# BIOMETRIC IMAGE PROCESSING AND RECOGNITION

*P. Jonathon Phillips*
*R. Michael McCabe*
Visual Image Processing Group
National Institute of Standards and Technology
Gaithersburg, MD 20899
e-mail: jonathon@nist.gov
e-mail: mccabe@nist.gov

*Rama Chellappa*
Dept. of Electrical Engineering
and Center for Automation Research
University of Maryland
College Park, MD 20742
Tel: 301-405-4526; fax: 301-314-9115
e-mail: rama@cfar.umd.edu

## Abstract

*Biometric-based identification and verification systems are poised to become a key technology, with applications including controlling access to buildings and computers, reducing fraudulent transactions in electronic commerce, and discouraging illegal immigration. There are at least eight image-based biometrics that are being actively considered. In image-based biometrics, the biometric signature is acquired as an image and the image is processed using techniques from computer vision, image understanding, and pattern recognition. We consider two promising image-based biometrics, faces and fingerprints. For each, we provide a critical assessment of the state of the art, suggest future research directions, and identify technological challenges.*

## 1 Introduction

In general terms, a biometric is observed data of a human that allows the identity of that person to be determined. Examples of biometrics actively being investigated are DNA, shape of the ear, faces, fingerprints, hand geometry, irises, pattern of keystrokes on a keyboard, signature, and speech. While each biometric has its own strengths and weaknesses, for a biometric to be effective it should have the following four properties: (1) universality, all members of population being identified should possess the biometric; (2) uniqueness, biometric signature should be different for all members of the population; (3) invariance, the signature should be invariant under the conditions that it will be collected; and (4) resistance, the biometric should be resistant to potential countermeasures.

Some biometrics can be collected without a person's knowledge (faces, speech, signature, keystrokes, and ears), others (fingerprints and hand geometry) require a person's cooperation, and for some (irises and DNA) this determination has not been made. This distinction is important. Potential applications for non-intrusive biometrics include remote surveillance and monitoring (border crossings and airport security), while potential applications for cooperative biometrics include access control and verifying transactions executed via electronic commerce.

Faces, fingerprints, irises, ears are intrinsically image based and require image processing, pattern recognition, and computer vision techniques to implement. Whereas, hand geometry, keystrokes, signature, and speech fall into the domains of signal processing and pattern recognition. Recent efforts have looked at combining multiple biometrics [3, 16] (audio-video, faces-fingerprints, etc).

Before proceeding further, we wish to point out two recognition subproblems, identification and verification. In identification, the algorithm is presented with a biometric image of an unknown person. The algorithm reports its best estimate of the identity of the unknown person from the database of known individuals. In a more general response, the algorithm will report a list of the most similar individuals in the database. The majority of the identification applications are in law enforcement, forensics, and intelligence. The applications include identifying faces from mug shots, surveillance images, newspapers, photographs, and images of deceased people.

In verification (also referred to as authentication), the algorithm is presented with a biometric image and

a claimed identity of the person. The algorithm either accepts or rejects the claim. Or, the algorithm can return a confidence measure of the validity of the claim. Verification applications include control access to apartment or secure buildings, verifying identities during point of sale transactions , and continuous verification of identity at computer terminals or in secure facilities.

In the pages allotted in the conference proceedings, we cannot discuss all the biometrics in detail. We will focus on faces and fingerprints. We discuss the state of the art and future research directions for these biometrics from both academic and industry perspective.

## 2 Face Recognition

For a face recognition system to be successfully deployed, it must be fully automatic. A fully automatic system detects and identifies/verifies a face in an image or video sequence without human intervention. Fully automatic face recognition systems generally have two components, recognition and detection. The recognition component identifies or verifies the face. Usually, a recognition module requires that the face be in a standard position. Currently, recognition is performed exclusively from a still image or a single video frame. The majority of the algorithms are projection view-based. In a view-based algorithm, the face is represented as a set of two-dimensional images. Thus, the variations of a face under different viewing conditions (those generated by changes in light or pose), are stored as separate images. In a projection-based algorithm, the image is projected onto a lower dimensional subspace, which is referred to as *face space*. Face space characterizes the differences among faces. Therefore, in a projection view-based algorithm, a face is represented as a set of points in face space. The similarity between faces is measured by a similarity distance in face space (the smaller the distance, the greater the similarity). In identification, a face is identified as the person in the database that minimizes the distance between the unknown face and the faces in the database. In verification, the claimed identity is accepted if the similarity between the presented face and the claimed face is below a set threshold.

Although over twenty five years of research has been carried out in this area [6], it is only in the last decade that tangible progress has been made towards developing fieldable algorithms [36]. Existing algorithms can be broadly classified as holistic [11, 20, 21, 34, 35, 37, 42] or feature-based [25, 26, 40]. Neural network methods have been demonstrated for face detection, but have not been scaled up to recognition applications.

Recognition algorithms have received the majority of the attention in the academic literature. However, developing algorithms to detect and locate a face is equally important and is as difficult to develop as recognition algorithms. To date, detection algorithms have been developed to either detect and locate the face in "mugshot" style images, or detect faces in images that contain multiple faces. In a mugshot style image, there is one face in the image and it occupies a majority of the pixels.

For mugshot style images, a detection algorithm locates the face in the image, detects a set of facial features, and then transforms the image into a standard geometric configuration. The basic detection method detects and locates the eyes and transforms the face so that the center of the eyes are on specified pixels (there exist more sophisticated geometric normalization methods).

Face detection and recognition components have been combined to yield fully automatic face recognition systems for mugshot style images [20, 21, 40]. Both fully automatic face recognition systems and recognition algorithms for mugshot style images have been tested by the independent FERET evaluation procedures (described later).

A number of techniques have been developed for detecting faces in "non-mugshot" images, a few notable approaches are neural networks-based [32, 33], support vector machines based [24], and Pfinder [41]. For still images, the neural network and support vector machine-based algorithms have been extensively tested in the laboratory. Performance points to good detection rates (between 78-90%) with acceptable false detection rates [24, 32, 33]. Pfinder allows for camera rotation and zoom. These systems have not been subjected to an independent evaluation procedure such as FERET, or large scale evaluations in operational environments such as airports and outdoor scenes. Simultaneously detecting faces in visible and infrared imagery could potentially improve performance [38]. By combining image stabilization [5, 23] and face detection algorithms, one should be able to perform "face detection on the move". One can then track a face in a scene using image mosaics.

There are two basic classes of face recognition systems, still and video. Video systems detect faces by a combination of motion, stereo, color, and facial pixel patterns, whereas, still images are restricted to color and facial pixel patterns [2, 7]. In still images, the face is detected, segmented, and passed to the recognition component. In video, the algorithm performs the same tasks as in still imagery, but also as the option of selecting the frame to process. Criteria for selecting the

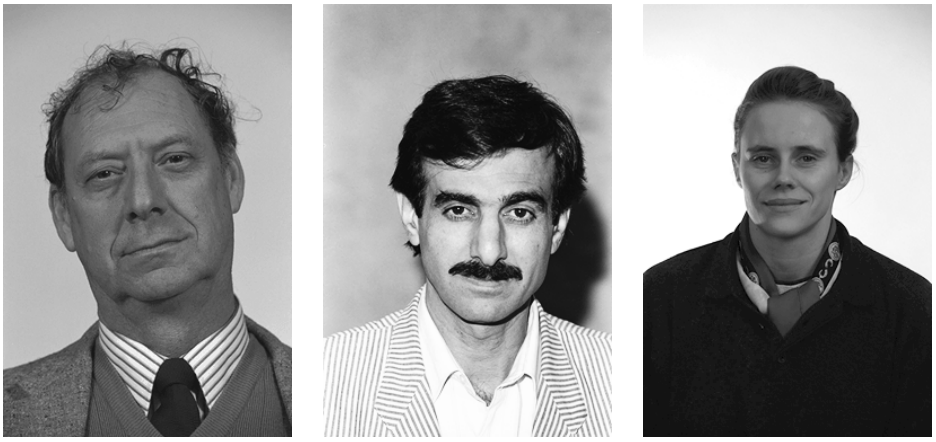

**Figure 1. Examples of images from the FERET database.**

**Table 1. Variations in performance on different categories of probes.**

| Probe Category | Gallery size | Probes scored | Average identification rate (%) |
|---|---|---|---|
| **FB** | 1196 | 1195 | 87 |
| Wearing glasses | 1196 | 176 | 75 |
| Duplicates | 1196 | 722 | 43 |
| **fc** | 1196 | 196 | 38 |
| Duplicates—18 months | 864 | 234 | 27 |

frame include, pose, lighting, and image quality.

The performance of face recognition algorithms has been measured by three FERET evaluations. The FERET evaluations are independently administered, with a set of facial images provided to researchers for development and a sequestered set of images for testing. The first two FERET evaluation were the Aug94 and Mar95 tests [29], and evaluated fully automatic algorithms. The most recent is the Sep96 test, which was administered in September 1996 and March 1997 (details of the Sep96 test can be found in Phillips et al. [27, 28, 31]) All FERET performance results quoted in this paper are from these papers. A total of twelve algorithms from six different groups were evaluated.

The Sep96 FERET evaluation procedure was designed to evaluate both fully automatic and recognition algorithms on identification and verification tasks. To obtain a robust assessment of the state-of-the-art in face recognition, strengths and weaknesses of individual algorithms, and identify future directions of research, the test was designed so that performance can be computed on multiple galleries and probe sets. (The *gallery* is the set of known individuals. An image of an unknown face presented to the algorithm is called a *probe*, and the collection of probes is called the *probe set.*)

The test consisted of images from the FERET database of facial images [29]. Some examples of images in the FERET database are in figure 1. The images in the database were collected in sets of images, where each set of images consisted 5-11 poses of an individual. Each set of images contained two frontal images that were taken within five minutes of each other. Different facial expression were requested for both frontal images. In the Sep96 FERET test, the gallery consisted of one frontal image of 1196 individuals. The face recognition algorithms were evaluated on five categories of probes. One category was the **FB** probe set, which consisted of the second frontal image from the same set as the gallery image. This probe set measured performance when there was at most five minutes between the acquisition of the gallery and probe images of a person. Performance on the **FB** probes provides an upper bound on the performance of state-of-the-art face recognition systems.

Another category of probes is all frontal duplicates.

A duplicate is an image taken in a different session (a different date) or taken under *special circumstances* (such as the subject was wearing glasses, different hair length, etc) than the gallery image.

In this paper we present summary scores. The summary scores reflect overall performance, but at the same time ignore may aspects of performance. For identification, we report the percentage of probes that are correctly identified. For **FB** probes the top algorithms had performance scores of 96%–98% and for duplicates 58%–60%.

For verification, the are two statistics that characterize algorithm performance [10, 31]. The first is the probability of rejecting a correct identity; i.e., if the reader presents herself to the system, the probability that the system will reject the claim. The second is the probability of accepting a false claim; i.e., the probability that the system accepts the first author's claim to be President Clinton. There is a tradeoff between the two types of error. At one extremes, the system can rejects all people; at the other, all claims are accepted. The operational performance point depends on the application. The equal error rate is the performance point where the two statistics are equal. For **FB** images the equal error rate is 1-2% and for duplicates is 15-18%.

In face recognition, performance is a function of the types of images in the gallery and probe set. To quantify this difference, we compare the difficulty of each probe category by averaging the identification rate for all ten recognition algorithms that took the Sep96 FERET test. Table 1 presents identification rates for five categories of probes. The first category consists of the second frontal images (**FB**) from the same set as the gallery image. The second category consists of images of people wearing glasses. The third category is the set of all duplicate probes. The fourth category consists of the frontal **fc** images (images taken on the same day, but with a different camera and lighting). This category shows the effect of a lighting change on performance. The final category consists of images that were taken at least a year apart. (The gallery in this category is a subset of the gallery in the other categories.) (Note: for the majority of the images with people wearing glasses were taken on the same day as the corresponding gallery images. None of the people in the gallery were wearing glasses.)

On the surface it may appear that performance needs to be improved for identification from mugshots. However, sketching out the details of a possible scenario presents a more optimist future. For a number identification scenarios, the output from the recognition system will be a list of the most similar individuals in the gallery to the probe. From the list a person will make the final decision. The question then becomes how many images is a person willing to look at. For duplicate probes in the Sep96 FERET test, the best algorithms report the correct answer in the top five images 70% of the time from a gallery of 1196 individuals. There is empirical evidence that performance changes linearly with the size of the gallery [31]. This means that for duplicates, potentially the correct answer will be in the top 50 images 70% of the time for a database of approximately 10,000 individuals. For many applications, using collateral information such as sex, race, approximate height, and age, it is possible to reduce the size of database that needs to be searched to at most 10,000.

Table 1 shows that future areas of research in face recognition include developing algorithms to handle variations due to lighting changes and when images of the same person are taken more than a year apart. In fact, for the majority of the algorithms, the performance on changes in lighting is the same as images of the same person taken at least 18 months apart.

Most discussions of performance of face recognition algorithms concentrate on performance of a single gallery for a category of probes. This means that relative performance among algorithms is based on a single estimate, and that this score is predictive of performance for an entire category of probes. Phillips et al. [22, 27, 31] show that performance is dependent on the images in the gallery and probe sets. This suggests that the effects of the images in the gallery and probe set on performance need to be further investigated. In fact, differences in performance among galleries and probe sets can overshadow difference in performance between algorithms [22].

For verification tasks on images taken closely in time under similar light conditions, the best face recognition algorithms are sufficiently mature to be field tested. However, for duplicate images more work still needs to be done. The evidence from the FERET tests is that the upper limits of performance from duplicates has not been reached. The main hurdle to the development of algorithms that can recognition faces in duplicate images is the availability of a large database of duplicates. The main advantage of faces is that they can be acquired non-intrusively. In numerous applications, non-intrusive methods are the only option.

## 3 Fingerprint recognition

Traditionally, inked fingerprint impressions recorded on 38mm x 40mm (1.5"x1.6") areas of a fingerprint card are scanned and processed by an automated fingerprint identification system (AFIS). However, the

(a) (b) (c)

**Figure 2. Examples of different classes of fingerprints. (a) right loop (b) whorl and (c) arch.**

inking process is being replaced by live-scan technology that relies on a moving light source and the principle of frustrated total internal reflection (FTIR) [15].

Images captured by the live-scan reader can be directly inputed into an AFIS for subsequent processing. Benefits of this new technology include the elimination of ink, determination of image quality before recording, multiple copies of the same image from a single scanning, and the immediate creation of a file containing the electronic fingerprint image.

Identification from fingerprints is a two stage process [12, 19]. The first stage, classification, performs a coarse classification of fingerprints into one of five classes. The second stage performs matching of details present in each fingerprint image. This two stage process is required so that efficient queries against databases of up to 50 million sets of fingerprints can be performed. Filtering based on fingerprint pattern classification from one or more fingers accomplishes this. Fingerprints can easily be classified into one of five classes or types: arch, tented arch, left loop, right loop, and whorl, see figure 2. Because of the large number of fingerprints in most databases, the classification accuracy of the algorithm used should approach 99%. There are at least four major approaches to automatic fingerprint classification. These are structural [17], syntactic [30], statistical [4], and the artificial neural network (ANN) [18, 39] approach. In an evaluation published about four years ago it was found that the ANN methods performed the best [4].

The second stage performs a search against the candidate fingerprints from the first stage using minutia. A minutiae is the point where a ridge either terminates or bifurcates into two or more ridges and is defined in terms of $x$ and $y$ coordinates, and ridge orientation angle. A fingerprint consists of ridges which are the raised portions of the skin and the valleys between the ridges. Figure 3 is an illustration of the method used to identify a minutiae. In the illustration the three hatched areas marked **A** represent ridges. The areas marked **B** are the valleys. At the tip of the middle ridge, the location of the minutiae is identified by its $x$ and $y$ coordinates. The orientation of the minutiae is the angle between the horizontal axis and the direction of the ridge.

The comparisons and matching of minutia between two fingerprint images are performed by an AFIS. In the first step, the AFIS rotates and translates the probe fingerprint into a standard position. After being rotated and translated, the similarity between probe and candidate fingerprints is a function of the geometric pattern of the minutia. A high score indicates a high probability of an identification. The fingerprint examiner is presented a minimum list of likely candidates based on these scores.

Due to the storage requirements of these uncompressed image files, they are generally not retained online. To reduce file size, the JPEG algorithm has been used to compress the fingerprint images. However, due to the 8x8 pixel tiling used in JPEG's DCT, blocking effects begin to appear as compression ratios exceed 8:1. A different approach, based on wavelet technology, was developed and adopted for use by the FBI. Rather than using an 8x8 pixel tile size, the wavelet scalar quantization (WSQ) algorithm globally compresses the image. This enables a compression ratio of 15:1 with minimal visual degradation in reconstructed images.

Because of nearly 20 years of operational experience by law enforcement agencies, we conclude that automated fingerprint identification is a mature technology. Law enforcement agencies have performed large scale evaluation of prototype AFISs. However, the results of these evaluations have not been publicly released. Whereas, for face recognition most of the

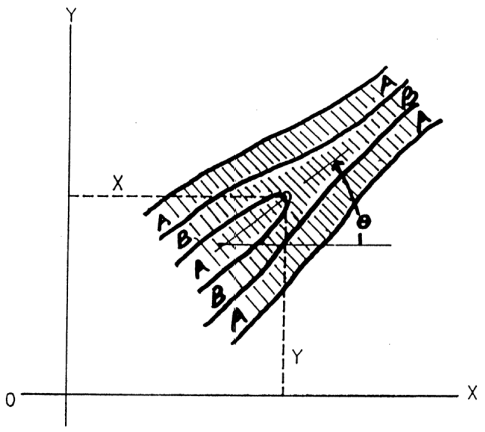


**Figure 3. Illustration of a minutiae**

underlying research has been performed in the academic arena; for fingerprints the significant technological progress has been made in the commercial arena. This is supported by the establishment of standards for inter-operability among vendors [1], fingerprint compression standard [13], and fingerprint image quality standard [14].

To date, fingerprints have been primarily used by law enforcement applications and for background clearances. With the availability of inexpensive live scanners, fingerprints are beginning to be applied to verification applications. In the past few years, new approaches to fingerprint readers have been introduced with some high volume prices targeted below $50 a unit. These new capture devices capable of capturing finger ridge structures are based on CCDs, CMOS, capacitance, thermal effects, and electrical fields. Some readers scan pieces of the finger as it is swiped across a narrow window and then electronically sews the pieces back together. Recently, keyboard manufacturers have begun integrating finger scanners directly into their units.

## 4 Discussion and future directions

In this paper we have primarily discussed face and fingerprints. A third visual biometric receiving considerable attention is iris [8, 9]. Recognition from irises is based on wavelet features derived from the texture patterns of the iris. Like fingerprints, the irises are phenotypical (features are a function of the interaction of genetics, environment, and development), whereas faces are genotypical (features are primarily genetically inherited). Iris recognition system is a three stage system requiring face detection, iris acquisition, and recognition.

To date the primary applications for biometrics has been identification in law enforcement. This market is saturated, with fingerprints being the established biometric. The future of biometrics is in verification application. Currently, for commercial applications, prototype verification systems are undergoing operational testing and demonstration.

The determination of which biometric, if any, is accepted by the community is a function of four factors: performance accuracy, cost of the system, losses prevented, and acceptance by the users. The first three factors are primarily economic. To be cost effective a system must save money or prevent losses in excess of the cost of installing and operating it. For example, operating a unit at all points of sales for consumer transaction maybe more expensive than the potential losses prevented. Whereas, for large financial transactions, operating a smaller number of systems maybe cost effective.

An equally important factor, especially for large scale applications, is the willingness of the user community to accept the biometric. For applications such as controlling access to sensitive information, acceptance of the biometric could be a condition of employment. However, for large scale use, such as with automatic teller machines (ATMs) or credit cards, overall societal and legal concerns will have to be addressed. In fact, this has the potential of becoming the biggest impediment to large scale use. For example, in the US, there is resistance to fingerprints because of criminal connotations. Also, there is the perception of "big brother" watching over us.

With the large scale use of biometrics, will come the large scale collection of information about society. Questions will arise about ownership, collection, distribution, and unauthorized transfer of this information. The answer to each of these questions has ethical and legal ramifications that need to be addressed by society.

In the law enforcement community there exist fingerprint standards to insure inter-operability among products and a best practices recommendations for mugshots (ANSI/NIST-ITL 1a-1997). For verification (faces, fingerprints, and irises), there currently do not exist standards. However, efforts are underway to establish application programming interfaces (APIs) for biometrics. A biometric API will make it possible to interchange biometrics at the system level (either changing between biometrics or implementations of the same biometric). This is different from law enforcement,

where the standards are at the sensor and image level.

For fingerprints, the technology has matured to a level suitable for incorporation into operational systems. For irises, the major research direction is acquiring the irises at increasely greater distances. To take full advantage of the potential of face recognition requires a number of technological advances. Among these are the ability to handle variations in lighting and pose, and detecting and tracking people in video. The extension to video will allow systems to continuously verify and monitor individuals. Thus, face recognition can operate with varying degrees of cooperation, from images acquired from cooperative subjects to images collected from surveillance operations. One of the most challenging problems in face recognition is to account for aging, which will be useful in recognizing missing and exploited children.

With the availability of digital cameras, face detection and recognition from video is attracting considerable attention. Several video systems that detect and identity faces in real-time on a windows Pentium platform are becoming available [2].

The interest in face recognition goes beyond biometrics. There are applications of face recognition to low bit video compression, human computer interfaces, gesture recognition, smart kiosks, and analysis/synthesis of faces. Because face recognition is a basic function of the human visual system, there is also significant interest in face recognition from neuroscience and pyschophysics.

An active area of research will be multi-modal recognition (use of two or more biometrics simultaneously). The interest in multi-modal biometrics is the potential to increase performance and make the system more robust. Multi-modal biometrics can be more robust because they can compensate for noise in one of the biometrics. It only makes sense to use multi-modal biometrics if the extra cost produces a corresponding increase in performance.

Over the last thirty years, image processing and understanding research have successfully produced commercial products in image compression and machine inspection. Over the last five years, numerous commercial products for biometrics have been marketed and demonstrated in private, military and government institutions. With the advent of electronic commerce on the Internet and World Wide Web, we believe that biometric driven technology will become ubiquitous. This presents enormous opportunities and challenges from technological, engineering, economic, legal, and societal points of view.

## References


[1] American National Standards Institute. *Data Format for the Interchange of Fingerprint Information.* ANSI/NIST-CSL 1-1993. American National Standards Institute, New York, 1995.

[2] J. Atick, P. Griffin, and A. N. Norman. face recognition from live video for real-world applications—now. *Advanced Imaging*, 10(5):58–62, May 1995.

[3] J. Bigun, G. Chollet, and G. Botgefors, editors. *1st Inter. Conf. on Audio- and video-based biometric person authentication.* LNCS 1206. Springer, Berlin, 1997.

[4] J. L. Blue, G. T. Candela, P. J. Grother, R. Chellappa, and C. L. Wilson. Evaluation of pattern classifiers for fingerprint and OCR applications. *Pattern Recognition*, 27(4):485–501, 1994.

[5] P. Burt and P. Anandan. Image stabilization by registration to a reference mosaic. In *Proc. DARPA Image Understanding Workshop*, pages 425–434, 1994.

[6] R. Chellappa, C. L. Wilson, and S. Sirohey. Human and machine recognition of face: A survey. *Proceedings of the IEEE*, 83:704–740, 1995.

[7] T. Darrell, G. Gordon, J. Woodfill, and M. Harville. A virtual mirror using real-time robust face tracking. In *3rd International Conference on Automatic Face and Gesture Recognition*, pages 616–621, 1998.

[8] J. Daugman. High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. PAMI*, 15(11):1148–1161, 1993.

[9] J. Daugman. Phenotypic versus genotypic approaches to face recognition. In P. J. Phillips, V. Bruce, F. F. Soulie, and T. S. Huang, editors, *Face Recognition: From Theory to Applications.* Springer-Verlag, Berlin, 1998.

[10] J. P. Egan. *Signal Detection Theory and ROC Analysis.* Academic Press, 1975.

[11] K. Etemad and R. Chellappa. Discriminant analysis for recognition of human face images. *J. Opt. Soc. Am. A*, 14:1724–1733, August 1997.

[12] Federal Bureau of Investigation. *The science of fingerprints.* U.S. Department of Justice, Federal Bureau of Investigation, Washington, DC, 1984.

[13] Federal Bureau of Investigation. *WSQ Gray-scale Fingerprint Image Compression Specification, Criminal Justice Information Services.* IAFIS-IC-0110V2. U.S. Department of Justice, Federal Bureau of Investigation, Washington, DC, 1993.

[14] Federal Bureau of Investigation. *Electronic Fingerprint Transmission Specification, Appendix F IAFIS Image Quality Specifications.* CJIS-RS-0010 (V6R2). U.S. Department of Justice, Federal Bureau of Investigation, Washington, DC, 1998.

[15] L. Gerhardt, D. Crockett, J. Attili, and A. Presler. Fingerprint imagery using frustrated total internal reflection. In *Proceedings of the 1986 International Carnahan Conference on Security Technology*, pages 251–255, 1986.

[16] L. Hong and A. Jain. Integrating faces and fingerprints for personal identification. In *Proc. 3rd Asian Conf. Computer Vision*, pages 16–23, 1998.

[17] A. Hrechak and J. McHugh. Automated fingerprint identification using structured matching. *Pattern Recognition*, 23:893–904, 1990.

[18] A. Jain, L. Hong, and R. Bolle. On-line fingerprint verification. *IEEE Trans. PAMI*, 19:302–314, 1997.

[19] H. Lee and R. E. Gaensslen, editors. *Advances in fingerprint technology*. Elsevier, New York, 1991.

[20] B. Moghaddam and A. Pentland. Probabilistic visual learning for object detection. *IEEE Trans. PAMI*, 17(7):696–710, 1997.

[21] B. Moghaddam, W. Wahid, and A. Pentland. Beyond eigenfaces: probablistic matching for face recognition. In *3rd International Conference on Automatic Face and Gesture Recognition*, pages 30–35, 1998.

[22] H. Moon and P. J. Phillips. Analysis of pca-based face recognition algorithms. In K. W. Bowyer and P. J. Phillips, editors, *Empirical Evaluation Techniques in Computer Vision*. IEEE Computer Society Press, Los Alamitos, CA, 1998.

[23] C. H. Morimoto and R. Chellappa. Fast electronic digital image stabilization. In *Proc. International Conference on Pattern Recognition*, volume 3, pages 284–288, 1996.

[24] E. Osuna, R. Frenad, and F. Girosi. Training support vector machines: An application to face detection. In *Proceedings Computer Vision and Pattern Recognition 97*, pages 130–136, 1997.

[25] P. Penev and J. Atick. Local feature analysis: a general statistical theory for object representation. *Network: Computation in Neural Systems*, 7(3):477–500, 1996.

[26] P. J. Phillips. Matching pursuit filters applied to face identification. *IEEE Trans. on Image Processing*, (in press) 1998.

[27] P. J. Phillips, H. Moon, P. Rauss, and S. Rizvi. The FERET evaluation methodology for face-recognition algorithms. In *Proceedings Computer Vision and Pattern Recognition 97*, pages 137–143, 1997.

[28] P. J. Phillips, H. Moon, S. Rizvi, and P. Rauss. The feret evaluation. In P. J. Phillips, V. Bruce, F. F. Soulie, and T. S. Huang, editors, *Face Recognition: From Theory to Applications*. Springer-Verlag, Berlin, 1998.

[29] P. J. Phillips, H. Wechsler, J. Huang, and P. Rauss. The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing Journal*, 16(5):295–306, 1998.

[30] K. Rao and K. Black. Type classification of fingerprints: A syntactic approach. *IEEE Trans. PAMI*, 2:223–231, 1980.

[31] S. Rizvi, P. J. Phillips, and H. Moon. A verification protocol and statistical performance analyis for face recognition algorithms. In *Computer Vision and Pattern Recognition 98*, (to appear) 1998.

[32] H. A. Rowley, S. Baluja, and T. Kanade. Neural network-based face detection. *IEEE Trans. PAMI*, 20:23–28, 1998.

[33] K.-K. Sung and T. Poggio. Example-based learning for view-based human face detection. *IEEE Trans. PAMI*, 20:39–51, 1998.

[34] D. Swets and J. Weng. Using discriminant eigenfeatures for image retrieval. *IEEE Trans. PAMI*, 18(8):831–836, 1996.

[35] M. Turk and A. Pentland. Eigenfaces for recognition. *J. Cognitive Neuroscience*, 3(1):71–86, 1991.

[36] H. Wechsler, P. J. Phillips, V. Bruce, F. F. Soulie, and T. S. Huang, editors. *Face Recognition: From Theory to Applications*. Springer-Verlag, Berlin, 1998.

[37] J. Wilder. Face recognition using transform coding of gray scale projection projections and the neural tree network. In R. J. Mammone, editor, *Artifical Neural Networks with Applications in Speech and Vision*, pages 520–536. Chapman Hall, 1994.

[38] J. Wilder, P. J. Phillips, C. Jiang, and S. Wiener. Comparison of visible and infrared imagery for face recognition. In *2nd International Conference on Automatic Face and Gesture Recognition*, pages 182–187, 1996.

[39] C. L. Wilson, G. Candela, P. J. Grother, C. I. Watson, and R. A. Wilkinson. Massively parallel neural network fingerprint classification system. Technical Report NISTIR 4880, National Institute of Standards and Technology, 1992.

[40] L. Wiskott, J.-M. Fellous, N. Kruger, and C. von der Malsburg. Face recognition by elasric bunch graph matching. *IEEE Trans. PAMI*, 17(7):775–779, 1997.

[41] C. R. Wren, A. Azarbayejani, T. Darrel, and A. P. Pentland. Pfinder: Real-time tracking of the human body. *IEEE Trans. PAMI*, 19:780–785, 1997.

[42] W. Zhao, R. Chellappa, and A. Krishnaswamy. Discriminant analysis of principal components for face recognition. In *3rd International Conference on Automatic Face and Gesture Recognition*, pages 336–341, 1998.