equivalent version to the one used for prototype selection, the minimization of the 1-NN error on the training set.

4.3 Results and Discussion

Figures 5, 6, 7 and 8 show error rates for different numbers of prototypes in the DS or features in the feature space. For the 1-NN all the features are used. For both baseline classifiers the training set used is consistent with the one used for the different DS. Note that the 1-NN with the up-scaled images (1-NN low-high feat.) correspond to a variant of baseline in LR FR, the so-called super resolution.

From the results we can see that the DS representation outperforms the feature representations. We think that LR images benefit from the relational

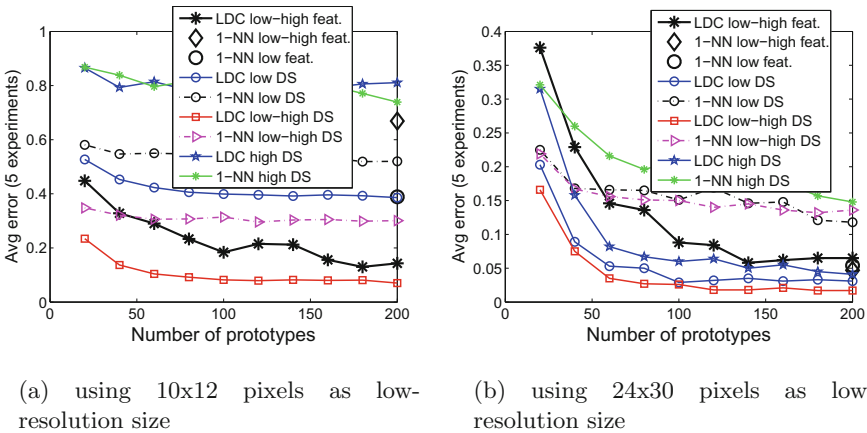




Fig. 5. Experimental results in ORL database




(a) using 10x12 pixels as low-resolution size

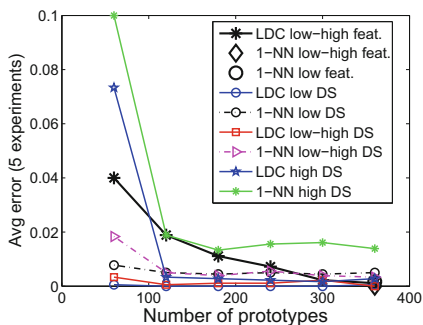


(b) using 24x30 pixels as low-resolution size

Fig. 6. Experimental results in Yale database



(a) using 10x12 pixels as low-resolution size



(b) using 24x30 pixels as low-resolution size

Fig. 7. Experimental results in Essex database

representation since features alone may not capture relevant information for discrimination. Comparisons with other objects can provide relevant information for discrimination since small details only present in high resolutions are not as influential as in a feature representation. The LR and high-resolution strategies perform poorly, while the best performing strategy is the low-high one. Especially the classification results with the LDC in the DS for this strategy are very promising.

The low-high strategy focuses on making the gallery images resemble the condition of the test images, since they are down-scaled and then up-scaled in the same way as the test images. In higher resolutions, the feature representation is able to capture the relevant information which is not possible for the LR case. Therefore, original high-resolution training images may be useful for comparing

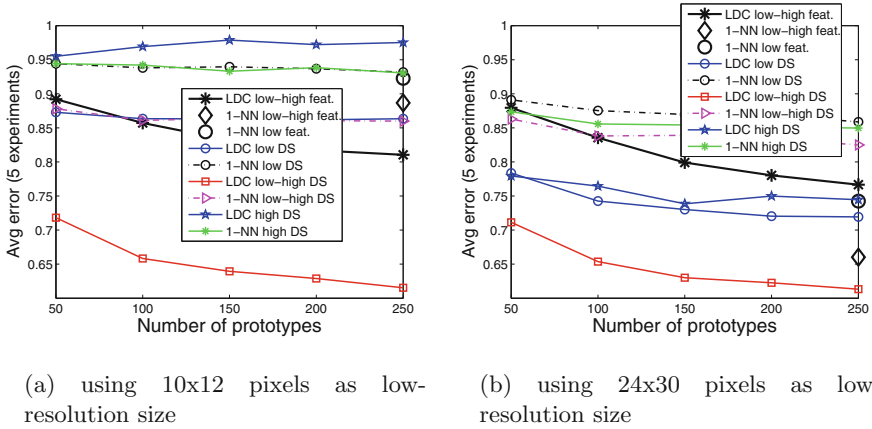


Fig. 8. Experimental results in LFW database

high to high resolution but they are definitely not good when the test images were originally of LR. We found that while the resolution of the test images increases, the classification results in the DS improve, especially when using high resolution training images.

Our results contradict those of Hennings et al. [2] where the authors found that the approach of matching in the low-resolution domain is better than applying SR when faces are of very low-resolution. What we found is that it is better to up-scale the test images and match them to the training images, instead of matching the original LR images. However, what is different in our approach is that we propose that the training images must also go through the same transformation process.

Note that the dissimilarity representations are very compact since the length of the final vectors is equal to the number of prototypes, and from the figures it can be seen that a small set of prototypes (e.g. cardinality equal to 200) is usually sufficient to obtain a good representation. This makes the approach suitable for large-scale and real-time recognition systems. This is also beneficial for representing a new test object since it implies that at test time only the dissimilarities with the small set of prototypes need to be measured. Note that we do not compare feature extraction methods because they would require the computation of dissimilarities with all the prototypes before performing the reduction for incoming test objects. This poses an extra computational cost that is avoided by our proposal.

5 Conclusions

In this paper we presented the reduced dissimilarity space (RDS) as an alternative representation for low-resolution face recognition. Different dissimilarity-based representations were compared with feature-based representations.

We found that using the down-scaled gallery and prototype images is counterproductive, while the strategies that up-scale the test images perform the best. However, there was a large difference between using the gallery or training images in their original high resolution and transforming them by first down-scaling and afterwards up-scaling them again. The proposed transformation outperformed using the gallery images in their original resolution. This is interesting since previous approaches focused on finding the best transformation for the low-resolution test images to resemble the high resolution images from the gallery, while we propose to also transform the gallery images to resemble the low-resolution test images.

The experiments showed that more discriminative information for classification can be obtained if the LR images are analyzed in the context of dissimilarities with other images. Note that, as our approach only assumes general dissimilarity measures, it can be used with any user-defined or learned metric. Dissimilarity measures computed directly on the images are desirable, however we did not find such a measure in the literature with good results and adopted an established dissimilarity for face recognition. In addition, our approach produces very compact representations which are suitable for large-scale and real-time recognition systems.

Future studies will be devoted to study metric learning approaches to create more discriminative dissimilarity measures or to improve the representation in the dissimilarity space. Furthermore, extending the dissimilarity space with additional dissimilarity measures [28] or prototypes from outside the training set [29] could be of interest. We believe that a learned representation using the dissimilarity representation as a starting point could improve the results even further. The design of robust measures for matching the images directly is also an interesting open issue.

References

1. Choi, J.Y., Ro, Y.M., Plataniotis, K.N.: Feature subspace determination in video-based mismatched face recognition. In: 8th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2008, pp. 1–6. IEEE (2008)
2. Hennings-Yeomans, P.H., Baker, S., Kumar, B.V.: Simultaneous super-resolution and feature extraction for recognition of low-resolution faces. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2008, pp. 1–8. IEEE (2008)
3. Wang, Z., Miao, Z., Wu, Q.J., Wan, Y., Tang, Z.: Low-resolution face recognition: a review. *Visual Comput.* **30**(4), 359–386 (2014)
4. Zhu, S., Han, M., Gong, Y.: Super resolution using gaussian regression, US Patent 7,941,004, 10 May 2011
5. Wright, J., Yang, A.Y., Ganesh, A., Sastry, S.S., Ma, Y.: Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(2), 210–227 (2009)
6. Li, B., Chang, H., Shan, S., Chen, X.: Coupled metric learning for face recognition with degraded images. In: Zhou, Z.-H., Washio, T. (eds.) *ACML 2009. LNCS*, vol. 5828, pp. 220–233. Springer, Heidelberg (2009)

7. Pekalska, E., Duin, R.P.W.: The Dissimilarity Representation for Pattern Recognition: Foundations and Applications (Machine Perception and Artificial Intelligence). World Scientific Publishing Co. Inc, River Edge (2005)
8. Pekalska, E., Duin, R.P.W., Paclík, P.: Prototype selection for dissimilarity-based classifiers. *Pattern Recogn.* **39**(2), 189–208 (2006)
9. Plasencia-Calaña, Y., García-Reyes, E., Orozco-Alzate, M., Duin, R.P.W.: Prototype selection for dissimilarity representation by a genetic algorithm. In: Proceedings of the 20th International Conference on Pattern Recognition, ICPR 2010, Washington, DC, USA, pp. 177–180. IEEE Computer Society (2010)
10. Orozco-Alzate, M., Duin, R.P.W., Castellanos-Domínguez, G.: A generalization of dissimilarity representations using feature lines and feature planes. *Pattern Recogn. Lett.* **30**(3), 242–254 (2009)
11. Pekalska, E., Duin, R.P.: Dissimilarity representations allow for building good classifiers. *Pattern Recogn. Lett.* **23**(8), 943–956 (2002)
12. Wolberg, G.: *Digital Image Warping*, vol. 10662. IEEE computer society press, Los Alamitos (1990)
13. Baker, S., Kanade, T.: Hallucinating faces. In: 2000 Proceedings of Fourth IEEE International Conference on Automatic Face and Gesture Recognition, pp. 83–88. IEEE (2000)
14. Zou, W.W., Yuen, P.C.: Very low resolution face recognition problem. *IEEE Trans. Image Process.* **21**(1), 327–340 (2012)
15. Li, B., Chang, H., Shan, S., Chen, X.: Low-resolution face recognition via coupled locality preserving mappings. *IEEE Signal Process. Lett.* **17**(1), 20–23 (2010)
16. Satta, R., Fumera, G., Roli, F.: Fast person re-identification based on dissimilarity representations. *Pattern Recogn. Lett.* **33**(14), 1838–1848 (2012)
17. Carli, A., Castellani, U., Bicego, M., Murino, V.: Dissimilarity-based representation for local parts. In: 2010 2nd International Workshop on Cognitive Information Processing (CIP), pp. 299–303. IEEE (2010)
18. Orozco-Alzate, M., Castellanos-Domínguez, C.: Nearest feature rules and dissimilarity representations for face recognition problems. *Face Recognition; International Journal of Advanced Robotic Systems*, Vienna, Austria, pp. 337–356 (2007)
19. Bunke, H., Riesen, K.: Graph classification based on dissimilarity space embedding. In: da Vitoria Lobo, N., Kasparis, T., Roli, F., Kwok, J.T., Georgiopoulos, M., Anagnostopoulos, G.C., Loog, M. (eds.) *SSPR & SPR 2008*. LNCS, vol. 5342, pp. 996–1007. Springer, Heidelberg (2008)
20. Li, Y., Duin, R.P., Loog, M.: Combining multi-scale dissimilarities for image classification. In: Proceedings of the 2012 21th International Conference on Pattern Recognition, ICPR 2012. IEEE Computer Society (2012)
21. Plasencia-Calaña, Y., Orozco-Alzate, M., Méndez-Vázquez, H., García-Reyes, E., Duin, R.P.W.: Towards scalable prototype selection by genetic algorithms with fast criteria. In: Fränti, P., Brown, G., Loog, M., Escolano, F., Pelillo, M. (eds.) *S+SSPR 2014*. LNCS, vol. 8621, pp. 343–352. Springer, Heidelberg (2014)
22. Chen, D., Cao, X., Wen, F., Sun, J.: Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification. In: Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2013, Washington, DC, USA, pp. 3025–3032. IEEE Computer Society (2013)
23. Samaria, F.S., *t, F.S.S., Harter, A., Site, O.A.: Parameterisation of a stochastic model for human face identification (1994)
24. Belhumeur, P.N., Hespanha, J.P., Kriegman, D.J.: Eigenfaces vs. fisherfaces: recognition using class specific linear projection. In: Buxton, B.F., Cipolla, R. (eds.) *ECCV 1996*. LNCS, vol. 1064, pp. 43–58. Springer, Heidelberg (1996)

25. Spacek., L.: Essex face recognition data. <http://cswww.essex.ac.uk/mv/allfaces/index.html>
26. Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst (2007)
27. Ahonen, T., Hadid, A., Pietikäinen, M.: Face recognition with local binary patterns. In: Pajdla, T., Matas, J.G. (eds.) ECCV 2004. LNCS, vol. 3021, pp. 469-481. Springer, Heidelberg (2004)
28. Plasencia-Calaña, Y., Cheplygina, V., Duin, R.P.W., García-Reyes, E.B., Orozco-Alzate, M., Tax, D.M.J., Loog, M.: On the informativeness of asymmetric dissimilarities. In: Hancock, E., Pelillo, M. (eds.) SIMBAD 2013. LNCS, vol. 7953, pp. 75-89. Springer, Heidelberg (2013)
29. Dinh, V.C., Duin, R.P.W., Loog, M.: A study on semi-supervised dissimilarity representation. In: International Conference on Pattern Recognition, pp. 2861-2864. IEEE (2012)