

Eigenspace Updating for Non-Stationary Process and Its Application to Face Recognition

Xiaoming Liu¹ Tsuhan Chen^{1*} Susan M. Thornton²

¹ Department of Electrical and Computer Engineering
Carnegie Mellon University
5000 Forbes Ave. Pittsburgh, PA 15213
xiaoming@andrew.cmu.edu tsuhan@cmu.edu

² Sonic Foundry, Inc.
12300 Perry Highway
Wexford, PA 15090
sthornton@sonicfoundry.com

Abstract

In this paper, we introduce a novel approach to modeling non-stationary random processes. Given a set of training samples sequentially, we can iteratively update an eigenspace to manifest the current statistics provided by each new sample. The updated eigenspace is derived more from recent samples and less from older samples, controlled by a number of decay parameters. Extensive study has been performed on how to choose these decay parameters. Other existing eigenspace updating algorithms can be regarded as special cases of our algorithm. We show the effectiveness of the proposed algorithm with both synthetic data and practical applications for face recognition. Significant improvements have been observed in recognizing face images with different variations, such as pose, expression and illumination variations. We also expect the proposed algorithm to have other applications in active recognition and modeling.

Keywords: Principal Component Analysis, Eigenspace Updating, Non-Stationary Process, Face Recognition.

1 Introduction

The principal component analysis (PCA) [1] has attracted much attention among image analysis researchers. The basic idea is to represent images or image features in a transformed space where the individual features are uncorrelated. The orthonormal basis functions for this space, called the eigenspace, are the eigenvectors of the covariance matrix of the images or image features. PCA gives the optimal representation of the images or image features in terms of the mean square error.

PCA has been used extensively by researchers in many fields, such as data compression [2], feature extraction [3], and object recognition [4]. One of the successful applications of PCA is introduced in [5] and later made popular by Turk and Pentland [6]. They projected a face image into an eigenspace that is

trained from all the images of multiple subjects, and performed face recognition in this eigenspace.

There are mainly two kinds of approaches to training the eigenspace in the literature. The first approach is to compute the eigenvectors given a set of training samples simultaneously, which we refer to as *batch training*. In this approach, PCA can be computationally intensive when it is applied to the image domain. The power method [7] is one approach to efficiently determining the dominant eigenvectors. Instead of determining all the eigenvectors, the power method obtains only the dominant eigenvectors, i.e., eigenvectors associated with the largest eigenvalues. Researchers have explored how to perform PCA more efficiently. Turk and Pentland [6] proposed to calculate the eigenvectors of an inner-product matrix instead of a covariance matrix, which is efficient in the case that the number of training samples is less than the dimension of the feature space. The second approach is to iteratively re-calculate the existing eigenvectors by taking the training samples one by one, which is called *eigenspace updating*, proposed for its computational efficiency compared to the batch training approach [8]. A few other researchers have proposed different

* Corresponding author. Tsuhan Chen, Department of Electrical and Computer Engineering, Carnegie Mellon University, 5000 Forbes Ave. Pittsburgh, PA 15213-3890, U.S.A. Tel.: +1-412-268-7536, fax: +1-412-268-3890.

eigenspace updating methods [9][10][11], and suggested some interesting applications, such as salient view selections [11].

PCA is originally created to model multidimensional random variables. When extended to modeling random *processes*, traditional PCA works well as long as the random process under consideration is stationary. For non-stationary random processes, PCA needs to be adapted to model the time-varying statistics. To some extent, existing eigenspace updating methods [8][9][10][11] already accomplish this implicitly, because they all try to compute eigenvectors iteratively as samples come in one by one. In this paper, we propose a new eigenspace updating method to consider non-stationary random processes explicitly by modifying these methods in the literature. In particular, our method puts more weights on recent samples than on older samples by using certain decay parameters, while traditional methods consider all samples equally and hence cannot effectively represent the most recent statistics of the data. Effectively these decay parameters function as forgetting factors for older samples. We have also studied how to choose the decay parameters to model the time-varying statistics in the least mean squares sense.

For decades human face recognition has been an active topic in the field of object recognition. A general statement of this problem can be formulated as follows: given still or video images of a scene, identify one or more persons in the scene using a stored database of faces [13]. There are mainly two kinds of face recognition systems: the feature matching-based approach and the template matching-based approach. In the latter, applying PCA to obtain a face model (also known as the eigenface approach [6]) plays a fundamental role. It has good performance for the case of frontal face recognition with reasonable constraints on illumination, expression variations, etc. However, in practical applications, when large variations, which may be due to aging, changes in expressions and poses, and variations caused by illumination, etc., appear in the test face images, the traditional PCA algorithm degrades quickly in performance. Although some methods in the literature work well for the specific variations being studied, their performance degrades rapidly when other variations are present [14]. In order to approach this general problem, we propose an updating-during-recognition scheme, which tries to make the recognition system more intelligent by learning the variations over time using the test images. In this paper, we utilize our eigenspace updating method to learn the time-varying statistics of the face images and eventually enhance the recognition performance. We use the individual PCA approach [15][16], instead of the universal PCA approach [6], as a baseline of our face recognition system.

1.1 Previous works

The eigenspace updating has a number of advantages. First, using the updating algorithm, we can determine the eigenvec-

tors more efficiently than the batch training approach [8]. Second, the updating algorithm allows the construction of the eigenspace via a procedure that uses less storage, so it renders feasible some previous inaccessible problems, such as the training of a huge image data set [10]. Third, the availability of the training data may be constrained in some applications, such as online training [11]. In that case, we have to iteratively perform PCA instead of waiting for all the training data to be available.

Murakami and Kumar [8] proposed the first eigenspace updating algorithm. They iteratively generated the covariance matrix or the inner-product matrix and calculated the eigenvectors whenever there is a new training sample. Chandrasekaran et al. [11] proposed an eigenspace updating algorithm by performing the Singular Value Decomposition (SVD) on the data matrix, instead of the covariance matrix. They showed the effectiveness of their algorithm in the 3D object representation from 2D images, which is useful in active recognition and exploration. Levy and Lindenbaum [9] also proposed an SVD-based eigenspace updating method by using QR decomposition to reduce computation and memory demand. They also pointed out the option of using forgetting factors for an image sequence. However, these three methods have limitations when used for classification because they assume that the samples have zero mean. Hall et al. [10] addressed this issue and proposed an eigenspace updating method where the mean is updated based on existing samples, and removed before PCA is performed on the covariance matrix. They showed that for classification, better performance could be obtained by their approach compared to those in [8][11]. Hall et al. [12] also proposed an algorithm to efficiently merge and split eigenspace models. All of these existing eigenspace updating methods, original designed to model the statistics of random vectors, also work for stationary random processes. They, however, cannot handle non-stationary processes to represent the time-varying statistics effectively, which is what our proposed algorithm tries to address.

Some researchers have tried to utilize the information provided by the new data when the system is being used to enhance the performance of face detection and tracking [17][18][19][20]. For example, Kurita et al. [17] iteratively updated the prior probabilities of the face location in the previous frames, which can guide and speed up the face detection on the current frame. Edwards et al. [18] described a method of updating the first order global estimation of the identity, which was integrated with an optimal tracking scheme. Wu et al. [19] proposed to build a subspace representation via the Gram-Schmidt orthogonalization procedure for the purpose of video compression. Weng et al. [20] proposed to incrementally derive discriminating features from training video sequences. Compared to these prior works, our work extends the idea of updating to face recognition and results in an updating-during-recognition scheme.

1.2 Paper outline

In Section 2, we introduce our eigenspace updating algorithms. Two algorithms aimed at different application scenarios are presented in detail. When PCA is applied to a high-dimensional image domain, we use the algorithm based on updating an inner-product matrix. Otherwise, an algorithm based on updating a covariance matrix can be used.

In our eigenspace updating method, the decay parameters play a key role on how well the time-varying statistics can be modeled. Thus in Section 3, we theoretically and experimentally show how to choose the decay parameters based on the knowledge of model statistics.

In Section 4, we address the issue of iteratively updating the individual eigenspace for face recognition. Given one test image, we can use it to update the eigenspace when we have high confidence for its recognition result. Also we propose to use a twin-subspace scheme to alleviate some limitations and enhance the face recognition performance.

Experimental results using eigenspace updating methods are presented in Section 5. We conduct experiments on face databases containing different variations, such as poses, expressions and illuminations. We show that better performance can be obtained in these applications by using our eigenspace updating method.

In Section 6, we discuss the related issues for our work, such as the video-based recognition and the high order statistical model for face sequences. We provide conclusions in Section 7. Also in the appendix we compare our mean estimation algorithm with the Kalman filter [21].

2 Eigenspace updating with decay

2.1 Updating based on the covariance matrix

Suppose there is a random process $\{\mathbf{x}_n\}$, where n is the time index, \mathbf{x}_n is a column vector in a d -dimensional space, of which we want to find the eigenspace. Each sample will be available sequentially over time. If this random process is stationary, we can estimate its mean by the following equation:

$$\hat{\mathbf{m}} = \frac{\mathbf{x}_n + \mathbf{x}_{n-1} + \cdots + \mathbf{x}_1}{n}$$

If $\{\mathbf{x}_n\}$ is a non-stationary random process, which implies that it has a time-varying mean $\hat{\mathbf{m}}_n$, we propose to estimate the mean at time n as:

$$\hat{\mathbf{m}}_n = \frac{\mathbf{x}_n + \alpha_m \mathbf{x}_{n-1} + \alpha_m^2 \mathbf{x}_{n-2} + \cdots}{1 + \alpha_m + \alpha_m^2 + \cdots} \quad (1)$$

where α_m is the decay parameter. It controls how much the previous samples contribute to the estimation of the current mean. Since α_m is in the range of 0 to 1, we have:

$$1 + \alpha_m + \alpha_m^2 + \cdots = \frac{1}{1 - \alpha_m} \quad (2)$$

Using (2) in (1), the resulting equation can be simplified to:

$$\hat{\mathbf{m}}_n = \alpha_m \hat{\mathbf{m}}_{n-1} + (1 - \alpha_m) \mathbf{x}_n \quad (3)$$

This equation reveals that based on the current sample and the previously estimated mean, we can obtain the new estimated mean in a recursive manner. How to choose α_m mainly depends on the knowledge of the random process. Note that α_m controls how fast we want to forget about the old samples. Therefore, if the statistics of the random process change fast, we choose a small α_m . If the statistics change slowly, a large α_m may perform better. In the next section, we will introduce how to choose these decay parameters based on the statistical knowledge of the samples.

After the mean of the random process has been estimated, we can estimate the covariance matrix, $\hat{\mathbf{C}}_n$, at each time n by:

$$\hat{\mathbf{C}}_n = \frac{(\mathbf{x}_n - \hat{\mathbf{m}}_n)(\mathbf{x}_n - \hat{\mathbf{m}}_n)^T + \alpha_v (\mathbf{x}_{n-1} - \hat{\mathbf{m}}_{n-1})(\mathbf{x}_{n-1} - \hat{\mathbf{m}}_{n-1})^T}{1 + \alpha_v + \alpha_v^2 + \cdots} + \alpha_v^2 (\mathbf{x}_{n-2} - \hat{\mathbf{m}}_{n-2})(\mathbf{x}_{n-2} - \hat{\mathbf{m}}_{n-2})^T + \cdots$$

where α_v is also a decay parameter, which is chosen based on how fast the covariance of a random process is changing. Now we can rewrite $\hat{\mathbf{C}}_n$ in a similar manner as $\hat{\mathbf{m}}_n$:

$$\hat{\mathbf{C}}_n = \alpha_v \hat{\mathbf{C}}_{n-1} + (1 - \alpha_v) (\mathbf{x}_n - \hat{\mathbf{m}}_n)(\mathbf{x}_n - \hat{\mathbf{m}}_n)^T \quad (4)$$

Since we obtain $\hat{\mathbf{C}}_n$ at time n , we can perform PCA for $\hat{\mathbf{C}}_n$ and obtain the corresponding eigenvectors. We keep N eigenvectors corresponding to the N largest eigenvalues. In the recursive updating process, we only need to store the mean vector $\hat{\mathbf{m}}_n$ and the covariance matrix $\hat{\mathbf{C}}_n$. All the previous training samples can be discarded.

2.2 Updating based on the inner-product matrix

In many applications, PCA is applied directly to the image domain, such as face recognition. Suppose the face image has a size of 32 by 32, then the covariance matrix of an image set

would be 1024 by 1024. It is very inefficient to store and update it using the algorithm introduced in Section 2.1. To solve this problem, we propose an updating algorithm based on the inner-product matrix.

Suppose at time n , we already have performed PCA for the random process at time $n-1$. Thus we have eigenvectors, $\boldsymbol{\varphi}_{n-1}^{(i)}$, and eigenvalues, $\lambda_{n-1}^{(i)}$, of the covariance matrix, $\hat{\mathbf{C}}_{n-1}$. We can write:

$$\hat{\mathbf{C}}_{n-1} = \lambda_{n-1}^{(1)} \boldsymbol{\varphi}_{n-1}^{(1)} \boldsymbol{\varphi}_{n-1}^{(1)T} + \lambda_{n-1}^{(2)} \boldsymbol{\varphi}_{n-1}^{(2)} \boldsymbol{\varphi}_{n-1}^{(2)T} + \dots + \lambda_{n-1}^{(d)} \boldsymbol{\varphi}_{n-1}^{(d)} \boldsymbol{\varphi}_{n-1}^{(d)T}$$

where eigenvalues, $\lambda_{n-1}^{(i)}$, have been sorted in the decreasing order and the superscript (i) indicates the order of eigenvalues. By retaining only the first Q eigenvectors (with the largest eigenvalues), we can approximate $\hat{\mathbf{C}}_{n-1}$ as

$$\hat{\mathbf{C}}_{n-1} \approx \lambda_{n-1}^{(1)} \boldsymbol{\varphi}_{n-1}^{(1)} \boldsymbol{\varphi}_{n-1}^{(1)T} + \lambda_{n-1}^{(2)} \boldsymbol{\varphi}_{n-1}^{(2)} \boldsymbol{\varphi}_{n-1}^{(2)T} + \dots + \lambda_{n-1}^{(Q)} \boldsymbol{\varphi}_{n-1}^{(Q)} \boldsymbol{\varphi}_{n-1}^{(Q)T} \quad (5)$$

The criteria for choosing Q vary, and depend on practical applications. We have tried three methods: (a) Fix Q to be a constant value; (b) Set a minimum threshold, and keep the first Q eigenvectors whose eigenvalues are larger than this threshold; (c) Keep the eigenvectors corresponding to the largest eigenvalues, such that a specific fraction of energy in the eigenvalue spectrum is retained. These methods result in different computational complexity for the updating algorithm.

Now we can use (3) to estimate the mean at time n . By substituting $\hat{\mathbf{C}}_{n-1}$ in (4) with (5), we obtain

$$\hat{\mathbf{C}}_n \approx \alpha_v \lambda_{n-1}^{(1)} \boldsymbol{\varphi}_{n-1}^{(1)} \boldsymbol{\varphi}_{n-1}^{(1)T} + \alpha_v \lambda_{n-1}^{(2)} \boldsymbol{\varphi}_{n-1}^{(2)} \boldsymbol{\varphi}_{n-1}^{(2)T} + \dots + \alpha_v \lambda_{n-1}^{(Q)} \boldsymbol{\varphi}_{n-1}^{(Q)} \boldsymbol{\varphi}_{n-1}^{(Q)T} + (1 - \alpha_v) (\mathbf{x}_n - \hat{\mathbf{m}}_n) (\mathbf{x}_n - \hat{\mathbf{m}}_n)^T$$

An equivalent formulation as above is that

$$\hat{\mathbf{C}}_n \approx \mathbf{B}_n \mathbf{B}_n^T$$

where

$$\mathbf{B}_n = \left[\sqrt{\alpha_v \lambda_{n-1}^{(1)}} \boldsymbol{\varphi}_{n-1}^{(1)} \quad \dots \quad \sqrt{\alpha_v \lambda_{n-1}^{(Q)}} \boldsymbol{\varphi}_{n-1}^{(Q)} \quad \sqrt{1 - \alpha_v} (\mathbf{x}_n - \hat{\mathbf{m}}_n) \right] \quad (6)$$

Based on the \mathbf{B}_n matrix, an inner-product matrix can be formulated as

$$\mathbf{A}_n = \mathbf{B}_n^T \mathbf{B}_n$$

Furthermore, \mathbf{A}_n can be described by the following equations:

$$\begin{aligned} (A_n)_{i,j} &= \alpha_v \sqrt{\lambda_{n-1}^{(i)} \lambda_{n-1}^{(j)}} \delta_{ij}; \quad i, j = 1, 2, \dots, Q \\ (A_n)_{i,Q+1} &= (A_n)_{Q+1,i} = \sqrt{\alpha_v (1 - \alpha_v) \lambda_{n-1}^{(i)}} (\mathbf{x}_n - \hat{\mathbf{m}}_n)^T \boldsymbol{\varphi}_{n-1}^{(i)}; \quad i = 1, \dots, Q \\ (A_n)_{Q+1,Q+1} &= (1 - \alpha_v) (\mathbf{x}_n - \hat{\mathbf{m}}_n)^T (\mathbf{x}_n - \hat{\mathbf{m}}_n). \end{aligned} \quad (7)$$

Since the matrix \mathbf{A}_n is usually a small matrix with the size of $Q+1$ by $Q+1$, we can determine its eigenvectors $\boldsymbol{\psi}_n$ by a direct method, which satisfies

$$\mathbf{A}_n \boldsymbol{\psi}_n^{(i)} = \mathbf{B}_n^T \mathbf{B}_n \boldsymbol{\psi}_n^{(i)} = \lambda_n^{(i)} \boldsymbol{\psi}_n^{(i)} \quad i = 1, 2, \dots, Q+1 \quad (8)$$

By pre-multiplying (8) with \mathbf{B}_n , we obtain the eigenvectors of matrix $\hat{\mathbf{C}}_n$ as follows:

$$\boldsymbol{\varphi}_n^{(i)} = \lambda_n^{(i)-\frac{1}{2}} \mathbf{B}_n \boldsymbol{\psi}_n^{(i)} \quad i = 1, 2, \dots, Q+1 \quad (9)$$

where the term $\lambda_n^{(i)-\frac{1}{2}}$ is to make the resulting eigenvector to be a unit vector. Now we summarize the iterative updating algorithm outlined in this section:

Initialization:

1. Given the first two samples $\mathbf{x}_0, \mathbf{x}_1$, estimate the mean, $\hat{\mathbf{m}}_1$, by (3), and construct the matrix $\mathbf{B}_1 = \left[\sqrt{\alpha_v} (\mathbf{x}_0 - \hat{\mathbf{m}}_1) \quad (\mathbf{x}_1 - \hat{\mathbf{m}}_1) \right]$
2. Based on (8) and (9), we can get the eigenvector, $\boldsymbol{\varphi}_1$, and the eigenvalue, λ_1 .

Iterative updating:

1. Get a new sample \mathbf{x}_n .
2. Estimate the mean, $\hat{\mathbf{m}}_n$, at time n by (3), and get the \mathbf{B}_n matrix from (6).
3. Form the matrix \mathbf{A}_n by (7) and calculate its eigenvectors, $\boldsymbol{\psi}_n$, and eigenvalues, λ_n , by a direct method.
4. Sort the eigenvalues λ_n , and retain Q corresponding eigenvectors.
5. Obtain the eigenvectors, $\boldsymbol{\varphi}_n$, at time n by (9).

We have mentioned three methods of choosing Q . If we use the second and the third methods, Q will increase as more and more training samples arrive till it reaches the intrinsic dimensionality of previous training samples. Due to the approximation in (5), among the Q eigenvectors, typically the first few eigenvectors are more precise than the others. Therefore, in practice if we need N eigenvectors for building an eigenspace, we would keep Q to be a number larger than N .

2.3 An example with synthetic data

In this section, we want to show that our updating algorithm can better model the statistics of a non-stationary random process, than the traditional eigenspace updating algorithms without decay parameters.

We generate 720 samples with a 2-dimensional Gaussian distribution, whose mean is zero and variances in the horizontal and vertical direction are 1 and 7 respectively. If we associate each sample with a time index, we can obtain a random process. For each random variable in this random process, we incrementally rotate it by a certain degree in the 2-dimensional space, and move its mean along the line, $x = y$. The first random variable rotates 0 degree, the last one rotates 90 degree, and all the others rotate in between. In other words, the synthetic data have the following statistics:

$$\begin{aligned} \begin{bmatrix} \mathbf{m}_x[n] \\ \mathbf{m}_y[n] \end{bmatrix} &= \begin{pmatrix} n/15 \\ n/15 \end{pmatrix} \\ \mathbf{C}_{xy}[n] &= \begin{pmatrix} 1 & 24\sin(n\pi/720) \\ 24\sin(n\pi/720) & 1+48\cos^2(n\pi/1440) \end{pmatrix} \end{aligned}$$

One example of this synthetic data is shown in Figure 1. We can see that the cluster of data keeps rotating and moving away from the origin over time.

We use the algorithm introduced in Section 2.1 to update the eigenspace, and show the estimation results of the mean compared to the ground truth in Figure 2. The traditional updating algorithm without decay parameters is also applied on the same data. We can see from Figure 2 that its estimation is much worse than our estimation. When the eigenspace is updated by a new random variable, we calculate the orientation of the first eigenvector with respect to the horizontal coordinate. Ideally the orientation should change from 90 degree to 0 degree according to the time coordinate. As shown in Figure 3, our algorithm can successfully estimate the statistic of the time-varying random process. However, if we apply the traditional method without decay parameters on the same data, the resulting orientation is around 45 degree because it considers all the previous samples equally. The traditional method works well in the beginning as it removes the mean as well. However, it quickly becomes worse because neither the mean nor the variance updating uses the decay parameters.




Figure 1 A synthetic random process.




Figure 2 Estimation of the mean for a random process.




Figure 3 Estimation of the variance for a random process.

3 Choosing the decay parameters

In the proposed eigenspace updating algorithm, we need to specify the decay parameters for both the mean estimation and the variance estimation. In practice, for a recognition system, there is usually a cross validation data available before the testing stage of the system. Thus based on the cross validation data, the optimal decay parameters specific to the application could be obtained by exhaustive search within the valid range, 0 and 1. If there is no cross validation data available, how do we determine decay parameters? We will answer it in this section.

Motivated by the Kalman filter, we model the time-varying mean and variance as autoregressive (AR) random processes

with certain parameters. Now the problem becomes, based on the models and the parameters, how do we find the optimal decay parameters without exhaustive search? We will address the mean estimation and the variance estimation separately.

3.1 Decay for mean estimation

Consider the model in Figure 4, where the sample, x_n , is a scalar and generated by an AR(1) random process plus a white noise, v_n .

$$x_n = m_n + v_n \quad (10)$$

where the observation noise, v_n , has zero mean and the variance of r . The AR(1) process is generated by the following equation:

$$m_n = \rho m_{n-1} + w_n \quad (11)$$

where the white noise, w_n , has zero mean and the variance of q . Based on the above two equations, we can see that x_n and m_n have the same mean. Given x_n , we can estimate its mean at each time instant by estimating the mean of m_n at that time instance, which is denoted as \hat{m}_n .




Figure 4 AR(1) model for observed samples

We consider that the mean is estimated via:

$$\hat{m}_n = \alpha_m \hat{m}_{n-1} + \beta_m x_n \quad (12)$$

where α_m and β_m can take on any values between 0 and 1. Note that (12) is basically the same as (3) without the constraint that $\beta_m = 1 - \alpha_m$. Removing this constraint allows us a more comprehensive study for choosing the decay.

Now the problem becomes to find the optimal α_m and β_m , which can make the \hat{m}_n as close to m_n as possible. Let us derive it by minimizing the estimation error, p_n .

$$p_n = E(m_n - \hat{m}_n)^2$$

By extending the above equation, we have

$$p_n = \rho^2 (1 - \beta_m)^2 p_{n-1} + (1 - \beta_m)^2 q + \beta_m^2 r + h^2 E(\hat{m}_{n-1})^2 - 2h\rho(1 - \beta)E(\hat{m}_{n-1}(m_{n-1} - \hat{m}_{n-1})) \quad (13)$$

where $h = \alpha_m - \rho(1 - \beta_m)$.

In order to have an explicit formulation for p_n , we let h be zero, i.e., α_m and β_m satisfy the following equation.

$$\rho = \frac{\alpha_m}{1 - \beta_m} \quad (14)$$

Thus p_n can be simplified as:

$$p_n = \alpha_m^2 p_{n-1} + (1 - \beta_m)^2 q + \beta_m^2 r \quad (15)$$

When n goes to infinity, p_n converges to:

$$p_{n \rightarrow \infty} = \frac{(1 - \beta_m)^2 q + \beta_m^2 r}{1 - \rho^2 (1 - \beta_m)^2} \quad (16)$$

By taking the derivative of $p_{n \rightarrow \infty}$ with respect to β_m to be 0, we obtain the optimal value for β_m :

$$\beta_m = \frac{r(\rho^2 - 1) - q + \sqrt{(r(1 - \rho^2) + q)^2 + 4\rho^2 qr}}{2r\rho^2} \quad (17)$$

Based on (14), the optimal value for α_m is the following:

$$\alpha_m = \frac{r(\rho^2 + 1) + q - \sqrt{(r(1 - \rho^2) + q)^2 + 4\rho^2 qr}}{2r\rho} \quad (18)$$

We can see that as ρ increases, both α_m and β_m increase. However, α_m increases faster than β_m since there is one more term of ρ in the denominator of β_m . This means that as the random process changes more and more slowly, i.e., ρ gets larger and larger, the previous estimate, \hat{m}_{n-1} , should contribute more and more to the current estimate.

In order to show the effectiveness of our choice of decay parameters, we perform an experiment based on synthetic data. First, we synthesize a set of random processes using the AR(1) process in (10) and (11). By taking these processes as observation samples and assuming we know the parameters in the AR(1) process, i.e., ρ , r and q , we can perform exhaustive search to find the optimal decay parameters which can generate the minimum estimation error. Also by using of (17) and (18), we calculate the decay parameters for our estimate. From Figure 5, we can see that our decay parameters are very close to the optimal decay resulting from exhaustive search.

All the above derivation works in the scalar case. When the sample, \mathbf{X}_n , is a vector, we can obtain the same results if we assume each element of \mathbf{X}_n is independent. In the appendix, we will compare the estimation performance between our estimate and other estimates, such as the Kalman filter and exhaustive search, in terms of estimation errors.




Figure 5 Optimal decay parameters vs. our estimated decay parameters (top: $r/q=3$; bottom: $r/q=1$)

3.2 Decay for variance estimation

Similar to the mean estimation, we study the case when the observation, x_n , is in the scalar form and it can be modeled as:

$$x_n = \sqrt{c_n} v_n \quad (19)$$

Here the white noise, v_n , has zero mean and variance of one. Thus c_n becomes the variance of x_n , and we assume c_n is generated from an AR(1) random process using (11).

Now the problem becomes given an observation sequence, x_n , estimate its variance, c_n , based on the parameter of the AR process, ρ , and the variance of the white noise w_n , q . Similar to our previous section, in our estimate, we use two parameters, α_v and β_v , to combine the information from the previous estimate, \hat{c}_{n-1} , and the current sample, x_n .

$$\hat{c}_n = \alpha_v \hat{c}_{n-1} + \beta_v x_n^2 \quad (20)$$

In order to find the optimal α_v and β_v , we can derive it by minimizing the estimation error. However, it turns out we need to make a very strict constraint in the derivation in order to obtain an explicit result. Experimentally, we found that the estimation performance of the derived result under this strict constraint is not satisfying. Thus we want to solve this problem by an empirical method.

Basically two parameters will affect the selection of decay parameters. The first one ρ , the parameter of the AR process, defines how fast the variance c_n changes over time. Since in most applications the variance does not change too fast, we study the case where ρ is between 0.6 and 1. The second one q , the variance of the white noise w_n , defines how much variability the variance itself will have over time. The larger w_n , the larger range the variance will vibrate. In our experiments, we change these two parameters, and observe the corresponding effect on the optimal decay parameters.

By fixing the above two parameters, we can synthesize the ground truth c_n and observation samples, x_n . The optimal decay parameters in the sense of minimal estimation error are obtained by exhaustive search. Now changing the ρ to be other values, synthesizing data and performing estimation many times, we found the optimal α_v and β_v actually change very little, which means they are basically unaffected by ρ . For a fixed q , by tuning different ρ , we can obtain both the mean and variance of the optimal α_v and β_v . By varying q from 25 to 10000, we can obtain four curves according to the above four statistics. We show the results in Figure 6, where the horizontal axis represents the square root of q . From this figure, we can see that even though the square root of q varies over a large range, the optimal decay parameters do not change significantly. The same experiment can be performed by fixing q and varying ρ . We plot the resulting optimal decay parameters according to different ρ in Figure 7. Thus a good choice of our estimate is to choose the mean of optimal decay parameters as the value of α_v and β_v , where $\alpha_v=0.85$ and $\beta_v=0.13$;

We now conduct an experiment to compare the estimation performance of different approaches. In Figure 8, we show the estimation performance of four approaches according to different ρ . The first one is exhaustive search by constraining that α_v and β_v sum to one. The second one is also exhaustive search, but both the optimal α_v and β_v are searched in the range of 0 and 1. The third one is the sample variance, which is calculated from all the samples with equally weighting. The

last one is our estimate where two fixed decay parameters are used. From this figure, we can see that α_v and β_v summing to one is not a bad constraint since the performance is only slightly worse than the unconstrained case. Also, when ρ is large, which indicates the variance changes slowly, our estimate works much better than the sample variance. But when the variance changes too fast, sample variance turns out to be better than ours because it is harder to estimate the variance for each time instance. Actually in practical applications, the variance tends to change slowly, where ρ is closer to 1. Similar to the mean estimate, when the samples are in the form of vectors, we can obtain the same results by extending our estimate to the vector form.




Figure 6 Optimal decay parameters for the different standard deviation of the white noise.




Figure 7 Optimal decay parameters for different ρ




Figure 8 Results of the variance estimation

4 Face recognition based on updating individual PCA

When applied to face recognition, the proposed eigenspace updating algorithm results in an updating-during-recognition scheme. That is, the eigenspace for each subject is updated by test images while each of them being recognized. There are two reasons for doing this. First, in many applications it is not feasible to capture many training images for each subject containing enough variations for statistical modeling of that subject. Usually only a few images under the normal condition are available for training. Thus, it would be better if more and more images of that subject are used to update its model during the testing stage. Secondly, people change their appearance over time. Even if there are many images available for training, the system may not recognize faces when a subject changes the appearance due to aging, expression, pose, and illumination changes. A recognition system that is able to learn the changing appearance of the subject and adapt to it can achieve better performance.

In using our updating method for face recognition, we assume the test images are from a face sequence and there is continuity between consecutive frames. In the next subsection, we introduce the scheme based on updating a single eigenspace model for each subject. Since this approach may suffer from slow learning, in Section 4.2 we also propose a twin-subspace scheme to alleviating this problem.

4.1 Single subspace updating scheme

Given a set of face images from K subjects for training, each subject has one individual eigenspace trained from his/her own images. When a test image arrives, it is projected into each individual eigenspace and assigned to the one that gives the minimal residue, which is defined by the difference between the test image and its projection in the eigenspace.

Now we need to decide whether to update the eigenspace model of the recognized subject, using the test image. First, by comparing the minimal residue with a pre-defined threshold, we can see whether the current model can represent the test image well. If it does, we do not perform updating since this test image does not bring enough new statistical information for the model. Second, we calculate the confidence measure as the difference between the residue of the second candidate and the residue of the top candidate. Then the confidence measure is compared with another pre-defined threshold. If the confidence measure is larger than the threshold, this test image is utilized to update the assigned eigenspace using our updating method. Basically the larger the confidence measure, the more confidence we have about the current recognition result. Thus as time goes on, the eigenspace will adapt to the most recent statistics of the subject's appearance, and be able to recognize more "new looking" images from that subject.

One risk in this approach is that sometimes the eigenspace model is not updated by the test images with new appearance because of not-high-enough confidence measures, while in the same time the test images keep showing new appearances. In this case, it is likely that the test images will not be correctly recognized because they show different appearance as the current model, which only represents out-of-date appearances. This is the problem with slow learning, i.e., the model does not learn fast enough in order to recognize the test images with new appearances. To alleviate this problem, we introduce the twin-subspace updating scheme in the next subsection.

4.2 Twin-subspace updating scheme

In this scheme, we train two subspaces, the static model and the dynamic model, for each subject. The static model is trained from the original training images of that subject, and the dynamic model is updated from the test images during the testing stage.

When one test image arrives, we calculate its residue to both subspaces for each subject. Then the smaller residue is considered as the distance between the test image and that subject. Eventually the test image is recognized as the subject with the minimal distance. The same as the previous section, we also make the decision of updating based on two thresholds. The first threshold filters out the test image without enough variations with respect to the current model. The second threshold is compared with the confidence measure, which is the differ-

ence between the top candidate and the second candidate in terms of distance. Test images with low confidence in its recognition result are rejected from being used for updating. In this scheme, only the dynamic model is updated by the test images, and the static model will never be changed once it is trained from the original training images. As illustrated in Figure 9, each one of the K subjects has two models, the static one, S_i , and the dynamic one, D_i . The residues $r_{1,1}, r_{2,2}, \dots, r_{k,2}$ are calculated as the distances between the test sample and each model. In this case Subject 2 is the recognition result because it has the minimal distance $r_{2,2}$. Suppose Subject k is the second candidate. The confidence measure $r_{k,2} - r_{2,2}$ is then utilized to decide whether this test sample will be used to update the dynamic model of Subject 2, D_2 .




Figure 9 Testing of the twin-subspace updating scheme

The main reason we propose this updating scheme is to capture different aspects of facial appearance. That is, the static model is used to capture the subject's more intrinsic appearance based on training images, while the dynamic model is used to capture the time-varying statistics of the appearance. The second reason is to deal with the slow learning problem. Because there are two models for each subject during testing, even the test image might not match well with the dynamic model because of the slow learning in the dynamic model, it is still possible that the static model will match with the test image and let the test image update the current dynamic model. Thus the dynamic model can learn the time-varying statistics and benefit future recognition.

In practice, many factors decide whether we should use the single subspace scheme or the twin-subspace scheme. For example, we should use the single subspace scheme when there is small amount of variations in the test images. Also, when we need to deal with the recurrent type of variations, such as pose, expression, illumination, and facial hair variations, we should use the twin-subspace scheme. While for

non-recurrent type of variation, such as aging, the single subspace scheme would be more proper because only the most recent statistics, which are captured by the dynamic model, are useful for future recognition.

5 Experimental results

We conduct experiments on face data sets that contain different variations, such as poses, illuminations and expressions. We will show that for all these variations, our algorithm can achieve much better performance than methods without updating, because we can model variations in a subject’s appearance over time and thus improve the recognition performance. The methods we compare with are the individual PCA method without updating, and traditional eigenspace updating without decay.

In practical applications of face recognition, the human face usually undergoes different kinds of variations, most of which come from the pose, expression, illumination and the combination of them. In order to show the effectiveness of our algorithm in dealing with these variations, experiments are conducted on data sets with these three types of variations.

5.1 Pose data set

We collect a face database with 20 subjects. Each subject has 10 training images. The test images for each subject come from a video sequence, where the subject continuously shows different poses. Both the test and training images are of 32 by 32 grayscale images. There are 210 test images within one sequence for each subject. In Figure 10, we show sample face images from six subjects in this data set, where the images in the same row belong to the same subject. A lot of pose variations can be observed from this data set. Also notice the registration error in some images. This is a very challenging data set for face recognition.

We show the experimental results in Figure 11. The horizontal axis shows the index of the test images, and the vertical axis shows the recognition error rate based on the number of test images so far. We perform experiments on different random orders of the test sequence, and show the average of them in the figure. Three algorithms have been tested on this data set. The first one is the individual PCA method, which works worst because there is no updating involving during the testing stage. The second is our updating method with dynamically estimated decay parameters, which has better performance than the individual PCA method. The third one is our twin-subspace method. It has significant improvement compared to the other two methods since it models the statistics more comprehensively for the changing appearance over time.

In the previous experiment, we do not use eigenvectors in constructing the eigenspace for each training subject. So basically only the mean is used for recognition. Because the number of eigenvectors will affect the recognition performance, we also perform experiments with different numbers of eigenvectors. Table 1 shows the recognition error rates with respect to different numbers of eigenvectors used in constructing the individual eigenspace. Among these three methods, our twin-subspace method has the best performance and the individual PCA works the worst. From this experiment we can see that a proper updating method will work better than a non-updating method in face recognition. Also the twin-subspace method is a promising approach to dealing with large variations, such as poses in this data set.



Figure 10 Sample images of face sequences showing different poses




Figure 11 Experimental results on the pose data set.

Table 1 Recognition error rate with different numbers of eigenvectors

Number of eigenvectors	0	2	4	6
Individual PCA method	27.88%	19.62%	16.43%	14.76%
Our method with dynamic decay	18.57%	10.67%	8.49%	6.87%
Our twin-subspace method	5.77%	4.83%	4.15%	3.98%

5.2 Expression data set

We collect another face database with 30 subjects. Each subject has 5 training images and 70 test images. Each image is of the size of 32 by 32 pixels. The test images for each subject come from a video sequence, where the subject shows varying expressions. The sample images from six subjects are shown in Figure 12. We use the same test scheme as the pose data set. The result is shown in Figure 13. We try both using fixed decay parameters and tuning the decay parameters dynamically according to the changing statistics over time. Here we use the AR(1) random process as the model for face sequences. In all three methods we only update the mean and do not use any eigenvectors.

From this experiment we see that updating methods with decay parameters have better performance than the updating method without decay. Also dynamically tuning decay parameters during the testing stage enhances the modeling of time-varying statistics and hence improves the recognition performance.



Figure 12 Sample images from the expression data set.




Figure 13 Experimental results on the expression data set.

5.3 PIE database

In this experiment, we use a subset of the CMU PIE database [25], which has 7 subjects. Each subject has 24 images, which have the size of 64 by 64 pixels, showing the same expression and pose while under continuous varying illuminations. We use 3 images for training and the remaining 21 images for testing. One eigenvector is used for building eigenspace for each subject. Part of the test images from one subject are shown in Figure 14. The experimental result shown in Figure 15 also indicates that our approach can achieve better performance compared to others.



Figure 14 Images of one subject from PIE database.




Figure 15 Experimental results on the PIE database.

6 Discussions

6.1 Frame-based and video-based recognition

While in Section 5 we treat each test image independently and perform a frame-based recognition, we can also do a video-based recognition in the following two applications scenarios. One is that we recognize the human from the video sequence in an online fashion and do not know when the subject will leave or another subject will come in. In this case, we need to know the recognition results up to the current frame immediately. Many online recognition and verification systems of human faces belong to this case. We call this scenario as online video. The other is that we could offline process the video content, such as indexing of the meeting records or ana-

lyzing surveillance videos, where we are interested in the recognition results after all the frames of one sequence have been captured. This is called offline video. We illustrate these scenarios in Figure 16.

For the online video, by using a face-tracking program [24], we can keep tracking human faces and crop the face region for recognition. With the face tracking, we also know whether the current frame and the previous frames belong to the same subject. An intuitive idea is to use majority voting to see which subject is mostly recognized among all the previous frames. Then a decision will be made on whether using the current frame to update the eigenspace. For the case of the offline video, we can still use the updating based on the majority voting in processing frames one by one. However, as shown in the third row of Figure 16, once a sequence has done with the recognition, we can use all the frames in this sequence to update the eigenspace of the most recognized subject, while this is not feasible in the online video case because it needs to store all the previous frames in one sequence. We have also performed experiments for both the online video and the offline video cases, and the result shows that for the same database, the recognition performance can be significantly improved by using the video-based recognition.




Figure 16 Three application scenarios.

6.2 AR(k) process for decay estimation

In Section 3, we solve the problem of determining decay parameters given the parameters of AR(1) process. However, in face recognition application, given a face sequence, how can we apply the theory in Section 3 on it, i.e., how can we determine whether a face sequence can be approximated by AR(1) or high order AR process; how do we estimate the parameters for an AR process?

The answer to the above questions involves two steps before applying our updating algorithm on the face sequence. One is

model selection. The other is model fitting. Model selection determines k in an AR(k) process. Model fitting estimates the parameters in a specific AR(k) process. There are many existing techniques to solve these two problems in the signal processing literature [26]. For example, model selection can be done by finding k where the synthesized AR(k) is similar to the original signal in terms of statistics. Table 2 shows the corresponding model parameters by assuming five face sequences as AR(k) random processes with k equals to 1 or 2. We found for most face sequences with expression variations, AR(1) is a good statistical model, while for face sequences showing pose variations, some of them might need AR(2) to model them.

Table 2 Model parameters for real face sequences

Face sequences	Expres- sion 1	Expres- sion 2	Expres- sion 3	Pose 1	Pose 2
ρ in AR(1)	0.9998	0.9998	0.9998	0.9975	0.9998
ρ_1 in AR(2)	0.9997	0.9997	0.9502	0.9982	0.9347
ρ_2 in AR(2)	0.1084	-0.0209	-0.0158	0.3283	0.5659

In Section 3, we assume that the observation samples come from the noised version of the AR(1) random process. However, what happen if the samples are actually intrinsic high order AR process, for example, AR(2) random process? In this case, how do we derive the relation between the model parameters, ρ_1, ρ_2, r, q , and the decay parameters, α_m, β_m ?

First of all, we can separate an AR(2) random process into two AR(1)-like random processes as follows:

$$H(Z) = \frac{m(Z)}{w(Z)} = \frac{1}{1 - \rho_1 Z^{-1} - \rho_2 Z^{-2}} = \frac{1}{(1 - p_1 Z^{-1})} \frac{1}{(1 - p_2 Z^{-1})}$$

$$= \frac{u(Z)}{w(Z)} \frac{m(Z)}{u(Z)} = H_1(Z) H_2(Z)$$

This is also illustrated in Figure 17. Now if we compare two AR(1)-like random processes with the AR(1) random process in Figure 4, we can see that u_n plays the similar role as w_n in Figure 4. Thus given an AR(2) random process, we can calculate the variance of u_n and treat it as q . Then by using q, r and p_2 in (17) and (18), we can obtain the decay parameters. Similarly for an AR(k) random process, we can separate it into two parts: an AR(k-1) random process and an AR(1)-like random process, where the former contributes the noise signal for the latter.

One difference in solving decay parameters for AR(1) and AR(2) is that, u_n is not a white noise while w_n being a white noise is one assumption in deriving (17) and (18). However, if w_n is not a white noise, there will be a small no-zeros terms in the right side of the objective function, (13). Thus the solution

provided in (17) and (18) will become sub-optimal for the AR(2) case. We have performed simulation on estimating the mean of an AR(2) random process, and we found the estimation error of our estimate is very close to the one from exhaustive search.




Figure 17 AR(2) model for observed samples

7 Conclusions and future works

In this paper, we introduce a novel approach to updating the eigenspace for non-stationary random processes. Given a new training sample, we iteratively update the eigenspace to manifest the current statistics provided by the new sample. The updated eigenspace is based more on the recent samples and less on the older samples. Extensive study has been performed on how to choose the decay parameters for our updating method. We show the effectiveness of our algorithm using both synthetic data and practical applications for face recognition. The experimental results indicate that the random processes in many practical applications are essentially non-stationary, which results in significantly improved performance by our updating method compared to other methods in the literature.

As a modeling tool, our eigenspace updating method can also be applied to other applications, for example, the detection of signal changing [22][23] and video coding. We have already applied it to the shot boundary detection [27] and the detection of facial expression changes. It is able to model the most recent statistics over time, and thus any change in signals can be detected from the residue between the new signal and the eigenspace.

In face recognition, many approaches have been proposed to deal with different variations. While each approach works well for the specific variation being studied, performance degrades rapidly when other variations are present. In practice, the test images usually undergo the mixture of variations, such as expressions, poses and illuminations. Trying to use a static model to cover all these variations is difficult. Using the pro-

posed updating method is one solution. Instead of trying to model all variations at once, we try to dynamically model only the most recent variations.

There are many interesting directions to be explored further. For example, in our updating method, the eigenvector expansion is truncated as in (5), which results in an approximate representation for the covariance matrix. Can we have a better estimation for the covariance by adding a diagonal term that accounts for the discarded residue? Also, as the eigenspace only provides a subspace representation for a data set, it lacks a probabilistic measure for the samples in the data set. Can we borrow the idea of probabilistic PCA [28] and update the eigenspace in a probabilistic framework, i.e., the resulting eigenspace can have a probabilistic interpretation for each sample? Further, while applying updating methods for classification, a good scheme to decide when to perform updating is very critical and requires more study. Finally, how to take advantage of the temporal information and perform the video-based recognition is also an interesting topic worth further study. In addition to updating the inner-product matrix or the covariance matrix, we can also extend our non-stationary updating algorithms, in particular the choice of decay parameters, to SVD-based eigenspace updating methods.

Appendix

In the Kalman filter, there are five iterative steps to perform the estimation. They can be described by the following five equations:

$$\hat{m}_n^- = \rho \hat{m}_{n-1} \quad (21)$$

$$p_n^- = \rho p_{n-1}^- + q \quad (22)$$

$$\hat{m}_n = \hat{m}_n^- + k_n (x_n - \hat{m}_n^-) \quad (23)$$

$$k_n = \frac{p_n^-}{p_n^- + r} \quad (24)$$

$$p_n = (1 - k_n) p_n^-$$

By extending (23) with (21), we get:

$$\hat{m}_n = \rho(1 - k_n) \hat{m}_{n-1} + k_n x_n \quad (25)$$

If we compare the above equation with our estimate (12), we can find that k_n corresponds to β_m , and $\rho(1 - k_n)$ corresponds to α_m in our estimate. By combining (22) and (24), the convergent formula of p_n when n goes infinity can be found. Since k_n only depends on p_{n-1}^- , q , and r , eventually we can obtain the convergent formulation for k_n , which is exactly the same as β_m in (17). Hence $\rho(1 - k_n)$ has the same formulation as α_m in (18). From this, we can see that our estimate turns out to be the convergent form of the Kalman filter.

We conduct the following experiment to show the estimation performance. Given different random processes synthesized by tuning different ρ in the AR(1) process, we can estimate the model parameters first. Then we can utilize the model parameters to derive the decay parameters for our updating method. For the AR(1) random process, we estimate model parameters as following:

$$R_{xx}(k) = E(x_n x_{n+k}) = E((m_n + v_n)(m_{n+k} + v_{n+k})) = E(m_n m_{n+k}) + E(v_n v_{n+k}) = R_{mm}(k) + R_{vv}(k)$$

$$R_{mm}(k) = \frac{q}{1 - \rho^2} \rho^{|k|}$$

Basically by calculating $R_{xx}(0)$, $R_{xx}(1)$, up to $R_{xx}(k)$, we can estimate q , r and ρ based on the above two equations. Since there are estimation errors in estimating these model parameters, we are interested in how our estimate performs based on these estimated model parameters.

Then these model parameters are fed into both the Kalman filter and our estimate to estimate the mean of the random process. Since we have the ground truth of the mean, we can also perform exhaustive search for the decay parameters as well. In Figure 18, we show the estimate errors of three different estimates and also the variance between the given observations and the ground truth. We can see that for different choices of ρ , our estimate performs better than the Kalman filter, especially in the region with large ρ .




Figure 18 Mean estimation based on estimated model parameters.

Acknowledgments

This work was supported by the U.S. Department of Commerce, National Institute of standards and Technology Program, Cooperative agreement Number 70NANB8H4076. The authors would like to thank the anonymous reviewers for insightful comments. Thanks to Prof. B.V.K. Vijaya Kumar for fruitful discussion. Thanks also go to a number of volunteers in the Electrical and Computer Engineering Department of Carnegie Mellon University for helping us in data collection.

References:

- [1] Y.T.Chien, K.S. Fu, On the generalized Karhunen-Loeve expansion, IEEE Transaction on Information Theory, 13 (3) (1967) 518-520.
- [2] A. Habibi P.A. Wintz, Image coding by linear transform and block quantization techniques, IEEE Transaction on Communication and Technology. Vol. CoOM-19, (1971) 948-956.
- [3] R. J. Wong, P. A. Wintz, Information extraction, SNR improvement, and data compression in multispectral imagery, IEEE Transaction on Communication and Technology. Vol. CoOM-21, (1973) 1123-1131.
- [4] E. Oja, Subspace methods of pattern recognition. Letchworth, Hertfordshire, England. New York: Wiley, 1983.
- [5] L. Sirovich, M. Kirby, Low dimensional procedure for the characterization of human faces, Journal of Optical Society of America, 4 (3) (1987) 519-524.
- [6] M. Turk, A. Pentland, Eigenfaces for Recognition. Journal of Cognitive Neuroscience. 3 (1) (1991) 71-86.
- [7] R.A. Horn, C.R. Johnson, Matrix Analysis. Cambridge University Press 1985.
- [8] H. Murakami, B.V.K.V. Kumar, Efficient calculation of primary images from a set of images. IEEE Transactions on Pattern Analysis and Machine Intelligence, 4 (5) (1982) 511-515.
- [9] A. Levy and M. Lindenbaum, Sequential Karhunen-Loeve Basis Extraction and its applications to Images, IEEE Transactions on Image processing, 9 (8) (2000) 1371-1374.
- [10] Peter M. Hall, David Marshall, Ralph R. Martin, Incremental Eigenanalysis for Classification. Department of Computer Science, University of Wales Cardiff, UK, Research Report Serious No: 98001, May 1998.
- [11] S. Chandrasekaran, B.S. Manjunath, Y.F. Wang, J. Winkeler, H. Zhang, An eigenspace update algorithm for image analysis, Graphical Models and Image Processing, Academic Press, 59 (5) (1997) 321-332.
- [12] Peter Hall, David Marshall, and Ralph Martin, Merging and Splitting Eigenspace Models, IEEE Transactions on Pattern Analysis and Machine Intelligence, 22 (9), (2000) 1042-1049.

- [13] R. Chellappa, C.L. Wilson, S. Sirohey, Human and machine recognition of faces: a survey. *Proceedings of the IEEE*, 83 (5) (1995) 705-741.
- [14] T. Sim, T. Kanade, Combining Models and Exemplars for Face Recognition: An Illuminating Example, *Proceedings of the CVPR 2001 Workshop on Models versus Exemplars in Computer Vision*, December 2001.
- [15] P.N. Belhumeur, J.P. Hespanha and D.J. Kriegman, Eigenfaces vs Fisherfaces: recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20 (7) (1997) 711–720.
- [16] X. Liu, T. Chen, B.V.K. Vijaya Kumar, Face Authentication for Multiple Subjects Using Eigenflow. Accepted for publication in *Pattern Recognition*, special issue on Biometric, November 2001.
- [17] T. Kurita, M. Tanaka, K. Hotta, H. Shimai, T. Mishima, Efficient face detection from news images by adaptive estimation of prior probabilities and ising search, *Proceedings of 15th International Conference on Pattern Recognition*, 2000, Volume: 2, (2000) 917-920.
- [18] G.J. Edwards, C.J. Taylor, T.F. Cootes, Learning to identify and track faces in image sequences, *Proceedings of Third IEEE International Conference on Automatic Face and Gesture Recognition*, 1998, (1998) 260-265.
- [19] Hsi-Jung Wu, D. Ponceleon, K. Wang, J. Normile, Tracking subspace representations of face images, *Proceeding of 1994 IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1994, Volume:5, (1994) 389-392.
- [20] Juyang Weng, C.H. Evans, Wey-Shiuan Hwang, An incremental learning method for face recognition under continuous video stream, *Proceedings of Fourth IEEE International Conference on Automatic Face and Gesture Recognition*, 2000, (2000) 251-256.
- [21] R. E. Kalman, A New Approach to Linear Filtering and Prediction Problems. *Transaction of the ASME--Journal of Basic Engineering*, March 1960, 35-45.
- [22] T. Otsuka, J. Ohya, Recognizing abruptly changing facial expressions from time-sequential face images, *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition*, 1998, (1998) 808-813.
- [23] H. J. Zhang, A. Kankanhalli, S. Smoliar, Automatic partitioning of full-motion video. *Multimedia systems*. 1 (1) (1993) 10-28.
- [24] F. J. Huang, T. Chen, Tracking of Multiple Faces for Human-Computer Interfaces and Virtual Environments. *Proceeding of IEEE International Conference on Multimedia and Expo*, New York, July 2000.
- [25] T. Sim, S. Baker, M. Bsat, The CMU Pose, Illumination, and Expression (PIE) Database of Human Faces, tech. report CMU-RI-TR-01-02, Robotics Institute, Carnegie Mellon University, January 2001.
- [26] James D. Hamilton, *Time Series Analysis*, Princeton University Press, 1994.
- [27] X. Liu, T. Chen, Shot Boundary Detection Using Temporal Statistics Modeling, *Proceedings of 2002 IEEE International Conference on Acoustics, Speech, and Signal Processing*, November 2001.
- [28] M. E. Tipping and C. M. Bishop, Mixtures of Probabilistic Principal Component Analyzers, *Neural Computation*, 11 (2) (1999) 443-482.

About the Author—XIAOMING LIU received the B.E. degree from Beijing Information Technology Institute, Beijing, China, and the M.E. degree from Zhejiang University, Hangzhou, China, in 1997 and 2000 respectively, both in Computer Science. Since January 2000, he has been working towards his Ph.D. degree in Electrical and Computer Engineering at Carnegie Mellon University, Pittsburgh, PA. His research interests are in the areas of image processing, pattern recognition and computer vision. He is a student member of the IEEE.

About the Author—TSUHAN CHEN received the Ph.D. degree in electrical engineering from the California Institute of Technology, Pasadena, California, in 1993. From August 1993 to October 1997, he worked in the Visual Communications Research Department, AT&T Bell Laboratories, Holmdel, New Jersey, and later at AT&T Labs-Research, Red Bank, New Jersey, as a senior technical staff member and then a principle technical staff member. Since October 1997, he has been with the Electrical and Computer Engineering Department, Carnegie Mellon University, as Professor. His research interests include multimedia signal processing and communication, audio-visual interaction, video coding and multimedia standards.

Dr. Chen co-founded and chaired the Multimedia Signal Processing Technical Committee of the IEEE Signal Processing Society. He serves as Associate Editor for *IEEE Trans. on Image Processing* and *IEEE Trans. on Multimedia*. He serves on the Steering Committee of *IEEE Trans. on Multimedia* and the Editorial Board of *IEEE Signal Processing Magazine*. He was a technical co-chair of the First IEEE Conference on Multimedia and Expo in 2000. He has recently been appointed as the Editor-in-Chief for *IEEE Transactions on Multimedia* for 2002-2004. Dr. Chen has published tens of technical papers and holds thirteen U.S. patents. Dr. Chen is a recipient of the National Science Foundation CAREER Award.

About the Author—SUSAN M. THORNTON received the B.S.E.E. degree in 1993 from the University of Notre Dame, Notre Dame, IN, and the M.Sc. and the Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1995 and 1999, respectively.

She is currently a Research Scientist with Sonic Foundry, Inc., Media Systems Division, Wexford, PA. Sonic Foundry Media Systems provides customized development of automated rich-media application systems that allow customers to deploy, manage, and distribute video content over corporate networks or the Web. Her research interests are in the areas of signal, image, and video processing, statistical signal processing, and pattern recognition.

Dr. Thornton was the recipient of a National Science Foundation Fellowship Award (1993-1996), as well as the winner of a prestigious Amelia Earhart Fellowship Award (1998-1999) from Zonta International. She is a member of IEEE and SPIE.