Extension of Robust Principal Component Analysis for Incremental Face Recognition

Haïfa Nakouri and Mohamed Limam
Institut Supérieur de Gestion, LARODEC Laboratory
University of Tunis, Tunis, Tunisia

Keywords: Image alignment, Robust Principal Component Analysis, Incremental RPCA.

Abstract: Face recognition performance is highly affected by image corruption, shadowing and various face expressions. In this paper, an efficient incremental face recognition algorithm, robust to image occlusion, is proposed. This algorithm is based on robust alignment by sparse and low-rank decomposition for linearly correlated images, extended to be incrementally applied for large face data sets. Based on the latter, incremental robust principal component analysis (PCA) is used to recover the intrinsic data of a sequence of images of one subject. A new similarity metric is defined for face recognition and classification. Experiments on five databases, based on four different criteria, illustrate the efficiency of the proposed method. We show that our method outperforms other existing incremental PCA approaches such as incremental singular value decomposition, add block singular value decomposition and candid covariance-free incremental PCA in terms of recognition rate under occlusions, facial expressions and image perspectives.

1 INTRODUCTION

In the last two decades, face recognition has been an active research area within the computer vision and the pattern recognition communities. Since an original input image space has a very high dimension, dimensionality reduction techniques are usually performed before classification. Principal Component Analysis (PCA) is one of the most popular representation methods for computer vision applications mainly face recognition. Usually, PCA is performed in the batch mode, where all training data are used to calculate the PCA projection matrix. Once the training data have been fully processed, the learning process stops. In case we want to incorporate additional data into an existing PCA projection matrix, the matrix has to be retained with all training data. Therefore, such system is hard to scale up. An incremental version for PCA is a straightforward solution to overcome this limitation.

Incremental PCA (IPCA) has been studied for more than two decades yielding many methods, which are specially useful when not all observations are simultaneously available. The aim of the IPCA approach is to do not consider all available observations more than once even when new data are eventually upcoming. New data can be used to incrementally update a previous computation. Such an approach reduces storage requirements and large problems become computationally feasible.

The performance of IPCA methods is evaluated with face recognition standard databases (Hall et al., 2000; Weng et al., 2003; Huang et al., 2009; Hall et al., 2002). However, one of their major drawbacks is that they cannot simultaneously handle large illumination variations, image corruptions and partial occlusions that often occur in real face data (e.g., self shadowing, hats, sunglasses, scarf, incomplete face data, etc), hence inducing important appearance variation. These image variations can be considered as outliers or errors regarding the original face image of one subject. Although classical PCA is effective against the presence of small Gaussian noise in the data, it is highly sensitive to even sparse errors of very high magnitude.

On the other hand, it is known that well-aligned face images of a person, under varying illumination, lie very close to a low-dimensional linear subspace. However, in practice, images deviate from this situation due to self shadowing, different angles and occlusions. Thus, we have a set of coherent images corrupted by essentially sparse errors. In order to efficiently extract low-rank face images from corrupted and distorted ones, we should first model those corruption factors and seek efficient ways to eliminate

Nakouri H. and Mohamed L..
Extension of Robust Principal Component Analysis for Incremental Face Recognition.
DOI: 10.5220/0004288305490555
In *Proceedings of the International Conference on Computer Vision Theory and Applications* (VISAPP-2013), pages 549-555
ISBN: 978-989-8565-47-1

them. The robust PCA (RPCA) (Wright et al., 2009) is a powerful tool to get rid off such errors and retrieve cleaner images potentially better suited for computer vision application, namely face recognition.

In this paper, we propose an incremental method for robust face recognition under various conditions based on RPCA. The proposed method handles both misalignment and occlusion problems on face images. In order to improve the recognition process and based on RPCA (Wright et al., 2009), we eliminate corruptions and occlusion in original face images. Besides, the incremental aspect of our face recognition method handles the memory constraint and computational cost of a large data set. To measure the similarity between a query image and a sequence of images of one person, we define a new similarity metric. To evaluate the performance of our method, experiments on the AR (Martinez and Benavente, 1998), ORL (Samaria and Harter, 1994), PIE (Sim et al., 2002), YALE (Belhumeur et al., 1997) and FERET (Phillips et al., 1998) databases and a comparison with other incremental PCA methods namely incremental singular value decomposition (SVD) (Hall et al., 2002), add block SVD (Brand, 2006) and candid covariance-free incremental PCA (Weng et al., 2003) are conducted. We also compare our method to a face recognition method based on batch robust PCA, denoted by face recognition RPCA (FRPCA) (Wang and Xie, 2010).

This paper is organized as follows. In Section 2, we introduce the RPCA method, incremental RPCA (IRPCA) and our face Recognition method, denoted by new incremental RPCA (NIRPCA). Finally, in Section 3, we present our experimental results.

2 FACE RECOGNITION BASED ON IRPCA

2.1 Robust Image Alignment by Sparse and Low-rank Decomposition

Peng et al., (Peng et al., 2010) proposed robust alignment by sparse and low-rank decomposition for linearly correlated images (RASL). It is a scalable optimization technique for batch linearly correlated image alignment. One of its objectives is to robustly align a dataset of human faces based on the fact that if faces are well-aligned, they show efficient low-rank structure up to some sparse corruptions. Even perfectly aligned images may not be identical, but at least they lie near a low-dimensional subspace (Basri and Jacobs, 2003). To the best of our knowledge, RASL is the first method that uses a trade-off between rank minimization and alignment of image data. Hence, the idea is to search for a set of transformations τ such that the rank of the transformed images becomes as small as possible and at the same time the sparse errors are compensated. Generally, the applied transformation is the 2D affine transform, where we implicitly assume that the face of a person is approximately on a plane in 3D-space.

2.2 Incremental Robust Principal Component Analysis (IRPCA)

RPCA algorithm is aimed to recover the low-rank matrix A from the corrupted observations $D = A + E$, where corrupted entries E are unknown and the errors can be arbitrarily large but assumed to be sparse. More specifically, in face recognition, E is a sparse matrix because it is assumed that only a small fraction of image pixels are corrupted by large errors (e.g., occlusions). Hence, being able to correctly identify and recover the low structure A could be very interesting for many computer vision applications namely face recognition.

We assume that we have m subjects and each one has n face images. Although RASL can give a very accurate alignment for faces (Peng et al., 2010), it is not applicable when the total number of images $m \times n$ denoted by l is very large. Wu et al., (Wu et al., 2011) proposed an extension to RASL from l to L where $L >> l$, by reformulating the problem using a "one-by-one" alignment approach. This incremental alignment can be summarized in three steps. First, l frames are selected to be aligned with batch RASL method producing a low-rank summary A^*. In the second step, the $(l+1)^{th}$ image is aligned with A^* which contains the information of the previously aligned l images. Finally, the second step is repeated for the remaining images, regardless of the size of the data set.

We denote by I_i^j, A_i^j, E_i^j the corrupted ,observed, face image, the original face image and the error of the j^{th} image of the i^{th} subject, respectively. Then, we have $I_i^j = A_i^j + E_i^j$, where i denotes the subject and j its corresponding image such that $i = 1,\ldots,m, \quad j = 1,\ldots,n$. Let:

$$vec : \mathbb{R}^{w\times h} \rightarrow \mathbb{R}^{(w\times h)\times 1}, \tag{1}$$

be a function which transforms a $w \times h$ image matrix into a $(w \times h) \times 1$ vector by stacking its columns to have $vec(I_i^j) = vec(A_i^j) + vec(E_i^j)$. Assuming that we have m subjects and each one has n images, we define for the i^{th} subject:

$$D_i := [vec(I_i^1)|\ldots|vec(I_i^n)] = A_i + E_i \tag{2}$$

where

$$A_i := [vec(A_i^1)|\ldots|vec(A_i^n)] \in \mathbb{R}^{(w\times h)\times n} \quad (3)$$

and

$$E_i := [vec(E_i^1)|\ldots|vec(E_i^n)] \in \mathbb{R}^{(w\times h)\times n}, \quad (4)$$

with $i = 1,\ldots,m$. D_i is formed by stacking the n image vectors of the i^{th} subject, A_i and E_i are the corresponding original images matrix and the error matrix, respectively. Since all images of the same person are approximately linearly correlated, A_i is regarded as a low-rank matrix and E_i is a large matrix but sparse. As proved in (Peng et al., 2010), the original face A_i can be efficiently recovered from the corrupted face image D_i. It is well known that if images are well-aligned, they should present a low-rank structure up to some sparse errors (e.g., occlusions) (Peng et al., 2010). Therefore, we search for a set of transformations, $\tau = \tau_1,\ldots,\tau_n$, such that the rank of the transformed images becomes as small as possible, and simultaneously images become as well-aligned as possible. Many works (Peng et al., 2010; Candès et al., 2011) prove that practical misalignment can be modeled as a certain transformation $\tau^{-1} \in \mathbb{G}$ acting on the two-dimensional domain of an image I. $\mathbb{G}$ is assumed to be a finite dimensional group that has a parametric representation, such as the similarity group $SE(2)\times\mathbb{G}_+$ or the $2D$ affine group $Aff(2)$. In this paper we assume that $\mathbb{G}$ is the affine group.

Consider that after performing a batch image alignment using the RASL method, we obtain the set of transformation τ, the low rank matrix A and the error matrix E. The IRASL algorithm is given in Algorithm 1.

2.3 The Proposed Face Recognition Algorithm based on IRPCA

In this section, we introduce a new face recognition algorithm called NIRPCA based on the one-by-one RASL method discussed in Section 2.2. We also define a new similarity metric, which is used for measuring the similarity between a query image and a sequence of images, and later used for our main face recognition application.

Given, m different subjects where each one has n training images I_i^j, $i = 1,\ldots,m$; $j = 1,\ldots,n$, we need to classify a query image I_{n+1}. The basic idea of the proposed algorithm is to recover the sparse error E_{n+1} of the test image I_{n+1} using approximated A_{n+1}. Once E_{n+1} is recovered, we use it to compute a similarity metric for our face recognition algorithm NIRPCA.

Let D be the observation matrix of each subject in the training set, and A its recovered low-rank matrix

Algorithm 1: Incremental robust alignment by sparse and low-rank decomposition.

- **INPUT:** Images $I_{n+1} \in \mathbb{R}^{w\times h}$, initial transformations τ_{n+1} in certain parametric group $\mathbb{G}$, weight $\mu > 0$
- **WHILE** not converged **DO**
  step1: compute Jacobian matrix w.r.t transformation τ_{n+1}:

  $$J \leftarrow \frac{\partial}{\partial\zeta}\left(\frac{vec(I_{n+1}\circ\zeta)}{\|vec(I_{n+1}\circ\zeta)\|_2}\right)\Big|_{\zeta=\tau_{n+1}};$$

  step2: wrap and normalize the image:

  $$I_{n+1}\circ\tau_{n+1} \leftarrow \left[\frac{vec(I_{n+1}\circ\tau_{n+1})}{\|vec(I_{n+1}\circ\tau_{n+1})\|_2}\right];$$

  step3: solve the linearized convex optimization:

  $$(x^*,\Delta\tau_{n+1}^*) \leftarrow \arg\min_{x,\Delta\tau_{n+1}} \frac{1}{2}\parallel I_{n+1}\circ\tau_{n+1} + J\Delta\tau_{n+1} - \tilde{A}x \parallel_2^2 + \mu \parallel x \parallel_1;$$

  step4: update transformation :

  $$\tau_{n+1} \leftarrow \tau_{n+1} + \Delta\tau_{n+1}^*;$$

- **END WHILE**
- **OUTPUT:** solution x, τ_{n+1}.

and E the error matrix. For the i^{th} subject in the training set, $i = 1,\ldots,m$, we have D_i, A_i and E_i as given in Equations (2), (3) and (4).

Let I_{n+1} be a new occluded face image that we need to classify. According to RASL method (Peng et al., 2010; Candès et al., 2011), this new observation can be decomposed as:

$$I_{n+1}\circ\tau_{n+1} = A_{n+1} + E_{n+1}, \quad (5)$$

where, τ_{n+1} is the transformation applied to corrupted image I_{n+1} to resolve the image misalignment. A_{n+1} is the occlusion-free image and E_{n+1} is the error image representing the occlusion. Our objective is to estimate τ_{n+1}, A_{n+1} and E_{n+1}. To solve Equation (5), we propose to use the low rank matrix A^* generated by the RASL method. As indicated in Algorithm 1, the one-by-one alignment approach proposed by (Wu et al., 2011) computes τ_{n+1} and x having the low rank matrix A^* and the corrupted face image I_{n+1}. Since A^* is a low-rank matrix, let $\tilde{A}$ denote the summary of low rank data resulting from batch RASL, such that $\tilde{A} \in \mathbb{R}^{m\times rank(A^*)}$, where the columns are equal to $rank(A^*)$, i.e., the independent columns of A^*. The vector x of dimension $rank(A^*)$ is represented as an approximation of the coefficients of the linear combination of $\tilde{A}$ and A_{n+1}. Hence, an approximation of A_{n+1} is obtained by the following equation

$$A_{n+1} = \tilde{A} * x. \quad (6)$$

Once we estimate A_{n+1} and by using Equation (5), we estimate the error vector E_{n+1} (standing for the occlusion on the I_{n+1} test image) and use it to compare the similarity between the test image and the stored images. Let the similarity metric be

$$M_i = \| E_i - E_{n+1} \|_2 , \qquad i = 1, \ldots, n \tag{7}$$

where M_i measures the similarity between the input corrupted test image I_{n+1} and the a class of images of the i^{th} subject. M_i is the Euclidean distance between the input image I_{n+1} and a class of images belonging to the i^{th} subject. If the query image I_{n+1} belongs to the i^{th} subject, D_i contains the image of the same subject so that the assumption that A_i is linearly correlated and the low-rank condition can be satisfied. In this case, the parameters of E_i are small, and then M_i should be small, otherwise, the value of M_i is relatively large. For face recognition, the test image I_{n+1} is recognized as the subject which has the smallest value of M_i. If the similarity M_i is greater than a given threshold α, the face image I_{n+1} is not recognized. It will be considered as a new subject and a new class subject A_{n+1} will be created and added to summary $\tilde{A}$. Otherwise, $M_i \leq \alpha$, and I_{n+1} is recognized as belonging to the i^{th} subject and we just should update the class of the i^{th} subject in the $\tilde{A}$ summary with A_{n+1}. The classification criterion α is set to be 0.5. We should also set at most 15 face images per subject to be stored in the $\tilde{A}$ summary in order to keep a sort of balance between the training faces. The NIRPCA algorithm is summarized in Algorithm 2. Figure 1 illustrates the recovered original images A and error images E of two different subjects using RPCA (see Section 2.1). In our case, RPCA is used to recover original (A) and error (E) images of the training set. As for the test set, we start using our NIRPCA method defined in Algorithm 2. Figure 2 shows an example of face recognition using NIRPCA method.

3 EXPERIMENT RESULTS AND DISCUSSION

In this section, we evaluate the performance of the proposed face algorithm NIRPCA based on the AR (Martinez and Benavente, 1998), ORL (Samaria and Harter, 1994), PIE (Sim et al., 2002), YALE (Belhumeur et al., 1997) and FERET (Phillips et al., 1998) databases. All testing images are grayscale and normalized. We precisely use the canonical size of the images rather than the original one, based on the eye corner locations. Table 1 provides information about these databases.

Algorithm 2 : Face Recognition based on IRASL (NIRASL).

- **INPUT:** Images $I_{n+1} \in \mathbb{R}^{w \times h}$, initial transformations τ^0_{n+1}, in certain parametric group $\mathbb{G}$, $\tilde{A}$ summary, weight $\mu > 0$
- **1.** (τ_{n+1}, x) = IRASL $(\tilde{A}, \tau^0_{n+1})$ (Algorithm 1)
- **2.** $A_{n+1} = \tilde{A} * x$
- **3.** $E_{n+1} = I_{n+1} \circ \tau_{n+1} - A_{n+1}$ $(Equation\,(5))$
- **4. FOR** each subject I_i

  $M_i = \| E_i - E_{n+1} \|_2, \qquad i = 1, \ldots, m \quad (Equation\,(7))$

  END FOR
- **5.** $M_{min} \leftarrow min(M_i), \qquad i = 1, \ldots, m$
- **6. IF** $M_{min} > \alpha$ **THEN**

  $\tilde{A} \leftarrow (\tilde{A} | A_{n+1})$
  $m \leftarrow m + 1$
  $n \leftarrow n + 1$

  ELSE

  $\tilde{A} \leftarrow (\tilde{A} | A_{n+1})$
  $n \leftarrow n + 1$

  END IF
- **OUTPUT:** Subject class of face image I_{n+1}.

Table 1: Face databases used for experiments.

Database	AR	ORL	PIE	YALE	FERET
Original size	64×64	92×112	640×486	320×243	80×80
Canonical size	50×50	90×90	110×130	120×140	80×80
Number of subjects	70	40	68	15	47
Number of total images	420	200	680	162	465

3.1 Comparison with Standard Incremental Face Recognition Methods

In this section, our method is tested on five different databases as shown in Table 2. For each database, $\frac{2}{3}$ of the images are used for the training set and the remaining $\frac{1}{3}$ of the images are randomly selected for the test set. In these databases, face images are captured under varied conditions such as illumination and shadowing levels, facial expression, different face perspectives and with/without occlusion (sunglasses, scarf, etc.). Besides, our method is compared to other well-known incremental face recognition methods,

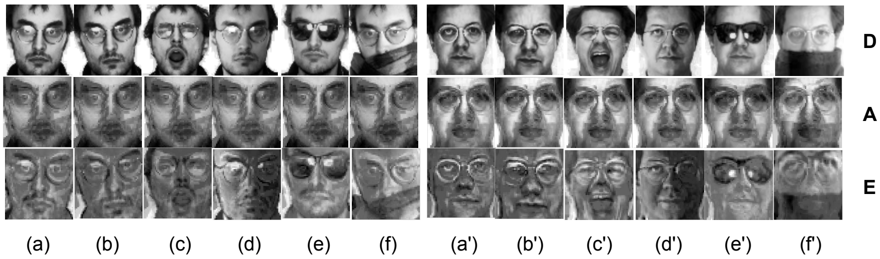

Figure 1: Images recovered by Robust PCA (RPCA), (a) - (f) are six images of two different subjects. D are original subjects, A and E are recovered low-rank and error faces. Sunglasses, scarf and face expressions are successfully removed. These images correspond to the training face images of the corresponding subjects.

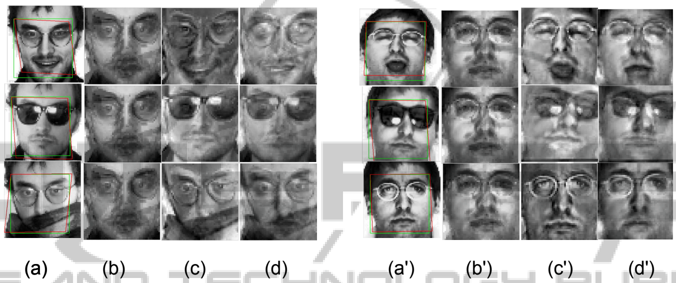

Figure 2: Face recognition using NIRPCA method for two different subjects. (a), (a'): test Face (on the canonical frame). (b), (b'): approximated low-rank matrix A using IRPCA. (c), (c'): recovered error matrix E. (d), (d'): the reconstructed face. The classification measure M_i computed for these test faces is less than α, when $\alpha = 0.5$. $M_a = 0.406, 0.428, 0.274$ respectively for each test face of subject (a) and $M_{a'} = 0.364, 0.394, 0.383$ respectively for each test face of subject (a'). Hence both test faces are correctly classified.

i.e., ISVD (Hall et al., 2002), ABSVD (Brand, 2006) and CCIPCA (Weng et al., 2003). In this section, occluded images in the AR database (with sunglasses or scarfs) are omitted in the evaluation. Performance on images with occlusions is considered in Section 3.3 and the accuracy rate is given in Table 2. These

Table 2: Accuracy rate of different incremental face recognition algorithms.

%	IPCA	ABSVD	CCIPCA	NIRPCA
AR	79.71	73.44	77.43	**86.42**
ORL	72.68	72.23	76.82	**77.39**
PIE	72.17	75.89	79.64	**88.72**
YALE	78.12	80.27	85.47	**90**
FERET	76.09	80.63	85.44	**89.58**

results indicate that our method achieves the best performance, specifically on images under different facial expressions, head positions and shadow levels. This can be explained by the application of IRPCA which recovers original face and removes shape and expression variations better than other algorithms. In fact, various angles of a face image or different head positions of a subject can be reduced to an image distortion problem. Our algorithm can solve this problem since IRPCA approximates the original image matrix A_{n+1}, the sparse error matrix E_{n+1} and a transformation τ_{n+1} at the same time and for each input image I_{n+1}. Thus applying the recovered affine transformation τ_{n+1} to the input image will generate a well-aligned and distortion-free face image.

3.2 Incremental Face Recognition with Different Numbers of Training Images

In this section, we evaluate the performance when the number of training images per subject varies. Based on the AR database, K images are randomly selected as the training set, and the remaining constitute the test set, where $K = 1, \ldots, 10$. Figure 3 shows the variation trend of different face recognition algorithms and NIRPCA, in terms of recognition rate. This experiment shows that beyond 3 images per subject NIRPCA achieves the best performance. In fact, with less than 4 images per subject, the D matrix, presented in Section 2.1, is itself a low-rank matrix, hence it cannot generate an exact and efficient error matrix E. Accordingly, approximated E_{n+1} and similarity measure are both inaccurate, which proves the slow growth of NIRPCA's accuracy rate with few images per subject.

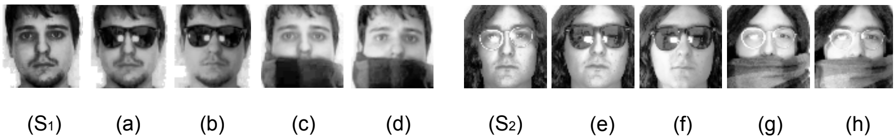

Figure 4: Some corrupted images with occlusions in the AR database of two subjects S_1 and S_2. Images (a), (b), (e) and (f) are occluded images with sunglasses. Images (c), (d), (g) and (h) are occluded images with a scarf.

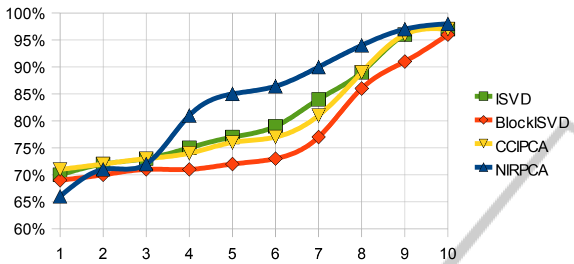

Figure 3: Accuracy rate of face recognition methods with different number of training images per subject.

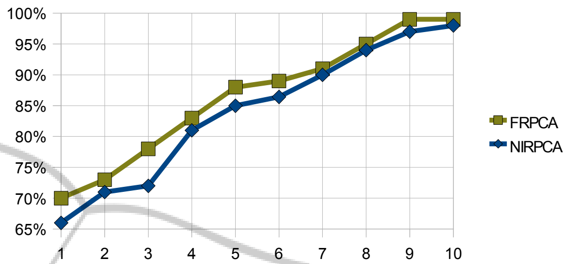

Figure 5: Accuracy rate variation of FRPCA and NIRPCA with different numbers of training images per subject.

3.3 Performance on Images with Occlusions

In practical face recognition applications, occlusions (e.g., sunglasses or scarfs on faces) could not be avoided. Robust and efficient face recognition algorithms should achieve good performance when faces are practically occluded. We use the occluded face images in the AR database to test the performance of different algorithms. In this experiment, 70 subjects are selected for the dataset, 5 occluded images and 6 non-occluded ones for each one. Some corrupted images with occlusions are shown in Figure 4.

These results show us that occluded images considerably reduce the performance of IPCA, ABSVD and CCIPCA, where the best recognition rate is only about 20%. Whereas a high recognition rate is still maintained by our algorithm NIRPCA as shown in Table 3. This is due to the robustness of our method regarding occlusions while other incremental algorithms cannot efficiently remove disruption caused by corruption on images.

Table 3: Accuracy rate of different incremental face recognition algorithms with occluded test faces.

%	IPCA	ABSVD	CCIPCA	NIRPCA
AR	28.89	24.36	22.27	**58.67**

3.4 Comparison with Batch FRPCA

Wang and Xie (Wang and Xie, 2010) presented a face recognition method, FRPCA, based on batch RPCA (Peng et al., 2010) presented in Section 2.1. Figure 5 shows the variation trend between the FRPCA method and NIRPCA with reference to various number of training sets. Results show that the accuracy of our method with different numbers of training face images is very close to that of FRPCA. This experiment shows high accuracy of the recovered low rank face image A_{n+1} approximated by Equation 6.

3.5 Runtime Comparison

In this section, we compare the runtime of NIRPCA with other IPCA-based face recognition methods namely ISVD, ABSVD and CCIPCA. This experiment is carried on 70 different subjects, 13 images for each subject from the AR database. Table 4 summarizes the runtime results. Although the recognition rate of NIRPCA is the best, its runtime is slower than that of CCIPCA due to iterative linearization of the convex optimization using the splitting Bregman method (Goldstein and Osher, 2009) (step 3 in Algorithm 1). On the other hand, the runtime of NIRPCA is faster than those of face recognition methods based on ISVD or ABSVD. When we decrease the number of principal components, the runtimes of ISVD and ABSVD are similar to that of NIRPCA , but their face recognition ratio decreases.

Table 4: Runtime of different incremental face recognition algorithms.

%	IPCA	ABSVD	CCIPCA	NIRPCA
Runtime (s)	96.89	135.41	26.53	60.29

4 CONCLUSIONS

In this paper, we proposed a new face incremental recognition method based on one-by-one RASL. Our method is robust to sparse corruptions on face images and performed experiments on different databases show its efficiency. The advantages of our method is that it handles many aspects of image variations such as face expression, image shadowing and various angles. Above all, unlike other existing incremental face recognition methods, the proposed method handles efficiently corrupted images, mainly the occluded ones. Experiment based on five different databases show that our proposed method has better accuracy rates. Further experiments can be extended to video images, so that it could be used in real face recognition applications. However, for video images, important image preprocessing work, such as face detection, should be done before the recognition step.

REFERENCES

Basri, R. and Jacobs, D. (2003). Lambertian reflectance and linear subspaces. *PAMI*, 25:218–233.

Belhumeur, P. N., Hespanha, J. P., and Kriegman, D. J. (1997). Eigenfaces vs. fisherfaces: Recognition using class specific linear projection.

Brand, M. (2006). Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and its Applications*, 415(1):20–30.

Candès, E., Li, X., Ma, Y., and Wright, J. (2011). Robust principal component analysis? *Journal of the ACM*, 58(3).

Goldstein, T. and Osher, S. (2009). The split bregman method for l1-regularized problems. *SIAM J. Img. Sci.*, 2(2):323–343.

Hall, P., Marshall, D., and Martin, R. (2000). Merging and splitting eigenspace models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22:1042–1049.

Hall, P., Marshall, D., and Martin, R. (2002). Adding and subtracting eigenspaces with eigenvalue decomposition and singular value decomposition. *Image Vision Comput.*, 20(13-14):1009–1016.

Huang, D., Yi, Z., and Pu, X. (2009). A new incremental pca algorithm with application to visual learning and recognition. *Neural Process. Lett.*, 30:171–185.

Martinez, A. and Benavente, R. (1998). The ar face database. Technical report.

Peng, Y., Ganesh, A., Wright, J., Xu, W., and Ma, Y. (2010). Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images. In *CVPR*, pages 763–770.

Phillips, P., Wechsler, H., Huang, J., and Rauss, P. (1998). The FERET Database and Evaluation Procedure for Face Recognition Algorithms. *Image and Vision Computing*, 16(5):295–306.

Samaria, F. S. and Harter, A. C. (1994). Parameterisation of a stochastic model for human face identification.

Sim, T., Baker, S., and Bsat, M. (2002). The cmu pose, illumination, and expression (pie) database.

Wang, Z. and Xie, X. (2010). An efficient face recognition algorithm based on robust principal component analysis. In *Proceedings of the Second International Conference on Internet Multimedia Computing and Service*, ICIMCS '10, pages 99–102, New York, NY, USA. ACM.

Weng, J., Zhang, Z. Y., and Hwang, W. (2003). Candid covariance-free incremental principal component analysis. *In PAMI*, 25:1034–1040.

Wright, J., Ganesh, A., Rao, S., Peng, Y., and Ma, Y. (2009). Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization. In *NIPS*, pages 2080–2088.

Wu, K. K., Wang, L., Soong, F. K., and Yam, Y. (2011). A sparse and low-rank approach to efficient face alignment for photo-real talking head synthesis. In *ICASSP*, pages 1397–1400.