A PDA-based Face Recognition System

Jie Yang, Xilin Chen, William Kunz
School of Computer Science, Carnegie Mellon University
Pittsburgh, PA 15213, USA
{yang+, xlchen, kunz+}@cs.cmu.edu

Abstract

In this paper, we present a PDA-based face recognition system as well as some of the associated challenges of developing a PDA-based face recognition system. We describe a prototype system built from an off the shelf PDA, and introduce algorithms for image preprocessing to enhance the quality of the image by sharpening focus, and normalizing both lighting condition and head rotation. We use a unified LDA/PCA algorithm for face recognition. The algorithm maximizes the LDA criterion directly without a separate PCA step, which eliminates the possibility of losing discriminative information due to a separate PCA step. We demonstrate effectiveness of these algorithms and the feasibility of this system by experimental results. This system has many applications including information retrieval and law enforcement.

1. Introduction

The palm-size personal digital assistant (PDA) is becoming ubiquitous among professionals. With continuous increase in computing power, memory, and accessories, a PDA can now provide more potential applications in our daily lives. For example, we are now able to turn a PDA into an amazing imaging productivity tool by simply plugging in a camera card. A straightforward application of this could be capturing images and adding a personal touch by recording voice with it, and then attach the images to e-mail to share with friends and colleagues. Another application might be to help people remember names and details. Many people struggle to remember names of people they have met before. A PDA with a camera and face recognition software can be used in a subtle and transparent way that will allow someone to use his/her PDA to recall a persons name and any relevant information about someone, thus forever doing away with the embarrassing, "Excuse me, could you tell me your name again?"

Automatic face recognition has been an active research area in the last two decades. The progress in this area can be found in review papers [3, 13] and the proceedings of the last four international conferences on Face and Gesture Recognition. In general, face recognition methods fall into two different categories: holistic template matching and geometric local feature-based schemes. Each type of system has certain advantages and disadvantages, even though both types of systems have some successfully applications to face recognition tasks. Therefore, we should carefully select different approaches based on the specific requirements of a given task. Taking into consideration various requirements, limitations, and applications of a PDA-based face recognition task, we have decided to us a holistic template based matching scheme. For example, we may have different number training samples, from one to many, for different people. We have developed efficient algorithms to deal with such problems in our previous research [18].

The previous work in system development has been mostly focused on using a fixed camera or a video surveillance camera. However, some recent publications have focused on face recognition systems on a wearable platform [4, 9, 11]. The system was built on the "WearCam" platform, which permits both hands to be free while the user wears at least one camera on the body in some manner and receives feedback from the head-mounted display. This kind of system requires special hardware. In addition, the image quality is generally poor since it requires a wide-angle lens (or fish eye lens), and it is difficult to position the camera without using hands. Faces are generally only recognizable under bright lighting conditions and from less than 10 feet away [4].

In this paper, we present a PDA-based face recognition system. The system takes advantage of human feedback to aid in processing. We discuss some of the associated challenges of developing a PDA-based face recognition system. We describe a prototype system built from an off the shelf PDA, and we introduce algorithms for image preprocessing to enhance the quality of the image by sharpening focus, and normalizing both lighting condition and head rotation. We use a unified LDA (Linear Discriminant Analysis)/PCA (Principal Components Analysis) algorithm for face recognition. The algorithm maximizes the LDA criterion directly without a separate PCA step, which eliminates the possibility of losing discriminative information due to a separate PCA step. We demonstrate effectiveness of these algorithms and the feasibility of this system through experiments.

2. Face recognition using a PDA

Unfortunately, although most automatic face recognition systems used today work very well under constrained conditions, they may fail under unconstrained conditions that can vary vastly. Our experiments indicate that face localization and alignment are crucial for a face recognition system. The task is very difficult for a system under highly variable conditions, but not for a human. With the flexibility of a PDA, a user can capture an image from the best position and help select facial features appropriately which can be used by the system to normalize faces. The system can then focus on the recognition task. As such, we have shifted some duties from computers to human users: the systems no longer attempt to do everything themselves, but instead, take advantages of human capabilities, which in turn is then capable of assisting the user by recognizing the face. The underlying assumptions to this approach and a possible solution come from incorporating the human operative into the process, which concept is different from both a fully automatic face recognition system and a "WearCam" type wearable face recognition system.

2.1 Challenges

A PDA-based face recognition system not only shares some common challenges that the previous face recognition systems have met, but also has some particular problems. Many challenges of a face recognition system come from substantial variations in appearance that faces undergo with changing illumination, orientation, scale, and facial expressions. Although a PDA with the attached camera could capture high quality pictures of a face under a normal condition, it may capture a poor quality image when a person is far away or the lighting conditions are poor. Problems with focus were the most common challenges in our experiments. Furthermore, depending on the rotations of the head and the position of a person relative to the lighting sources, the illumination of the face changes dramatically. This means we can observe the full range of shade variations even though the overall lighting conditions in the room remain constant over the time. Moreover, people constantly move their heads and change their facial expressions. Given the dynamic nature of the environment where we use a PDA-based face recognition system, the image that we capture can be significantly different from the image in the database. Last but not least, a PDA has limited resources including:

- **Limited computing power:** Computing power of PDAs is much less powerful than that of desktop or laptop computers. Currently, the most powerful CPU for the Palm OS based system is 66MHz DragonBall CPU, while the most powerful CPU for WIN CE based system is a 206MHz StrongARM CPU. In addition, all CPUs of current PDAs don't have hardware supporting float point computation in order to reduce power consumption and chip size. Float point computations on a PDA are implemented by software through a float emulation library. For a floating task, a 200MHz integer CPU may be even slower than a 50MHz CPU with float point computing component.
- **Limited memory:** For a variety of reasons, all PDAs are configured with limited memory. A typical Palm OS based system is with 8-16MB memory, and a WIN CE based system is with 16-64MB memory. This space is for storage and application programs. The memory size is crucial for a computer vision application because some algorithms require a huge space.
- **Limited size of display:** All PDAs use a small display. Most of the Palm OS PDAs use 160 x 160 displayer (SONY is the only exception, which uses the 320x320 or 320x480 displayer), and all of the Palm PC and Pocket PC use 320x240 display, which is the quarter of standard VGA. A small sized display imposes more constraints in user interface design.

We will address these challenges in both algorithm development and system implementation.

2.2 Approaches

We shall now address the challenges in a PDA-based face recognition system in both image preprocessing and recognition algorithms.

2.2.1 Image peprocessing

Almost all the cameras attached to a PDA are very simple and therefore cannot guarantee high quality images. Blurry images are one of the most common problems. The problem comes mainly from the structure of the simple camera, which does not have auto focus capabilities. To compensate for such a disadvantage, it uses a small size diaphragm to approximate a pinhole camera and sets its focus at the most commonly used distance. However, the camera would be out of focus when the object is far beyond the focus range, because it is impossible to implement a true pinhole camera. In addition, a small size diaphragm requires longer exposure time during which unexpected motion can happen. All these factors can cause a blurred image.

A blurred image can be sharpened using image processing algorithms. In fact, a blurred image can be considered the result of a clear image passing through a linear shift-invariant system; i.e.,

$$g(x,y) = f(x,y) \otimes h(x,y), \tag{1}$$

where $f(x,y)$ is the clear image, $h(x,y)$ is the kernel function, $\otimes$ is convolution operator, and $g(x,y)$ is the blurred image.

We can obtain a clear image from the blurred image using a deconvolution operation, if $h(x,y)$ is known. A practical method is to use a Gaussian model to approximate $h(x,y)$. The variance of the model can be estimated from training samples from the camera. Considering the limited resources on a PDA, we used an iteration method [2] to implement the deconvolution. The problem then becomes how to solve the following iteration equations when g and h are known,

$$\begin{cases} f_0(x,y) = T\left[g(x,y)\right] \\ f_i(x,y) = \lambda g(x,y) + q(x,y) \otimes T\left[f_{i-1}(x,y)\right] \quad (i>0) \end{cases}, \tag{2}$$

where T is some constrained operators, and

$$q(x,y) = \delta(x,y) - \lambda h(x,y), \ \delta(x,y) = \begin{cases} 1 & x = y \\ 0 & x \neq y \end{cases}.$$

We use normalization methods to deal with variations in scale, head rotations, and lighting. In a PDA-based application, a user has a certain freedom to find an optimal position to capture an image. For example, a user can try to capture a front face instead of a side face. However, we cannot always capture every person's head always up straight. Furthermore, there is no guarantee that a user can hold the PDA perfectly horizontal. This will result in a head rotation on the image plane as shown in Figure 1(a). Such a rotation will decrease the performance of a face recognition system. The alignment can be made from facial features. Figure 1(b) shows the result of normalized image in scale and rotation from the image of Figure 1(a) using the locations of the irises.

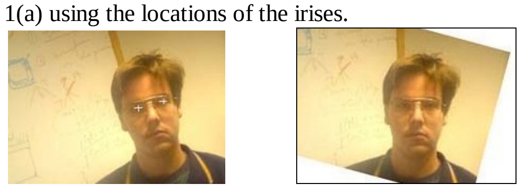

(a) (b)

Figure 1 An example of rotated head and its recovery

In order to handle changes in light conditions, we can use an adaptive method to modify intensity of a face image. The idea is to select a face image with a good lighting condition and use its intensity histogram as the standard histogram. For every other input image, we normalize its histogram to match the standard histogram. After normalization, all input images will have approximately the same mean and variance for their intensity histograms. This method can be extended to sub-images of a face when multiple lighting sources exist, e.g., when side-illumination occurs in the image.

2.2.2 Confidence masure

Confidence measure is another important factor for a face recognition system. Unfortunately, there is no formal way to judge the recognition results. For a face recognition system, the image quality is a primary factor that influences the result. We evaluate the recognition result using a quality estimation model. The model considers several aspects of an image, such as entropy, noise, mean value, and contrast. The entropy measures the complexity of an image, but complexity alone is not a sufficient measurement. For example, a random synthesized image is complex but meaningless to people. Such a case can be detected by measuring noise in an image.

We can use the second order derivative to measure noise in the image based on the assumption that the image is piecewise in most areas and the edge is notable. We can further measure the dynamic scope of the image intensity from the mean and the contrast of the image. We then use the following criterion to measure the quality of an image.

$$Quality = \begin{cases} 10\log\dfrac{E^2}{\mu_n} & \text{if } \mu < \mu_{\min} \text{ or } \mu > \mu_{\max} \text{ or } Con < Th_{Con} \\ 20\log\dfrac{E^2}{\mu_n} & \text{otherwise} \end{cases}, \qquad (3)$$

where E is the entropy, μ_n is the noise mean, μ is the mean of the image and Con is the contrast, the $\mu_{\min}$, $\mu_{\max}$ and Th_{Con} are the thresholds which can be set experimentally.

2.2.3 Recognition agorithm

As mentioned in the previous section, a holistic template matching-based method is more suitable for applications performed by a PDA-based face recognition system. Among various approaches, techniques based on PCA [7], popularly called *eigenfaces* [12, 15], have played a fundamental role in dimensionality reduction and demonstrated excellent performance. Many different methods have been used for face recognition, such as the Euclidean distance [15], Bayesian [10] and LDA [1, 5, 8, 14, 19]. Unlike the PCA that encodes information in an orthogonal linear space, the LDA encodes discriminatory information in a linear separable space of which bases are not necessarily orthogonal. Researchers have demonstrated that the LDA based algorithms outperform the PCA algorithm for many different tasks [1, 19]. However, the standard LDA algorithm has difficulty processing high dimensional image data. PCA is often used for projecting an image into a lower dimensional space or so-called face space, and then LDA is performed to maximize the discriminatory power. In those approaches, PCA plays a role of dimensionality reduction and form a PCA subspace. Relevant information might be lost due to an inappropriate choice of dimensionality in the PCA step [20]. However, LDA can be used not only for classification, but also for dimensionality reduction. We employ a unified LDA/PCA algorithm [18] for PDA-based face recognition. The algorithm maximizes the LDA criterion directly without a separate PCA step. This eliminates the possibility of losing discriminative information due to a separate PCA step. The algorithm is equivalent to the PCA approach in the special case where total scatter matrix is used in the numerator of Fisher's criterion, and each inner-class variance for each person is zero. This algorithm provides a flexible way for handling a variable number of training samples for each person.

The basic idea of LDA is to find a linear transformation such that feature clusters are most separable after the transformation. This can be achieved through scatter matrix analysis [6]. For an *M*-class problem, the class separability can be measured by established criterion. A commonly used one is the ratio of the determinant of the between-class scatter matrix S_b of the projected samples to the within-class scatter matrix S_w of the projected samples.

$$J(A) = \arg\max_{A} \frac{\left|AS_bA^T\right|}{\left|AS_wA^T\right|}, \qquad (4)$$

where A is an $m \times n$ matrix with ($m \leq n$). The most frequently used LDA algorithm in practice is based on simultaneous diagonalization [6]. The basic idea of the algorithm is to find a matrix A that can simultaneously diagonalize both S_w and S_b. The simultaneous diagonalization algorithm involves matrix inversion. To our knowledge, most algorithms require that the within-class scatter matrix be S_w non-singular, because the algorithms diagonalize S_w first. Such a procedure breaks down when the within-class scatter matrix S_w becomes singular. This can happen when the number of training samples is smaller than the dimension of the sample vector. This is the case for most face recognition tasks. For example, a small size of image of 64x64 turns into a 4096-dimensional vector when vectorized. The solution to this problem is to perform two projections [1, 5, 14, 20]:

- Perform PCA to project the *n*-dimensional image space onto a lower dimensional sub-space;
- Perform discriminant projection using LDA.

The null space of S_w may contain useful information if the projection of S_b is not zero in that direction, but the null space of S_b can be safely discarded. Almost all the LDA algorithms diagonalize S_w first. This requires S_w to be non-singular because the procedure involves inversion. However, the simultaneous diagonalization algorithm can start from either matrix of two symmetric matrices. In other words, we can diagonalize S_b first instead of S_w. If we begin diagonalization from S_b, we need to keep S_b non-singular. It will not lose any useful information if we remove the null space from S_b. This leads a direct LDA algorithm to obtain an exact LDA solution without a separate dimensionality reduction step. In fact, the algorithm is a "unified" algorithm with some previous face recognition algorithms if we modify the Fisher's criterion. In fact, there are other variants of Fisher's criterion [4]. For example, if we use $S_t = S_b + S_w$ to replace S_b in Equation (4), the first step of the new algorithm performs PCA function exactly as other "PCA+ LDA" algorithms did in a separated procedure. In an extreme case where each class has only one sample, the new algorithm will get the same result as a PCA algorithm. Therefore, the new algorithm is a "unified PCA + LDA" algorithm. Interested readers can obtained the detailed algorithm from [18].

3. A prototype system

The current PDA market is shared by Palm OS based PDA and WIN CE based PDA. Some PDAs can attach a camera. For example, Palm M100/M105 can use the KODAK PALMPIX camera. SONY CLIE NR 70 even has an embedded camera. Among WIN CE based systems, HP JORNADA 54X and 56X serials can attach a camera, and CASIO PDAs can use its JK-710DC Digital camera card, which can be used by almost all CASIO PDAs. Having considered computing power, memory size, and availability of camera as an accessory, we have chosen HP Jornada Pocket PC as a platform for developing a prototype system. The HP JORNADA 568 is configured with a 206MHz StrongARM CPU and 64MB RAM which is shared by storage and program space. It provides a CF type I slot which can attach a camera as an imaging device. It provide a 320x240 LCD with 65536 colors. Figure 2 shows an HP Jornada Pocket PC with an HP pocket camera.

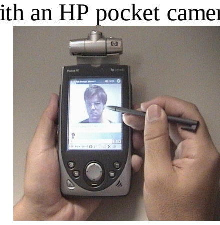

Figure 2 HP Jornada pocket PC with HP pocket camera

The system consists of two parts: an interface and a face recognition module. The interface is to let a user to select a face and facial features for face alignment. The face recognition module contains image preprocessing and face recognition algorithms discussed in the previous section. The hairstyle can play an important role in a global matching algorithm such as PCA and LDA. In order to eliminate influence of the hairstyle, the system removes hair by a facial area mask based on locations of irises.

The system runs on Windows CE environment. We developed the system using Microsoft embedded development tools on a desktop computer and then download to a pocket PC. In order to make the system run at real-time, we have employed many engineering solutions. For example, we have tried to avoid float point computation multiplication, and memory copy. We can convert a float point computation into an integer computation by normalizing the operands into an integer under certain operator and ensuring no overflow for both operand and result. In some cases, we can convert multiplication into addition or table lookup. We can reduce memory copy by directly access a file. In our application, files are stored in memory, and thus they can be accessed very fast. Instead of copying a file into the application, we can access it from file system directly. To implement this, we just need to maintain a file pointer. In this way, we can avoid to copy a database or image file.

Simple usage involves a user capturing an image of the subject. Because of the non-intrusive nature of a PDA, this can be transparently to the subject. The user then circles the face of the person to be recognized. This brings up a zoomed in version of the person's face where the user then draws a line between the subjects pupils. Having the user identify the face and irises has many advantages. A human has near 100% accuracy in recognizing contours of the face and location of the eyes, tasks in which even the best software heuristics are still far from perfect. In this current prototype system, the face is decomposed using the method explained earlier and compared to the facial database, which is stored on the PDA. Depending on the number of dimensions kept, the PDA is capable of storing thousands of faces. Future incarnations of this project might involve reading the dimensional data from a centralized database via a wireless connection. This would permit a nearly infinite number of faces in the database.

4. Experimental results

We have performed many different experiments to study image processing and face recognition algorithms and demonstrate feasibility of the hardware. We use public available databases for algorithm evaluations to make it comparable with other research. We use images taken by the hardware of our prototype system to demonstrate feasibility of the system, and we will make those images available to interested researchers.

(a) (b) (c)

Figure 3 An example of image enhancement

4.1 Image processing

We have tested the effectiveness of the image processing algorithms. Figure 3(a) is an example of a blurred image taken by the HP pocket camera. We applied Equation (2) to enhance the quality of the image. Figure 3(b) and (c) show the results of the enhanced image after 30 passes with $\sigma = 0.5$ and $\sigma = 0.3$ separately. The improvements are obvious.

4.2 Recognition

We have evaluated recognition accuracy of the recognition algorithm using two different databases: ORL (Olivetti-Oracle Research Lab) database (ORL Website) and Yale face database (Yale Face Database). The ORL dataset consists of 400 frontal faces: 10 tightly-cropped images of 40 individuals with variations in pose, illumination, facial expression (open/closed eyes, smiling/not smiling) and accessories (glasses/no glasses). We reduced the number of dimension to 39 and tested recognition accuracy vs. number of training samples. We varied the number of training samples from 1 to 9. n ($n = 1,2,\ldots,9$) randomly-selected images for each individual in the dataset were placed in the training set, and the remaining images were used for testing. Ten runs for each of n samples were performed with different, random partitions between training and testing images. The best recognition rate for the new algorithm is about 95%, which is compatible with the best result obtained by other researchers on the same test set using different algorithms.

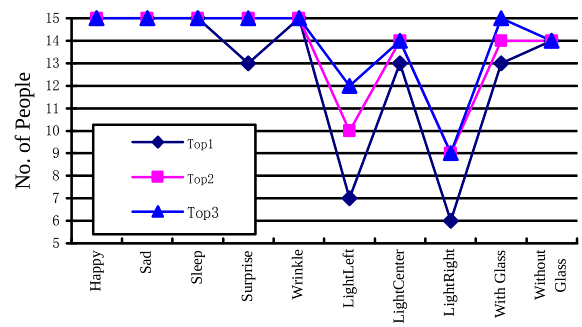

Figure 4 Experiments of the Yale face database

The new LDA algorithm has the capability of handling a small size of training data. We have proved that the algorithm is equivalent to the eigenface (PCA) approach in the special case where each person has only one sample in the training set. The feasibility of the new algorithm has been demonstrated by experimental results (Yang 2000). We used the Yale face database to test the robustness of the system against variation. The Yale face database contains images from 15 different people. Each person has 11 pictures with variations in lighting condition, emotion, and with/without glasses. We used the normal face as the training sample for each person and the rest of 10 pictures as the testing samples. Figure 4 depicts the results. The results indicate that the system is robust against variations in displayed emotion. However, it performs poorly for side lighting source changes. This is an inherent problem of a holistic template matching-based method, because the algorithm encodes the intensity of the whole face.

4.3 Pocket camera

We have tested the face recognition system using the images taken from the HP pocket camera used for the prototype system. Considering the applications of the system, we performed two specific tests: effective distance of the camera, and recognition accuracy. To test the effective distance of the system, we took pictures at different distances and then performed the recognition. We took images at 10, 15, 20, 25, and 30 feet both indoors and outdoors. The system performed very well for images within 20 feet, and worked fine for images taken at 25 feet. But the performance decreased significantly after 25 feet. This can be illustrated numerically using Equation (3). Table 1 shows the evaluation results of a set of images at different distances.

Table 1 The evaluation results of a set of images

	Mean	Contrast	Noise	Entropy	Quality
10 feet	120	66	7.5499	4.5719	20.3664
15 feet	110	68	8.1118	4.5048	18.3396
20 feet	122	52	7.8167	4.2939	17.1624
25 feet	128	52	7.4440	4.2104	17.3541

To test recognition accuracy, we collected data from our lab using this system. We collected a total 116 pictures for 24 people under various lighting conditions and in different poses. We randomly selected one image from each person for training and the rest for testing. The recognition accuracy is 81.25%, if the system considers only top one choice. In a PDA-based based face recognition task, a user can help to make decision from an *N*-best list provided by the system.

Figure 5 Chandra Levy pictures found from Internet

Figure 6 The composed pictures by Washington police

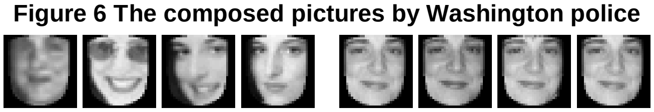

(a) The train samples (b) The test images

Figure 7 The masked faces

We have also tested flexibility of the system. A good example is the application of the recognition system to the missing federal intern Chandra Levy's case, which once attracted much public attention. In order to investigate the case, Washington police composed pictures showing what she might look like. We used the images found in the Internet as training images (Figure 5) and the images published by the police (Figure 6) as the test images. Figure 7 shows the masked faces. It is clear that the mask has effectively removed any influences of the hairstyle. We added the training set to our database and then tested the system. As the result, all test images have been recognized successfully by the system. This example suggests another application of the proposed system for law enforcement.

5. Conclusions

Technology development in palm-size PDA's has made it possible to take face recognition systems in a new direction. A PDA-based face recognition system can take advantage of having a human-in-the-loop and enhances human capability of recognizing other people. This paper has presented techniques for developing a PDA-based face recognition system and demonstrated the feasibility of such a system through experiments. We are further improving the robustness of the system and working to apply this technology to new fields. The applications of a PDA based face recognition system include extending human memory and law enforcement.

6. Acknowledgments

The authors would like to our colleagues in the Interactive Systems Labs for their inspiring discussions, support and helps on data collection. This Research is partially supported by the National Science Foundation under Grant No. IIS-0121560, and a gift from Microsoft Research.

References

[1] P. Belhumeur, J. Hespanha, and D. Kriegman. Eigenfaces vs. Fisherfaces: recognition using class specific linear projection. IEEE PAMI, 19(7):711-720, 1997.

[2] P. H. Van Cittert. Zum einfluss der splatbreite auf die intensitat-swerteilung in spektrallinien II. *Z . für* Physik, 69: 298-308, 1931.

[3] R. Chellappa, C. Wilson, and S. Sirohey. Human and machine recognition of faces: A survey. Proc. of the IEEE, 83(5):705-740, 1995.

[4] B. Clarkson and A. Pentlan, Predicting Daily Behavior via Wearable Sensors, working paper, MIT Media Lab, July, 2001.

[5] K. Etemad and R. Chellappa. Discriminant analysis for recognition of human face images. Journal of the Optical Society of America A, 14(8):1724-1733, 1997.

[6] K. Fukunaga. Introduction to Statistical Pattern Recognition (Second Edition). Academic Press, 1990.

[7] M. Kirby and L. Sirovich. Application of the karhunen-loeve procedure for the characterization of human faces. IEEE PAMI, 12(1):103-108, January 1990.

[8] C. Liu and H. Wechsler. Enhanced fisher linear discriminant models for face recognition. Proc.14th ICPR, Vol. 2, pp. 1368-1372, 1998.

[9] S. Mann. Wearable, tetherless computer--mediated reality: Wearcam as a wearable face--recognizer, and other applications for the disabled. TR 361, MIT Media Lab, Cambridge, MA, USA, 1996.

[10] B. Moghaddam and A. Pentland. Probabilistic visual learning for object recognition. IEEE PAMI, 19(7):696-710, 1997.

[11] A. Pentland, Wearable intelligence, *Scientific American Presents*, 9(4), pp. 90-95, 1998.

[12] A. Pentland, B. Moghaddam, and T. Starner. View-based and modular eigenspaces for face recognition. Proceedings of the 1994 Conference on Computer Vision and Pattern Recognition, pages 84-91, Seattle, WA, 1994.

[13] A. Samal and P. Iyengar. Automatic recognition and analysis of human faces and facial expressions: A survey. Pattern Recognition, 25(1):65-77, 1992.

[14] D. Swets and J.Weng. Using discriminant eigenfeatures for image retrieval. IEEE PAMI, 18(8):831-836, August 1996.

[15] M. Turk and A. Pentland. Eigenfaces for recognition. Journal of Cognitive Neuroscience, 3(1):72-86, 1991.

[16] ORL web site: http://www.cam-orl.co.uk.

[17] Yale face database: http://cvc.yale.edu/projects/yalefaces/yalefaces.html

[18] J. Yang, H. Yu, W. Kunz, An efficient LDA algorithm for face recognition," International Conference on Automation, Robotics, and Computer Vision, Singapore, December, 2000.

[19] W. Zhao, R. Chellappa, and N. Nandhakumar. Empirical performance analysis of linear discriminant classifiers. Proceedings of the 1998 Conference on Computer Vision and Pattern Recognition, pages 164-169, Santa Barbara, CA, 1998.

[20] W. Zhao, R. Chellappa, and P. Philips. Subspace Linear Discriminant Analysis for Face Recognition. Technical Report CAR-TR-914. Center for Automation Research, University of Maryland, 1999.