

Impact of Face Registration Errors on Recognition

E. Rentzeperis, A. Stergiou, A. Pnevmatikakis and L. Polymenakos
Athens Information Technology, Autonomic and Grid Computing,
Markopoulou Ave., 19002 Peania, Greece
{eren, aste, apne, lcp}@ait.edu.gr
<http://www.ait.edu.gr/research/RG1/overview.asp>

Abstract. Face recognition systems detect faces in moving or still images and then recognize them. However, face detection is not an error-free process, especially when designed for real-time systems. Thus the face recognition algorithms have to operate on faces that are not ideally framed. In this paper we analyze quantitatively the impact of face detection errors on six different face recognition algorithms. Hence, we propose a matching of face recognition algorithms with face detector performance, which can be used for a system based on the expected performance of the face detector.

1 Introduction

Face recognition has many applications in security and human-machine interfaces. It can be applied either like other biometric recognition approaches, by requiring the person to pose in a very controlled way, or unobtrusively, on still or moving images of the person. In the first approach, the face is manually extracted and, if necessary, normalized, prior to the application of any face recognition algorithm. The second approach is much more interesting, but poses a serious problem: the automatic detection and normalization of the face.

Many different face recognition algorithms have been reported in the literature [1-8] and have been tested on different face databases [8-11], adhering or not to a standard evaluation methodology [11-13]. In most of these reports, the performance of the different algorithms is assessed on manually extracted and normalized faces. Comparisons, when attempted with a common database and methodology, only assess which method performs best under a particular type of impairment (pose, illumination or expression variation) [13,14].

When all these algorithms are used in the context of a fully-automatic face recognition system, then the faces presented to the recognizer are not manually extracted. The automatic face detection process involves finding the face, locating in

Please use the following format when citing this chapter:

Rentzeperis, Elias, Stergiou, Andreas, Pnevmatikakis, Arstodemos, Polymenakos, Lazaros, 2006, in IFIP International Federation for Information Processing, Volume 204, Artificial Intelligence Applications and Innovations, eds. Maglogiannis, I., Karpouzis, K., Bramer, M., (Boston: Springer), pp. 187-194

it some features and based on these, geometrically normalizing the face to match some template [15-17]. This is the face registration process. Face detection on still images given ample processing time is a problem more or less solved [16], resulting to small registration errors. For real-time operation and under unconstrained lighting and pose conditions, face registration becomes very difficult [15], and the registration errors can grow significantly.

Hence there is a need to compare the different face recognition algorithms in terms of their robustness to face registration errors. Here we address this need by assessing the impact of face registration errors on six different face recognition algorithms. We use the evaluation methodology in [13] to obtain the performance of the algorithms, and then perform statistical analysis of the results to obtain the optimum method for different face detection accuracy ranges.

This paper is organized as follows: In Section 2, the face recognition algorithms to be compared are outlined. Then, in Section 3 the evaluation methodology and the experimental results are presented. Finally, in Section 4 the conclusions are drawn.

2 Face recognition methods

The impact of face registration errors on six face recognition algorithms is compared. The algorithms are the Eigenfaces [1], Fisherfaces [3], Elastic Bunch Graph Matching (EBGM) [4], pseudo Two-Dimensional Hidden Markov Models (2D-HMM) [2], correlation filters [5] and Laplacianfaces [6]. These algorithms have been implemented and extensively tested by the authors [18-20] on a variety of standard face databases and on a novel video database [13]. They have been incorporated into a real-time face recognition system [15] operating in the SMART lab [21], allowing us to determine the best algorithm for the system according to the performance of the detection system. In the rest of this section, these algorithms are briefly introduced.

2.1 Eigenfaces

Eigenfaces [1] is a linear subspace projection algorithm that uses Principal Component Analysis (PCA). As no class label information is used in PCA, the projection is estimated in an unsupervised manner. After linear projection, the resulting recognition space is of much lower dimension. The PCA feature vectors are robust to noise and minor head rotations, but not to illumination changes [3,18].

Since its introduction in 1991, the eigenface technique has seen many modifications [3]. In [3,18] the influence of distance metrics and eigenvector selection on PCA performance is analysed. Eigenvector selection consists of discarding a few eigenvectors with larger eigenvalues and/or some of those with the smallest. This is attributed to the empirical observation that the discarded eigenvectors with the larger eigenvalues encode direction-of-illumination changes. Discarding three such eigenvectors is shown in [3,18] to greatly enhance PCA performance.

2.2 Fisherfaces

PCA maximizes the total scatter of the training vectors while reducing their dimensions. It is optimum in the mean-squared error sense for representation in the resulting subspace, but offers no guarantee of optimality for classification. Linear Discriminant Analysis (LDA) on the other hand does take into account class labels and maximizes the between-class scatter under the constraint that the within-class scatter is minimized. This results to compact clusters for each class, as far as possible from each other. This projection is optimum for classification and is supervised. A PCA+LDA combination, termed Fisherfaces, was introduced in [3] and was proven robust to illumination changes [3,18].

2.3 Elastic Bunch Graph Matching

EBGM [4] assumes that the positions of certain facial features are known for each training image. The image regions around these features are convolved with 40 complex 2D Gabor kernels. The resulting 80 coefficients constitute the Gabor jet for each facial feature. The Gabor jets for all facial features are grouped in a graph, the Face Graph, where each jet is a node and the distances between facial features are the weights on the corresponding vertices. The information in the Face Graph is all that is needed for recognition; the image itself is discarded. All Face Graphs from the training images are combined in a stack-like structure called the Face Bunch Graph (FBG). Each node of the FBG contains a list of Gabor jets for the corresponding facial feature from all training images, and the vertices are now weighted with the average distances across the training set. The positions of the facial features in the testing images are unknown; EBGM estimates them based on the FBG. Then a Face Graph can be constructed for each testing image based on the estimated positions. The Face Graph of each testing image is compared with the FBG to determine the training image it is most similar to, according to some jet-based metric. In [4], a number of such metrics is proposed, most of which can also be used for the feature estimation step. Our results in [19] indicate that Displacement Estimation Local Search is the best choice for facial feature localization; for the actual identification stage, Displacement Estimation Grid Search yields the best recognition rate.

2.4 Pseudo Two-Dimensional Hidden Markov Models

Face recognition using HMM is based on approximating blocks from the face image with a chain of states of a stochastic model [2]. For the pseudo 2D HMM the image blocks are extracted by scanning the face from left to right, top to bottom with an overlap both in horizontal and vertical direction. Pixel intensities do not lead to robust features, as they are susceptible to illumination changes and other detrimental effects. A transformation like the 2D Discrete Cosine Transform attenuates those distorting effects, leading to better performance. A pseudo 2D HMM model of hidden states is obtained by linking left-right 1D HMM models with vertical super-states. For the training of each class the Baum-Welch algorithm is used. In the recognition phase the class that gives the highest value for the probability of the observation sequence of the testing image, given the class model, is considered the

most likely to be the true identity of the testing face. Our results in [20] indicate that close cropping the faces and using mixture of three Gaussians for the states enhance performance.

2.5 Correlation Filters

Face recognition can be performed by cross correlating a face image with a suitable filter and processing the output. Many correlation filters have been proposed [5]; amongst them, the Minimum Average Correlation Energy filter (MACE) is reported to perform best. It reduces the large sidelobes by minimizing the average correlation energy plane while at the same time satisfying the correlation peak constraints at the origin. These constraints result in the correlation plane close to zero everywhere except at a location of a trained object, where a sharp peak appears. For recognition, the output plane is searched for the highest point and that value, as well as the values of its surrounding points, is used to determine the class that the face belongs to. We have implemented and tested the MACE correlation filter algorithm in [20].

2.6 Laplacianfaces

The Eigenfaces and Fisherfaces algorithms both use linear projection to a subspace aiming to preserve the global structure of the face. Laplacianfaces [6] on the other hand, is an algorithm that uses optimal linear approximations to the eigenfunctions of the Laplace-Beltrami operator. By aiming to preserve the local structure of the face, Laplacianfaces attempts to attenuate the unwanted variations resulting from changes in lighting, facial expression and pose. In that, Laplacianfaces shares many of the properties of nonlinear projection algorithms. We have implemented and tested the Laplacianfaces algorithm in [20].

3 Experimental results

The evaluation methodology proposed in [13] is followed to assess the performance of the face recognition algorithms outlined in the previous section. In particular, the HumanScan database [22] is utilized. This database offers 20 feature points on each face; including the two eye centers. The eye centers are used to register the faces, whereas all the rest are utilized in the training images for EBGM. To introduce the needed registration errors, Gaussian noise is added to the co-ordinates of the eye centers. Then the noisy co-ordinates are used in the face normalization process.

The variance of the added noise in each co-ordinate is controlled so that the RMS error is a given percentage of the eye distance. This way, different versions of the normalized database are obtained, with the RMS eye perturbation ranging from zero (ideal normalization) to 7% of the eye distance. To get a feeling of the meaning of eye perturbation range selected for the comparison of the algorithms, the RMS eye perturbation that results from the face detection scheme presented in [15] is between 4.7% (full scale HumanScan faces) and 7.2% (HumanScan faces decimated to 10 pixels eye distance). Hence, depending on the resolution, state-of-the-art face

detection systems are expected to have an RMS eye perturbation in the upper half of the selected testing range.

The effect of the eye perturbation on the faces is to shift, scale and rotate them in a random manner. Examples of the erroneously registered faces are shown in Fig. 1. While some of them are registered moderately well, there are others where the registration errors have caused only a tilted part of the face to lie in the template or significant background portion to enter the template.




Fig. 1. Effect of 4% RMS eye perturbation on 21 randomly selected HumanScan faces from a single person

The algorithms are tested with 5 training faces per person. This is a moderate number for face recognition systems that are intended for smart room or human-machine interfaces, but a rather large one for security applications. The 400 runs described in [13] are carried out, for each of the six algorithms and each of the eight different degrees of eye perturbation, allowing for a statistically sound analysis of the results. The median value for each group of 400 runs is reported in Fig. 2. Note that two variants of the eigenfaces algorithm are tried: the classic one and that without the three eigenvectors that correspond to the largest eigenvalues (PCAw/o3).

From the results depicted in Fig. 2, the correlation filters algorithm is the most sensitive to face registration errors, followed by the Laplacianfaces. The rest of the methods are close up to 2% RMS eye perturbation, a value that is very optimistic for real-time face detectors. This is indicated in the boxplots [23] of Fig. 3, where a test of the statistical significance of the different median values is performed. The PCA variants and pseudo-2D HMM have statistically similar performance; LDA and EBGM are somewhat better, with LDA being the winner at 5% significance level. At 3% RMS eye perturbation, the subspace projection methods start deteriorating. Of these, the first to deteriorate is the PCAw/o3. This is to be expected, since the first few eigenvectors encode only illumination changes (hence unwanted within class variation) only if the registration is ideal. This is also the reason that PCAw/o3 performs badly in face databases without ideal face registration, with pose variations and with pronounced expression changes [13]. At large RMS eye perturbations, above 5% RMS, such as those expected in real-time detectors operating on low resolution faces, the performance of supervised and unsupervised subspace projection becomes comparable; Eigenfaces even exceed Fisherfaces in performance

at 7% RMS, indicating ill-training. This is again to be expected; five training images are not able to capture the within class variation introduced by the large registration errors.




Fig. 2. Probability of misclassification of the face recognition algorithms under varying percentage of RMS eye perturbation and 5 training faces per person




Fig. 3. Boxplots of the PMC of the linear subspace projection algorithms, the pseudo-2D HMM and the EBGM for 2% RMS eye perturbation and 5 training faces per person. The difference in the median values (horizontal lines) is statistically significant at the 5% significance level only if the notches of the boxes do not overlap

The pseudo-2D HMM and EBGM algorithms perform noticeably better than the rest at 3% RMS or higher eye perturbation. Both deteriorate only minimally up to 3% RMS, with EBGM being marginally better up to 4% RMS. After that EBGM deteriorates fast, while pseudo-2D HMM retains its performance remarkably well. The robustness of EBGM up to 4% RMS is due to the search for the feature points inherent to the recognition algorithm. This search can compensate for medium displacements of the features from the expected positions. As for the exceptional robustness of the pseudo-2D HMM algorithm to face registration errors, it is due to the stochastic nature of the model of the face. The transitions between the states (overlapping blocks of the face image) are probabilistic, and can account for the perturbation of the feature locations. This is the reason for the success of the pseudo-2D HMM algorithm in the ORL database [2] it was tested when introduced, where no particular care has been taken in registering the ORL faces.

4 Conclusions

Since in most applications, face recognition algorithms do not perform stand-alone but within detection and recognition systems, the face registration errors resulting from imperfect face detections become important. In this paper we have quantitatively analyzed the impact of face registration errors on six face recognition algorithms and one of their variants, to establish ranges of face detection accuracy where some of them are optimal. These are summarized in Table 1. The experimentally established robustness of the various face recognition algorithms is also backed-up by an analysis of the features of the algorithms.

Table 1. Optimum face recognition algorithms depending on face registration accuracy

Face registration accuracy (% of eye distance, RMS)	Optimum algorithm
[0,2]	LDA
(2,4]	EBGM
> 4	2D HMM

References

1. Turk, M., Pentland, A.: Eigenfaces for Recognition. *J. Cognitive Neuroscience.* (1991) 71-86
2. Samaria, F., Harter, A.: Parametrisation of a Stochastic Model for Human Face Identification. 2nd IEEE Workshop on Applications of Computer Vision. (1994) 138-142
3. Belhumeur, P., Hespanha, J., Kriegman, D.: Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection. *IEEE Trans. Pattern Analysis and Machine Intelligence.* 7 (1997) 711-720
4. Wiskott, L., Fellous, J-M., Krueger, N., Malsburg, C.: Face Recognition by Elastic Bunch Graph Matching. In: Jain, L.C. et al. (eds.) *Intelligent Biometric Techniques in Fingerprint and Face Recognition.* CRC Press (1999) 355-396

5. Xie, C., Vijaya Kumar, B. V. K. , Palanivel, S., Yegnanarayana, B.: A Still-to-Video Face Verification System Using Advanced Correlation Filters. *International Conference on Biometric Authentication*. (2004) 102-108
6. He, X., Yan, S., Hu, Y., Niyogi, P., Zhang, H.-J.: Face Recognition Using Laplacianfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence*. 3 (2005) 328-340
7. Bartlett, M., Movellan, J., Sejnowski, T.: Face Recognition by Independent Component Analysis. *IEEE Trans. Neural Networks*. 6 (2002) 1450-1464
8. Yang, J., Frangi, A., Yang, J.-Y., Zhang, D., Jin, Z.: KPCA Plus LDA: A Complete Kernel Fisher Discriminant Framework for Feature Extraction and Recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*. 2 (2005) 230-244
9. Georgiades, A., Belhumeur, P., Kriegman, D.: From Few to Many: Illumination Cone Models for Face Recognition under Variable Lighting and Pose. *IEEE Trans. Pattern Analysis and Machine Intelligence*. 6 (2001) 643-660
10. Liu, X., Chen, T., Vijaya Kumar, B.V.K.: On Modeling Variations For Face Authentication. *International Conference on Automatic Face and Gesture Recognition*, (2002) 369-374
11. Philips, P., Moon, H., Rizvi, S., Rauss, P.: The FERET Evaluation Methodology for Face-Recognition Algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence*. 10 (2000) 1090-1104
12. Phillips, P. et al.: Overview of the Face Recognition Grand Challenge. *CVPR*. (2005) 947-954
13. Pnevmatikakis, A., Polymenakos, L.: A Testing Methodology for Face Recognition Algorithms. *2nd Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms*. Edinburgh. (2005)
14. Li, S. Z., Lu, J.: Face Detection, Alignment and Recognition. In Medioni, G., Kang, S. (eds.) *Emerging Topics in Computer Vision*. Prentice-Hall (2004) 455
15. Pnevmatikakis, A., Polymenakos, L.: An Automatic Face Detection and Recognition System for Video Streams. *2nd Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms*. Edinburgh. (2005)
16. Hsu, R.-L., Abdel-Mottaleb, M., Jain, A. K.: Face Detection in Color Images. *IEEE Trans. Pattern Analysis and Machine Intelligence*. 5 (2002) 696-706
17. Stiefelhagen, R., Yang, J., Waibel, A.: Tracking eyes and monitoring eye gaze. *Workshop on Perceptual User Interfaces*. Banff, Canada. (1997)
18. Pnevmatikakis, A., Polymenakos, L.: Comparison of Eigenface-Based Feature Vectors under Different Impairments. *Int. Conf. Pattern Recognition*. (2004) 296-300
19. Stergiou, A.: Elastic Bunch Graph Matching Face Recognition: Performance and Comparison with Subspace Projection Methods. *MSc Thesis, Athens Information Technology*. (2004)
20. Rentzeperis, E.: A Comparative Analysis of Face Recognition Algorithms: Hidden Markov Models, Correlation Filters and Laplacianfaces vs. Linear subspace projection and Elastic Bunch Graph Matching. *MSc Thesis, Athens Information Technology*. (2005)
21. Soldatos, J., Polymenakos, L., Pnevmatikakis, A., Talantzis, F., Stamatias, K., Carras, M.: Perceptual Interfaces and Distributed Agents supporting Ubiquitous Computing Services. *Eurescom*. (2005)
22. Jesorsky, O., Kirchberg, K., Frischholz, R.: Robust Face Detection Using the Hausdorff Distance. in Bigun, J., Smeraldi, F. (eds.): *Audio and Video based Person Authentication*. Springer (2001) 90-95
23. McGill, R., Tukey, J. W., Larsen, W. A.: Variations of Boxplots. *The American Statistician*. (1978) 12-16