

Semi-supervised PCA-Based Face Recognition Using Self-training

Fabio Roli and Gian Luca Marcialis

Dept. of Electrical and Electronic Engineering, University of Cagliari
Piazza d'Armi, 09123 Cagliari, Italy
{roli, marcialis}@diee.unica.it

Abstract. Performances of face recognition systems based on principal component analysis can degrade quickly when input images exhibit substantial variations, due for example to changes in illumination or pose, compared to the templates collected during the enrolment stage. On the other hand, a lot of new unlabelled face images, which could be potentially used to update the templates and re-train the system, are made available during the system operation. In this paper a semi-supervised version, based on the self-training method, of the classical PCA-based face recognition algorithm is proposed to exploit unlabelled data for off-line updating of the eigenspace and the templates. Reported results show that the exploitation of unlabelled data by self-training can substantially improve the performances achieved with a small set of labelled training examples.

1 Introduction

Face recognition based on Principal Component Analysis (PCA) operates in two distinct stages: the enrolment stage and the recognition, or authentication, stage [1]. In the enrolment stage, a set of face images is acquired for each user. Then PCA is applied to the enrolled face images to compute the eigenspace associated to the selected eigenvalues. For each user, the enrolled images are projected to the eigenspace and a face template, often computed as the mean of the projected faces, is stored in a gallery. In order to account for variations in the appearance of a user, multiple templates, associated, for example, to different poses, can be stored in the user's gallery. In the recognition stage the input image is projected to the above eigenspace and the system associates the identity of the nearest template to this image. As Uludag et al. pointed out [2], the large intra-class variability of face images, due for example to changes in illumination or pose, can make the templates acquired during the enrolment stage poorly representative of the images to be recognized, so resulting in poor recognition performances. Increasing the size of the galleries of users' templates does not necessarily solve the problem, as the intra-class variability of face images is often due to aging, appearance, expression, and illumination changes which cannot be captured during a single enrolment stage over a short period of time [3]. On the other hand, a lot of new unlabelled face images are made available during the system operation over the time. These new data may be

exploited to update the templates and re-train the system. It is reasonable to hypothesize that exploiting unlabelled test data to update the eigenspace and the templates may improve the performances on new test data. Liu et al. showed that, in PCA-based face recognition, unlabelled test data can be exploited to update the eigenspace and improve the system's performance [3]. The potential benefits of the automatic, or semi-automatic, updating of the galleries of the users' templates using labelled and unlabelled data have been pointed out in [4-6].

The design of a face recognition system using a small set of labelled faces, collected during the initial enrolment stage, and a large batch of unlabelled face images, collected during the system operation, can be naturally regarded as a problem of semi-supervised learning. In fact, semi-supervised learning deals with the design of recognition systems using both labelled (possibly few) and unlabelled training examples [7]. However, the use of semi-supervised learning methods for face recognition has been poorly investigated so far, the only exception being the work of Balcan et al. [8].

The goal of this paper is to give a contribution to the development of semi-supervised face recognition systems. To this end, the use of a well known semi-supervised learning method, namely, the self-training method, is proposed to develop a semi-supervised version of the standard PCA-based face recognition algorithm. Reported experimental results show that the exploitation of a batch of unlabelled test images by self-training can substantially improve the performance of a PCA-based face recognition system on new test data in comparison with the ones achievable with a small set of labelled training examples.

2 Self-training for Semi-supervised PCA-Based Face Recognition

In this section, first the main steps of the classical supervised PCA-based face recognition method are summarized. Then a semi-supervised version of this method, based on the self-training technique, is proposed.

2.1 Supervised PCA-Based Face Recognition

Enrolment Stage

For each user, a set, usually small, of reference images is acquired. These images constitutes the training set $D_l=(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_l})$ containing n_l labelled data, where each vector $\mathbf{x}_i, i=1 \dots n_l$, represents a face image. An identity label $I_k, k=1 \dots K$, is associated to each vector. Then PCA is applied to the enrolled face images by the following steps:

Principal components computation

The covariance matrix is computed for the data set D_l . The eigenvalues and eigenvectors of the covariance matrix are computed. A set of eigenvalues with the highest variance is selected. The eigenvectors associated to the selected eigenvalues are named "eigenfaces" and define the matrix W of the "principal components".

Data transformation

The data in D_l are projected to the above eigenspace using the principal components

matrix W : $y_i = W^t(x_i - \mu)$, $i=1 \dots n_l$, where $\mu = \frac{1}{n_l} \sum_{i=1}^{n_l} x_i$ is the mean vector of data in D_l .

Template creation

For each user, a face template, often computed as the mean of the projected faces, is stored in a gallery.

Recognition Stage

In the so called “closed set” scenario, the input image is projected to the above eigenspace by the matrix W and the system associates the identity of the nearest template to this image.

2.2 Semi-supervised PCA-Based Face Recognition

Given a set D_l (usually, small) of labeled data, and a set D_u (usually, large) of unlabelled data, semi-supervised methods aim to design recognition systems using both sets. Several semi-supervised methods have been proposed so far, based on expectation-maximization algorithms, self-training, co-training, active learning, transductive learning, and graph-based techniques. We refer the reader to [7] for an overview on semi-supervised learning methods. For the purposes of this work, we summarize here the so called self-training method. In self-training a classifier is initially trained using the labeled data set D_l . This classifier is then used to assign pseudo-class labels to a subset of the unlabelled examples in D_u , and such pseudo-labeled data are added to D_l . Usually, the unlabelled data classified with the highest confidence are selected to increase D_l . Then the classifier is re-trained using the increased data set D_l . As the convergence of this simple algorithm can not be guaranteed in general, the last two steps are usually repeated for a given number of times or until some heuristic convergence criterion is satisfied.

Due to its easy use, we chose self-training as technique to develop a semi-supervised version of the classical PCA-based face recognition algorithm. The main steps of the algorithm developed are summarized in Figure 1. In the enrolment stage a set D_l of labelled images is collected and used to compute the matrix W of the principal components. Data are then projected to the eigenspace defined by W , and, for each user, a face template is computed as mean of the projected faces. During the on-line recognition stage an unlabelled batch of data D_u , to be used in the semi-supervised stage, is collected over a given period of time; to this end, recognition labels obtained as system’s outputs are obviously disregarded. It should be noted that the designer should select a period of time that allows collecting a *representative* batch of data. In our experiments, the data set D_u contains the same identity classes, with the same priors, of the set D_l . The semi-supervised stage is then performed off-line. It is assumed that the recognition system carries out this stage either when it is not operating (e.g., during the night) or using a separate processing unit which allows carrying out it in parallel with the recognition stage. The semi-supervised stage starts by projecting the unlabelled data to the eigenspace defined by the matrix W computed in the enrolment stage. Then the semi-supervised cycle goes through N

iterations. For each iteration, a pseudo-label is assigned to each data in D_u . The pseudo-label coincides with the label of the nearest template. For each identity class, a set P_k is defined which contains all the data pseudo-labelled as belonging to the class I_k . Then, for each identity class, the face image pseudo-labelled with the highest confidence, that is, the one nearest to the class template, is selected from the set P_k , and this image is added to the training set D_l and removed from the set D_u . Therefore, only one pseudo-labelled image per class, the most confident one, is added to the user's gallery during every iteration of the semi-supervision cycle. It is worth noting that the use of a less conservative confidence threshold, which allows adding more than one class example per iteration, could be investigated with the goal to speed up the learning. The increased training set is then used to update the eigenspace by re-computing the principal components matrix W . The labelled and unlabelled data are projected to the updated eigenspace. Finally, the class templates are updated using the augmented training set. In our algorithm the templates are simply the "mean" faces, but more sophisticated methods, based, for example, on clustering, could be used for template update [2]. In our experiments we performed ten iterations of the semi-supervised algorithm. Other stopping criteria could be investigated and used. For example, the iterations could be stopped when no unlabelled data can be added to the users' galleries. To sum up, after the initial enrolment stage, the system goes through on-line recognition phases, with collection of unlabelled data, and off-line semi-supervised phases to update the eigenspace and the templates. Performances on new test data are expected to improve after each semi-supervised stage in comparison with the ones achievable using the previous, non-updated, eigenspace and templates.

3 Experimental Results

3.1 Data Set and Goal of Experiments

The goal of our experiments was to evaluate the capability of the developed semi-supervised algorithm to exploit a batch of unlabelled test images, collected during a given session of the system operation, in order to improve the performance on novel test data in comparison with the ones achievable using the initial training data. To this end, we carried out experiments with the AR data set [9]. This data set contains frontal view faces with different facial expressions, illumination conditions, and occlusions (sun glasses and scarf). Each person participated in two acquisition sessions, separated by two weeks time. Each session is made up of seven images per person. We selected 100 subjects (50 males and 50 females), and manually cropped face images and, after histogram stretching and equalization, resized them at 40x40 pixels. Various subsets of first session images were used for the enrolment stage and the collection of unlabelled data set D_u . Second session images were always used as separate test set.

This experimental setting simulates well the acquisition of unlabelled data during a given session of system's operation, and the performance evaluation, after the semi-supervised stage, on new test data belonging to a separate acquisition session.

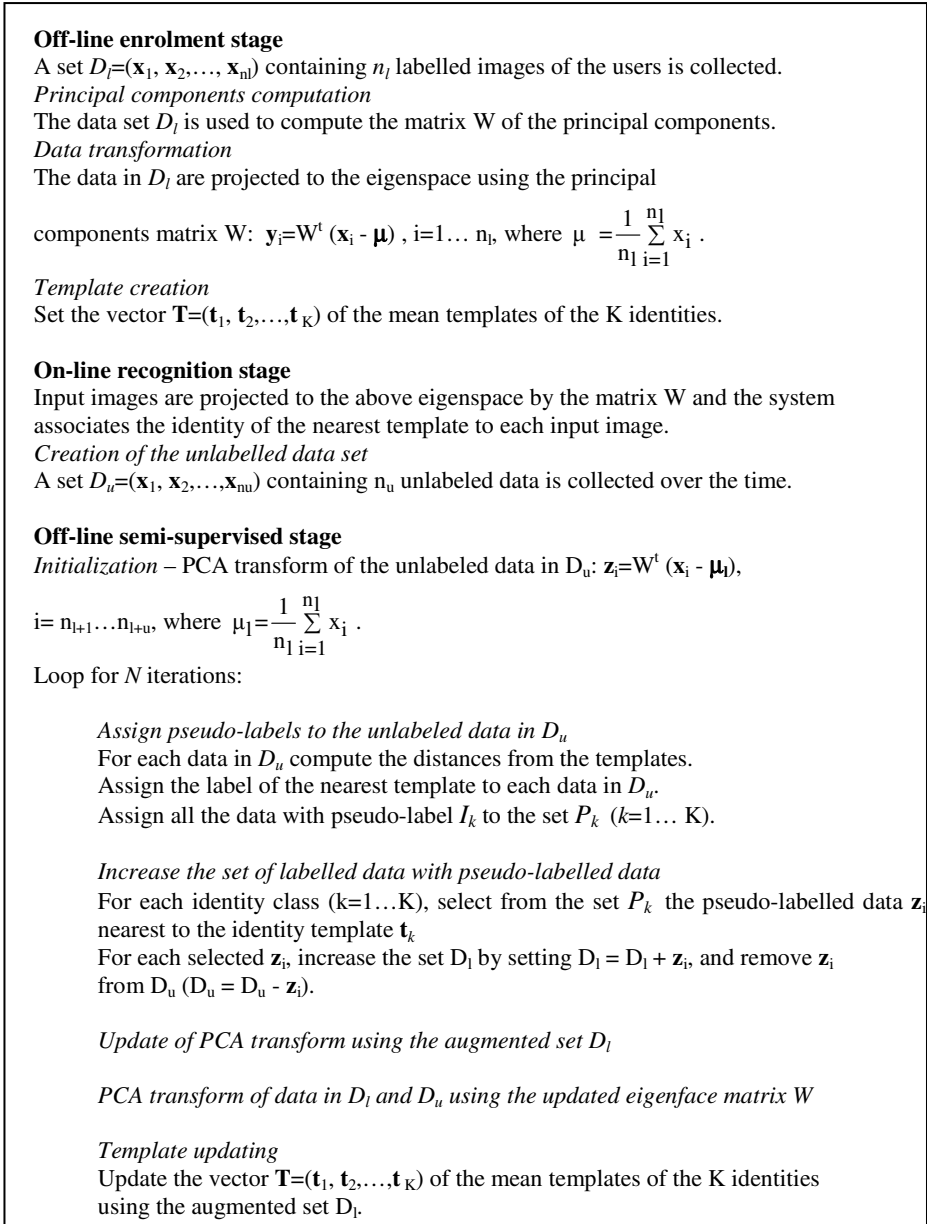


Fig. 1. The semi-supervised PCA-based algorithm using self-training

3.2 Results

We assessed performances for different numbers of templates created during the enrolment stage. Reported performances are averaged on five different trials. For each

trial, we randomly selected one, two or three images of the first session as templates. The remaining images of the first session were used as unlabelled data set D_u . Second session images were always used as separate test set. Figure 2 shows the percentage accuracy on the test set and the unlabelled data set D_u averaged on five trials. For this experiment, only one face template per person was used. The remaining first-session images (six per person) were used as unlabelled data set D_u . Performances are shown as function of the number of unlabelled data added to the training set during the ten iterations of the semi-supervised stage. Around one-hundred pseudo-labeled data were added during every iteration. Figure 2 shows that the accuracy on second-session images used as test set is very low, around 18%, when unlabelled data are not used. The average accuracy increases substantially with the number of unlabelled data exploited by self-training. The maximum accuracy obtained for test data, around 62%, is anyway low due to the use of a single template per person and the large differences between first and second session images. But the increase of accuracy from 18% to 62% shows the benefits of the exploitation of unlabelled data. The accuracy is much higher for the unlabelled data, as the set D_u contains images of the first session which are more similar to the initial templates. It should be noted the practical interest of results obtained on the unlabelled data set. Unlabelled data are input data collected during a given period of time of the system's operation. Figure 2 shows that the initial accuracy on such batch of data is low, around 50%. After the semi-supervised phase based on self-training the accuracy increases to 89%. This results points out that the semi-supervision process can allow to improve the recognition results previously achieved on a batch of input data. For example, in a video surveillance scenario such a system re-training could allow improving the identification results stored in the data base the day before.




Fig. 2. Average accuracy on the test and unlabelled data sets as function of the number of unlabelled data used in the semi-supervised algorithm of Figure 1

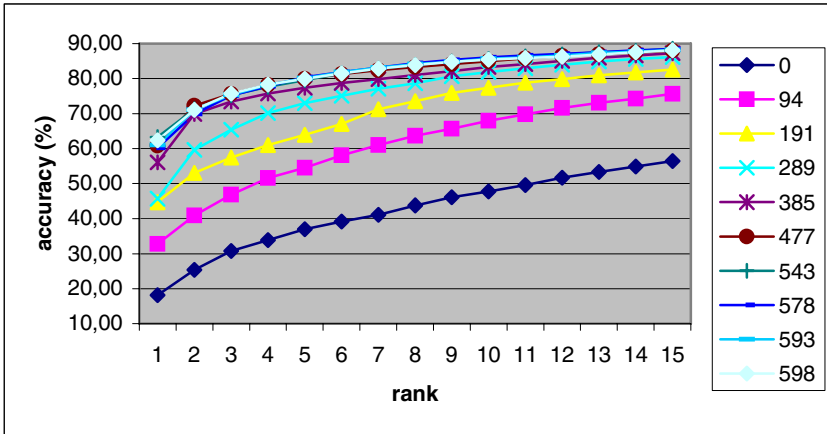


Fig. 3. Rank-order curves for the test data set. Each curve refers to an iteration of the semi-supervised cycle and is labeled with the number of unlabelled data which were pseudo-labeled and added to the training set.




Fig. 4. Average accuracy on the test set as function of the number of iterations of the semi-supervised phase and the number of templates

We also assessed performances in terms of the so called rank-order curves, that is, we assessed the percentage accuracy, averaged on five trials, achieved by considering the fifteen template faces nearest to the input face (Figure 3). Figure 3 clearly shows the improvement of accuracy with the increase of the number of unlabelled data added to the training set during the ten iterations of the semi-supervised cycle.

To investigate how the performance of the semi-supervised algorithm depends on the number of templates collected during the enrolment stage, we performed experiments with one, two and three images of the first session as templates. Figure 4 depicts the three accuracy curves related to the different numbers of templates. Each curve provides the percentage average accuracy as function of the number of iterations

of the semi-supervised phase. Around one hundred unlabelled data are added to the training set during every iteration. As one could expect, the greatest benefit of the use of unlabelled data is obtained when only one template per class is used.

Finally, we analyzed how the galleries of users' templates were updated and increased by our semi-supervised algorithm. Figure 5 depicts four examples of the update of users' galleries. For each gallery, the first image on the left is the initial training image used as face template. The remaining images are the unlabelled images which were pseudo-labeled and added to the gallery during the ten iterations of our semi-supervised algorithm. Due to the strict confidence threshold used in our algorithm (i.e., only one pseudo-labelled image per class, the most confident one, is added to the user's gallery during every iteration), images very similar to the template are initially added to the galleries. Then images which exhibit variations of expressions are added, and, in some cases, this causes wrong images to be added to the gallery. It is worth noting that different number of images can be added to different users' galleries.

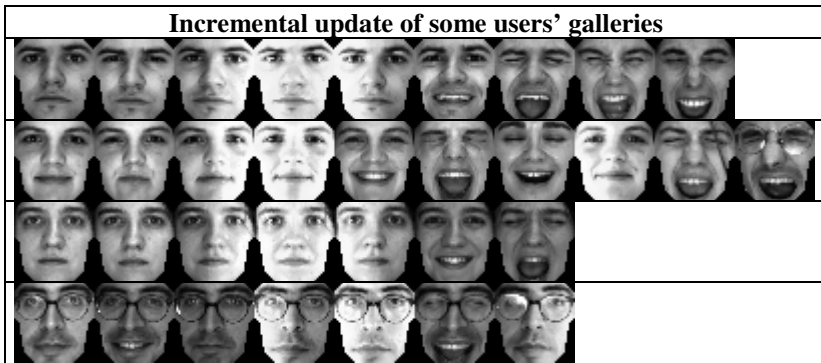


Fig. 5. Examples of the incremental update of users' galleries during the ten iterations of the semi-supervised algorithm. The first image on the left is the initial training image used as face template.

4 Conclusions

Performances of face recognition systems based on principal component analysis strongly depends on how much the face images collected during the enrolment stage are representative of face images to be recognized. Unfortunately, representative face images are difficult to be captured during a single enrolment stage over a short period of time. More representative images might be collected during the system operation over the time. The exploitation of such unlabelled data naturally demands for semi-supervised face recognition systems. Accordingly, we developed and assessed by experiments a semi-supervised version, based on self-training, of the classical PCA-based face recognition algorithm. Reported results show that the exploitation of a batch of unlabelled images by self-training can substantially improve the performance of a PCA-based face recognition system on new test data in comparison with the ones achievable with a small set of labelled training examples. Although final conclusions

cannot be drawn on the basis of this work, we believe that a first step towards the development of semi-supervised face recognition systems has been done. As directions for our future work, the use of other semi-supervised learning methods will be investigated. In addition, as the good performances of the proposed system, based on a simple self-training mechanism, are, in a sense, surprising, the conditions under which it can be expected to work well will be analysed further by experiments and, if possible, theoretically.

References

- [1] M. Turk and A. Pentland, Eigenfaces for Face Recognition, *Journal of Cognitive Neuroscience*, 3 (1) 71-86, 1991.
- [2] U. Uludag, A. Ross and A. K. Jain, Biometric template selection and update: a case study in fingerprints, *Pattern Recognition*, Vol. 37, No. 7, pp. 1533-1542, July 2004.
- [3] X. Liu, T. Chen, S.M. Thornton, Eigenspace updating for non-stationary process and its application to face recognition, *Pattern Recognition*, Special issue on Kernel and Subspace Methods for Computer Vision, 2003, 1945–1959.
- [4] R. Sukthankar, R. Stockton, Argus: the digital doorman, *IEEE Intelligent Systems*, March/April 2001, pp. 14-19.
- [5] K. Okada, C. von der Malsburg, Automatic video indexing with incremental gallery creation: integration of recognition and knowledge acquisition, *Proceedings of ATR Symposium on Face and Object Recognition*, pages 153-154, Kyoto, July 19-23, 1999.
- [6] K. Okada, L. Kite, C. von der Malsburg, An adaptive person recognition system, *Proceedings of the IEEE International Workshop on Robot-Human Interactive Communication*, pages 436-441, Bordeaux/Paris, 2001.
- [7] X. Zhu, Semi-supervised learning literature survey, Technical report, Computer Sciences TR 1530, Univ. Wisconsin, Madison, USA, Jan. 2006.
- [8] M. F. Balcan, A. Blum, P.P. Choi, J. Lafferty, B. Pantano, M.R. Rwebangira, X. Zhu, Person Identification in Webcam Images: An Application of Semi-Supervised Learning, *ICML2005 Workshop on Learning with Partially Classified Training Data*, Bonn, Germany, 7 August 2005.
- [9] A. Martinez, R. Benavente, The AR Face Database. *CVC Technical Report #24*, June, 1998.