Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

Human Recognition Using Biometrics: An Overview

Arun Ross[1] and Anil K. Jain[2]

[1] Lane Department of Computer Science and Electrical Engineering, West Virginia University, Morgantown, WV 26506 arun.ross@mail.wvu.edu

[2] Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824 jain@cse.msu.edu

Abstract

Establishing identity is becoming critical in our vastly interconnected society. Questions such as"Is she really who she claims to be?," "Is this person authorized to use this facility?," or "Is he in the watchlist posted by the government?" are routinely being posed in a variety of scenarios ranging from issuing a drivers license to gaining entry into a country. The need for reliable user authentication techniques has increased in the wake of heightened concerns about security and rapid advancements in networking, communication, and mobility. Biometrics, described as the science of recognizing an individual based on his or her physical or behavioral traits, is beginning to gain acceptance as a legitimate method for determining an individuals identity. Biometric systems have now been deployed in various commercial, civilian, and forensic applications as a means of establishing identity. This paper presents an overview of biometrics and discusses some of the pertinent terminology necessary to understand this technology. The importance of information fusion in the context of biometrics is also highlighted. It is becoming increasingly apparent that multibiometric systems will play a significant role in dispensing the role of identity management in the 21st century.

1 Introduction

Establishing the identity of an individual is of paramount importance in our highly networked society. The overarching task of any identity management system is the ability to determine or validate the identity of its users prior to granting them access to the resources protected by the system. Traditionally, passwords (knowledge-based schemes) and ID cards (token-based schemes) have been used to validate the identity of an individual intending to access the services offered by an application (e.g., online banking) or a facility (e.g., theme parks). However, such mechanisms for user authentication have several limitations. For example, passwords can be divulged to unauthorized users resulting in a breach of security; moreover, simple passwords can be easily guessed by an intruder while complex passwords can be difficult to recollect for a legitimate user. ID cards, on the other hand, can be misplaced, forged or stolen thereby undermining the security afforded by the system. Thus, it is necessary to utilize alternate methods of authentication that are not merely based on *what you know* or *what you have* but, rather, on *who you are*.

Biometrics is the science of establishing the identity of an individual based on the inherent physical or behavioral traits associated with the person [29, 4, 33]. Biometric systems utilize fingerprints, iris, face, hand geometry, palmprint, finger vein structure, gait, voice, signature, keyboard typing pattern, etc. in order to recognize a person (Figure 1). A typical biometric system operates by capturing the biometric trait of a person via an appropriately designed acquisition module and comparing the recorded trait with the biometric samples (or templates) in a database in order to determine the identity of the person (*identification*) or to validate a claimed identity (*verification*). For example, a face biometric system captures the face image of an individual, extracts a feature set from the segmented face, compares this feature set against the templates stored in the database and renders a decision regarding the identity of the individual. Thus, a generic biometric system may be viewed as a pattern recognition system in which the raw biometric data (or signal) constitutes the input pattern that is assigned a class label. In an identification system, the class label pertains to the identity of the individual while in a verification system the class label is a match (genuine) or a non-match (impostor). In both modes of operation, a *reject* label is emitted when the system is unable to determine a valid class.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

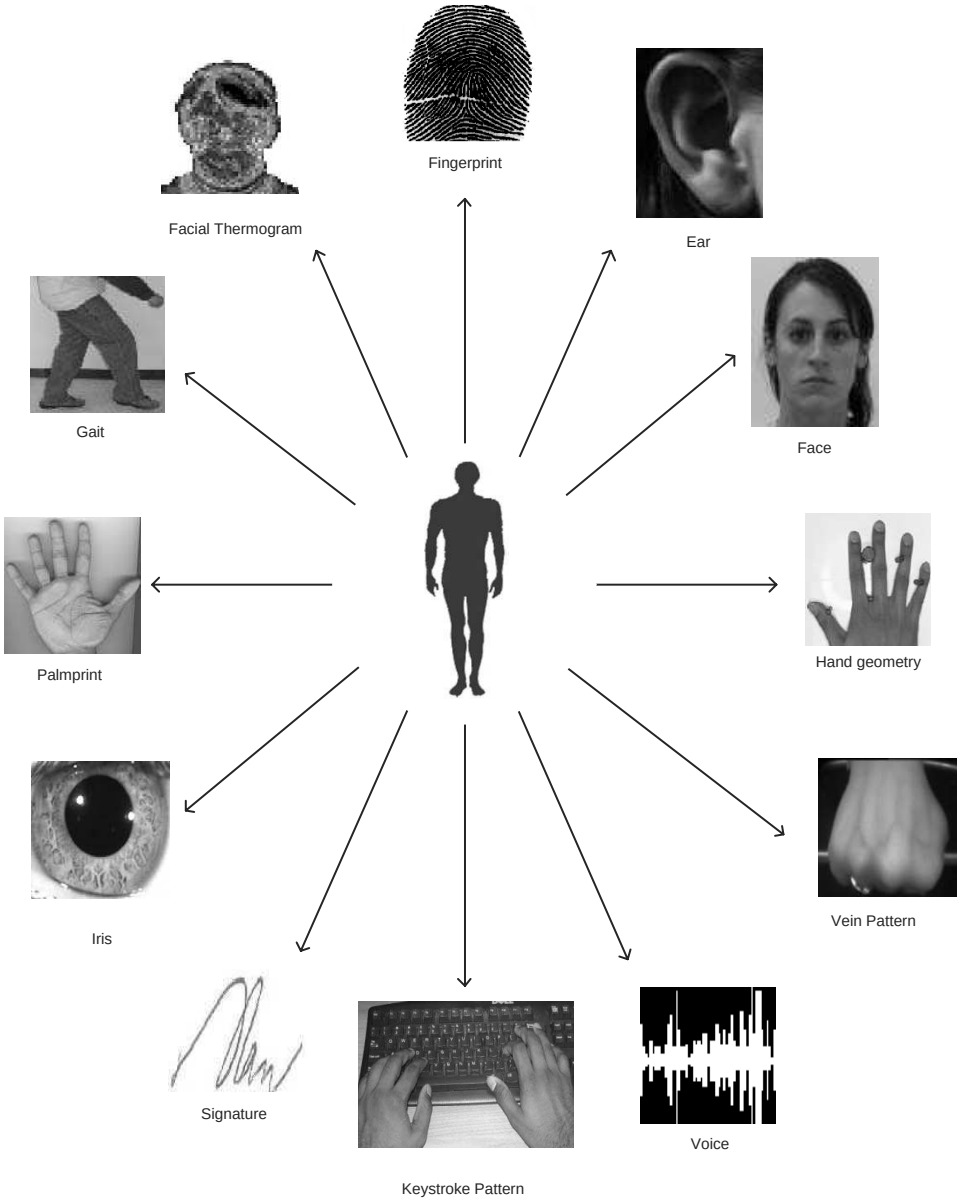

Fig. 1. Examples of biometric traits that can be used for authenticating an individual's identity.

The need for establishing identity in a reliable manner has spurred active research in the field of biometrics [66]. The deployment of biometric systems in border security programs (e.g., US-VISIT[3]), criminal investigations (e.g., IAFIS[4]), logical access points (e.g., computer login) and surveillance applications (e.g., face recognition in public spaces) further underscores the importance of designing and implementing large-scale authentication systems that can consistently render the correct decision under various operational scenarios. Furthermore, as the number of enrolled subjects increases over time, it is imperative that the matching accuracy of these systems is not compromised. Indeed, the problem of biometric recognition may be viewed as a *Grand Challenge*, given the expectations of high matching accuracy, ease of usability and efficient scalability in a variety of applications accessed by different segments of the general population [31].

[3] United States Visitor and Immigrant Status Indicator Technology
[4] Integrated Automated Fingerprint Identification System

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

2 Operation of a biometric system

A biometric system is essentially a pattern recognition system that acquires biometric data from an individual, extracts a salient feature set from the data, compares this feature set against the feature set(s) stored in the database, and executes an action based on the result of the comparison. Therefore, a generic biometric system can be viewed as having four main modules: a sensor module; a quality assessment and feature extraction module; a matching module; and a database module. Each of these modules is described below.

1. **Sensor module**: A suitable biometric reader or scanner is required to acquire the raw biometric data of an individual. To obtain fingerprint images, for example, an optical fingerprint sensor may be used to image the friction ridge structure of the fingertip. The sensor module defines the human machine interface and is, therefore, pivotal to the performance of the biometric system. A poorly designed interface can result in a high failure-to-acquire rate and, consequently, low user acceptability.
2. **Quality assessment and feature extraction module**: The quality of the biometric data acquired by the sensor is first assessed in order to determine its suitability for further processing. Typically, the acquired data is subjected to a signal enhancement algorithm in order to improve its quality. However, in some cases, the quality of the data may be so poor that the user is asked to present the biometric data again. The biometric data is then processed and a set of salient discriminatory features extracted to represent the underlying trait. For example, the position and orientation of minutia points (local ridge and valley anomalies) in a fingerprint image are extracted by the feature extraction module in a fingerprint-based biometric system. During enrollment, this feature set is stored in the database and is commonly referred to as a *template* .
3. **Matching and decision-making module**: The extracted features are compared against the stored templates to generate match scores. In a fingerprint-based biometric system, the number of matching minutiae between the input and the template feature sets is determined and a match score reported. The match score may be moderated by the quality of the presented biometric data. The matcher module also encapsulates a decision making module, in which the match scores are used to either validate a claimed identity or provide a ranking of the enrolled identities in order to identify an individual.
4. **System database module**: The database acts as the repository of biometric information. During the enrollment process, the feature set extracted from the raw biometric sample (i.e., the template) is stored in the database (possibly) along with some biographic information (such as name, Personal Identification Number (PIN), address, etc.) characterizing the user. The data capture during the enrollment process may or may not be supervised by a human depending on the application. For example, a user attempting to create a new computer account in her biometric-enabled workstation may proceed to enroll her biometrics without any supervision; a person desiring to use a biometric-enabled ATM, on the other hand, will have to enroll her biometrics in the presence of a bank officer after presenting her non-biometric credentials.

The template of a user can be extracted from a single biometric sample, or generated by processing multiple samples. Thus, the minutiae template of a finger may be extracted after mosaicing multiple samples of the same finger. Some systems store multiple templates in order to account for the intra-class variations associated with a user. Face recognition systems, for instance, may store multiple templates of an individual, with each template corresponding to a different facial pose with respect to the camera. Depending on the application, the template can be stored in the central database of the biometric system or be recorded on a token (e.g., smart card) issued to the individual.

In the face recognition literature, the raw biometric images stored in the database are often referred to as *gallery images* while those acquired during authentication are known as *probe images*. These are synonymous with the terms *stored images* and *query/input images*, respectively.

3 Performance of a biometric system

Unlike password-based systems, where a *perfect* match between two alphanumeric strings is necessary in order to validate a user's identity, a biometric system seldom encounters two samples of a user's biometric trait that result in exactly the same feature set. This is due to imperfect sensing conditions (e.g., noisy fingerprint due to sensor malfunction), alterations in the user's biometric characteristic (e.g., respiratory ailments impacting speaker recognition), changes in ambient conditions (e.g., inconsistent illumination levels in face recognition) and variations in the user's

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

interaction with the sensor (e.g., occluded iris or partial fingerprints). Thus, the distance between two feature sets originating from the same biometric trait of a user is typically non-zero (a distance score of zero would indicate that the feature sets are identical). In fact, a perfect match might be an indication that a replay attack is being launched against the system.

The variability observed in the biometric feature set of an individual is referred to as *intra*-class variation, and the variability between feature sets originating from two different individuals is known as *inter*-class variation. A useful feature set exhibits small intra-class variation and large inter-class variation.

A similarity match score is known as a *genuine* or *authentic* score if it is a result of matching two samples of the same biometric trait of a user. It is known as an *impostor* score if it involves comparing two biometric samples originating from different users. An impostor score that exceeds the threshold η results in a false accept (or, a false match), while a genuine score that falls below the threshold η results in a false reject (or, a false non-match). The *False Accept Rate (FAR)* (or, the False Match Rate (FMR)) of a biometric system can therefore be defined as the fraction of impostor scores exceeding the threshold η. Similarly, the *False Reject Rate (FRR)* (or, the False Non-match Rate (FNMR)) of a system may be defined as the fraction of genuine scores falling below the threshold η. The *Genuine Accept Rate (GAR)* is the fraction of genuine scores exceeding the threshold η. Therefore,

$$GAR = 1 - FRR. \tag{1}$$

Regulating the value of η changes the FRR and the FAR values, but for a given biometric system, it is not possible to decrease both these errors simultaneously. When a large number of genuine and impostor scores is available, one could *estimate* the probability density functions of the two sets of scores in order to analytically derive the FAR and FRR. Let $p(s|genuine)$ and $p(s|impostor)$ represent the probability density functions (or, probability distributions) of the score s under the genuine and impostor conditions, respectively. Then for a particular threshold, η,

$$FAR(\eta) = \int_{\eta}^{\infty} p(s|impostor)ds, \tag{2}$$

$$FRR(\eta) = \int_{-\infty}^{\eta} p(s|genuine)ds. \tag{3}$$

If the match score represents a distance or dissimilarity value, then $FAR(\eta)$ and $FRR(\eta)$ may be expressed as follows:

$$FAR(\eta) = \int_{-\infty}^{\eta} p(s|impostor)ds, \tag{4}$$

$$FRR(\eta) = \int_{\eta}^{\infty} p(s|genuine)ds. \tag{5}$$

Figure 2 illustrates the genuine and impostor distributions corresponding to a face biometric system. The similarity scores, in this case, are taken from the NIST BSSR1 database and originate from a matcher identified as `Face-G`.

The FAR and FRR at various values of η can be summarized using a Detection Error Tradeoff (DET) curve ([48]) that plots the FRR against the FAR at various thresholds on a *normal deviate* scale and interpolates between these points (Figure 3(a)). When a linear, logarithmic or semi-logarithmic scale is used to plot these error rates, then the resulting graph is known as a Receiver Operating Characteristic (ROC) curve ([18]). In many instances, the ROC curve plots the GAR (rather than the FRR) against the FAR (see Figure 3(b) and (c)). The primary difference between the DET and ROC curves is the use of the normal deviate scale in the former.

It is important to note that the occurrence of false accepts and false rejects is not evenly distributed across the users of a biometric system. There are inherent differences in the "recognizability" of different users. Doddington et al. [15] identify four categories of biometric users based on these inherent differences. Although this categorization (more popularly known as Doddington's zoo) was originally made in the context of speaker recognition, it is applicable to other biometric modalities as well.

1. Sheep represent users whose biometric feature sets are very distinctive and exhibit low intra-class variations. Therefore, these users are expected to have low false accept and false reject errors.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

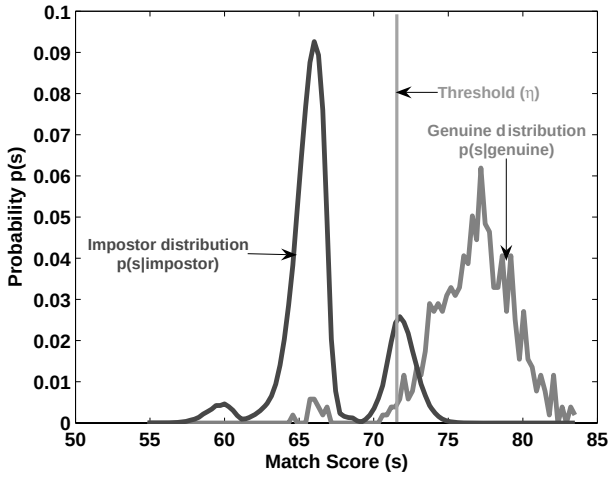

Fig. 2. The genuine and impostor distributions corresponding to the `Face-G` matcher in the NIST BSSR1 database. The threshold, η, determines the FAR and FRR of the system. Note that given these two distributions, the FAR and the FRR cannot be reduced *simultaneously* by adjusting the threshold.

2. Goats refer to users who are prone to false rejects. The biometric feature sets of such users typically exhibit large intra-class variations.
3. Lambs are users whose biometric feature set overlaps extensively with those of other individuals. The biometric feature sets of these users have low inter-class variations. Thus, a randomly chosen user (from the target population) has a high probability of being accepted as a lamb than as a sheep. The false accept rate associated with these users is typically high.
4. Wolves indicate individuals who are successful in manipulating their biometric trait (especially behavioral traits) in order to impersonate legitimately enrolled users of a system. Therefore, these users can increase the false accept rate of the system.

Doddington et al. [15] discuss the use of statistical testing procedures to detect the presence of goats, lambs and wolves in a voice biometric system. A combination of the F-test, Kruskal Wallis test and Durbin test is used to establish the occurrence of these categories of users in the 1998 NIST database of speech segments that was used in the evaluation of speaker recognition algorithms (`http://www.nist.gov/speech/tests/spk/1998/current_plan.htm`).

Besides the two types of errors (viz., false accept and false reject) indicated above, a biometric system can encounter other types of failures as well. The *Failure to Acquire (FTA)* (also known as Failure to Capture (FTC)) rate denotes the proportion of times the biometric device fails to capture a sample when the biometric characteristic is presented to it. This type of error typically occurs when the device is not able to locate a biometric signal of sufficiently good quality (e.g., an extremely faint fingerprint or an occluded face image). The FTA rate is also impacted by sensor wear and tear. Thus, periodic sensor maintenance is instrumental for the efficient functioning of a biometric system. The *Failure to Enroll (FTE)* rate denotes the proportion of users that cannot be successfully enrolled in a biometric system. User training may be necessary to ensure that an individual interacts with a biometric system appropriately in order to facilitate the acquisition of good quality biometric data. This necessitates the design of robust and efficient user interfaces that can assist an individual both during enrollment and recognition.

There is a tradeoff between the FTE rate and the perceived system accuracy as measured by FAR/FRR. FTE errors typically occur when the system rejects poor quality inputs during enrollment; consequently, if the threshold on quality is high, the system database contains only good quality templates and the perceived system accuracy improves. Because of the interdependence among the failure rates and error rates, all these rates (i.e., FTE, FTC, FAR, FRR) constitute important performance specifications of a biometric system, and should be reported during system evaluation along with the target population using the system.

The performance of a biometric system may also be summarized using other single-valued measures such as the Equal Error Rate (EER) and the d-prime value. The EER refers to that point in a DET curve where the FAR equals the FRR; a lower EER value, therefore, indicates better performance. The d-prime value (d') measures the separation

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

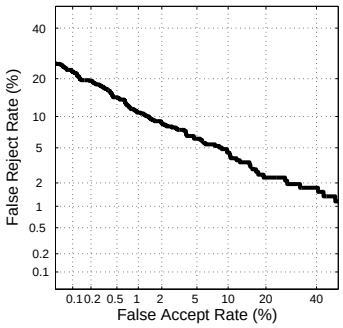

(a)

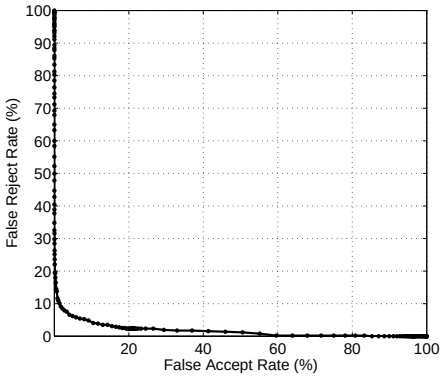

(b)

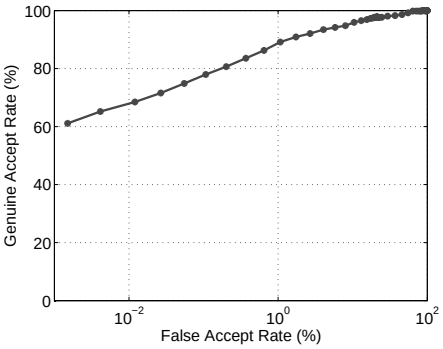

(c)

Fig. 3. The performance of a biometric system can be summarized using DET and ROC curves. In this example, the performance curves are computed using the match scores of the `Face-G` matcher from the NIST BSSR1 database. The graph in (a) shows a DET curve that plots FRR against FAR in the normal deviate scale. In (b) a ROC curve plots FRR against FAR in the linear scale, while in (c) a ROC curve plots GAR against FAR in a semi-logarithmic scale.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

between the means of the genuine and impostor probability distributions in standard deviation units and is defined as,

$$d' = \frac{\sqrt{2} \mid \mu_{genuine} - \mu_{impostor} \mid}{\sqrt{\sigma^2_{genuine} + \sigma^2_{impostor}}},$$

where the μ's and σ's are the means and standard deviations, respectively, of the genuine and impostor distributions. A higher d-prime value indicates better performance. If the genuine and impostor distributions indeed follow a normal (Gaussian) distribution with equal variance (a very unlikely situation in the practical biometric domain), then d' reduces to the normal deviate value [72].

In the case of identification, the input feature set is compared against all templates residing in the database in order to determine the top match (i.e, the best match). The top match can be determined by examining the match scores pertaining to all the comparisons and reporting the identity of the template corresponding to the largest similarity score. The *identification rate* indicates the proportion of times a previously enrolled individual is successfully mapped to the correct identity in the system. Here, we assume that the question being asked is, "Does the top match correspond to the correct identity?" An alternate question could be, "Does any one of the top k matches correspond to the correct identity?" (see [51]). The rank-k identification rate, R_k, indicates the proportion of times the correct identity occurs in the top k matches as determined by the match score. Rank-k performance can be summarized using the Cumulative Match Characteristic (CMC) curve ([51]) that plots R_k against k, for $k = 1, 2, \ldots M$ with M being the number of enrolled users. The relationship between CMC and DET/ROC curves has been discussed by [23] and [5].

The biometric of choice for a particular application is primarily dictated by the error rates and failure rates discussed above. Other factors such as the cost of the system, throughput rate, user acceptance, ease of use, robustness of the sensor, etc. also determine the suitability of a biometric system for an application.

4 Biometric characteristics

A number of biometric characteristics are being used in various applications. Each biometric has its pros and cons and, therefore, the choice of a biometric trait for a particular application depends on a variety of issues besides its matching performance. Jain et al. [29] have identified seven factors that determine the suitability of a physical or a behavioral trait to be used in a biometric application.

1. **Universality**: Every individual accessing the application should possess the trait.
2. **Uniqueness**: The given trait should be sufficiently different across individuals comprising the population.
3. **Permanence**: The biometric trait of an individual should be sufficiently invariant over a period of time with respect to the matching algorithm. A trait that changes significantly over time is not a useful biometric.
4. **Measurability**: It should be possible to acquire and digitize the biometric trait using suitable devices that do not cause undue inconvenience to the individual. Furthermore, the acquired raw data should be amenable to processing in order to extract representative feature sets.
5. **Performance**: The recognition accuracy and the resources required to achieve that accuracy should meet the constraints imposed by the application.
6. **Acceptability**: Individuals in the target population that will utilize the application should be willing to present their biometric trait to the system.
7. **Circumvention**: This refers to the ease with which the trait of an individual can be imitated using artifacts (e.g., fake fingers), in the case of physical traits, and mimicry, in the case of behavioral traits.

No single biometric is expected to effectively meet all the requirements (e.g., accuracy, practicality, cost) imposed by all applications (e.g., Digital Rights Management (DRM), access control, welfare distribution). In other words, no biometric is *ideal* but a number of them are *admissible*. The relevance of a specific biometric to an application is established depending upon the nature and requirements of the application, and the properties of the biometric characteristic. A brief introduction to some of the commonly used biometric characteristics is given below:

1. **Face**: Face recognition is a non-intrusive method, and facial attributes are probably the most common biometric features used by humans to recognize one another. The applications of facial recognition range from a static,

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

controlled "mug-shot" authentication to a dynamic, uncontrolled face identification in a cluttered background. The most popular approaches to face recognition [43] are based on either (i) the location and shape of facial attributes, such as the eyes, eyebrows, nose, lips, and chin and their spatial relationships, or (ii) the overall (global) analysis of the face image that represents a face as a weighted combination of a number of canonical faces. While the authentication performance of the face recognition systems that are commercially available is reasonable [60], they impose a number of restrictions on how the facial images are obtained, often requiring a fixed and simple background with controlled illumination. These systems also have difficulty in matching face images captured from two different views, under different illumination conditions, and at different times. It is questionable whether the face itself, without any contextual information, is a sufficient basis for recognizing a person from a large number of identities with an extremely high level of confidence. In order that a facial recognition system works well in practice, it should automatically (i) detect whether a face is present in the acquired image; (ii) locate the face if there is one; and (iii) recognize the face from a general viewpoint (i.e., from any pose).

2. **Fingerprint**: Humans have used fingerprints for personal identification for many decades. The matching (i.e., identification) accuracy using fingerprints has been shown to be very high [76, 45]. A fingerprint is the pattern of ridges and valleys on the surface of a fingertip whose formation is determined during the first seven months of fetal development. It has been empirically determined that the fingerprints of identical twins are different and so are the prints on each finger of the same person [46]. Today, most fingerprint scanners cost less than US $50 when ordered in large quantities and the marginal cost of embedding a fingerprint-based biometric in a system (e.g., laptop computer) has become affordable in a large number of applications. The accuracy of the currently available fingerprint recognition systems is adequate for authentication systems in several applications, particularly forensics. Multiple fingerprints of a person (e.g., ten-prints used in IAFIS) provide additional information to allow for large-scale identification involving millions of identities. One problem with large-scale fingerprint recognition systems is that they require a huge amount of computational resources, especially when operating in the identification mode. Finally, fingerprints of a small fraction of the population may be unsuitable for automatic identification because of genetic factors, aging, environmental or occupational reasons (e.g., manual workers may have a large number of cuts and bruises on their fingerprints that keep changing).
3. **Hand geometry**: Hand geometry recognition systems are based on a number of measurements taken from the human hand, including its shape, size of palm, and the lengths and widths of the fingers [79]. Commercial hand geometry-based authentication systems have been installed in hundreds of locations around the world. The technique is very simple, relatively easy to use, and inexpensive. Environmental factors such as dry weather or individual anomalies such as dry skin do not appear to adversely affect the authentication accuracy of hand geometry-based systems. However, the geometry of the hand is not known to be very distinctive and hand geometry-based recognition systems cannot be scaled up for systems requiring identification of an individual from a large population. Furthermore, hand geometry information may not be invariant during the growth period of children. In addition, an individual's jewelry (e.g., rings) or limitations in dexterity (e.g., from arthritis), may pose challenges in extracting the correct hand geometry information. The physical size of a hand geometry-based system is large, and it cannot be embedded in certain devices like laptops. There are authentication systems available that are based on measurements of only a few fingers (typically, index and middle) instead of the entire hand. These devices are smaller than those used for hand geometry, but still much larger than those used for procuring certain other traits (e.g., fingerprint, face, voice).
4. **Palmprint**: The palms of the human hands contain pattern of ridges and valleys much like the fingerprints. The area of the palm is much larger than the area of a finger and, as a result, palmprints are expected to be even more distinctive than the fingerprints [78, 39]. Since palmprint scanners need to capture a large area, they are bulkier and more expensive than the fingerprint sensors. Human palms also contain additional distinctive features such as principal lines and wrinkles that can be captured even with a lower resolution scanner, which would be cheaper [17]. Finally, when using a high-resolution palmprint scanner, all the features of the hand such as geometry, ridge and valley features (e.g., minutiae and singular points such as deltas), principal lines, and wrinkles may be combined to build a highly accurate biometric system.
5. **Iris**: The iris is the annular region of the eye bounded by the pupil and the sclera (white of the eye) on either side. The visual texture of the iris is formed during fetal development and stabilizes during the first two years of life (the pigmentation, however, continues changing over an extended period of time [75]). The complex iris texture carries very distinctive information useful for personal recognition [13]. The accuracy and speed of currently deployed iris-based recognition systems is promising and support the feasibility of large-scale identification systems based

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

on iris information. Each iris is distinctive and even the irises of identical twins are different. It is possible to detect contact lenses printed with a fake iris (see [12]). The hippus movement of the eye may also be used as a measure of liveness for this biometric. Although early iris-based recognition systems required considerable user participation and were expensive, the newer systems have become more user-friendly and cost-effective [54, 20]. While iris systems have a very low False Accept Rate (FAR) compared to other biometric traits, the False Reject Rate (FRR) of these systems can be rather high [28].

6. **Keystroke**: It is hypothesized that each person types on a keyboard in a characteristic way. This biometric is not expected to be unique to each individual but it may be expected to offer sufficient discriminatory information to permit identity verification [50]. Keystroke dynamics is a behavioral biometric; one may expect to observe large intra-class variations in a person's typing patterns due to changes in emotional state, position of the user with respect to the keyboard, type of keyboard used, etc. The keystrokes of a person could be monitored unobtrusively as that person is keying in information. This biometric permits "continuous verification" of an individual's identity over a session after the person logs in using a stronger biometric such as fingerprint or iris.
7. **Signature**: The way a person signs her name is known to be a characteristic of that individual [52, 42]. Although signatures require contact with the writing instrument and an effort on the part of the user, they have been accepted in government, legal, and commercial transactions as a method of authentication. With the proliferation of PDAs and Tablet PCs, on-line signature may emerge as the biometric of choice in these devices. Signature is a behavioral biometric that changes over a period of time and is influenced by the physical and emotional conditions of the signatories. Signatures of some people vary substantially: even successive impressions of their signature are significantly different. Further, professional forgers may be able to reproduce signatures that fool the signature verification system [25].
8. **Voice**: Voice is a combination of physical and behavioral biometric characteristics [7]. The physical features of an individual's voice are based on the shape and size of the appendages (e.g., vocal tracts, mouth, nasal cavities, and lips) that are used in the synthesis of the sound. These physical characteristics of human speech are invariant for an individual, but the behavioral aspect of the speech changes over time due to age, medical conditions (such as common cold), emotional state, etc. Voice is also not very distinctive and may not be appropriate for large-scale identification. A text-dependent voice recognition system is based on the utterance of a fixed predetermined phrase. A text-independent voice recognition system recognizes the speaker independent of what she speaks. A text-independent system is more difficult to design than a text-dependent system but offers more protection against fraud. A disadvantage of voice-based recognition is that speech features are sensitive to a number of factors such as background noise. Speaker recognition is most appropriate in telephone-based applications but the voice signal is typically degraded in quality by the communication channel.
9. **Gait**: Gait refers to the manner in which a person walks, and is one of the few biometric traits that can be used to recognize people at a distance. Therefore, this trait is very appropriate in surveillance scenarios where the identity of an individual can be surreptitiously established. Most gait recognition algorithms attempt to extract the human silhouette in order to derive the spatio-temporal attributes of a moving individual. Hence, the selection of a good model to represent the human body is pivotal to the efficient functioning of a gait recognition system. Some algorithms use the optic flow associated with a set of dynamically extracted moving points on the human body to describe the gait of an individual [55]. Gait-based systems also offer the possibility of tracking an individual over an extended period of time. However, the gait of an individual is affected by several factors including the choice of footwear, nature of clothing, affliction of the legs, walking surface, etc.

Most biometric systems that are presently in use, typically use a single biometric trait to establish identity (i.e., they are unibiometric systems). For example, the Schiphol Privium scheme at Amsterdam's Schipol airport employs iris scan smart cards to speed up the immigration process; the Ben Gurion International Airport at Tel Aviv employs automated hand geometry-based identification kiosks to enable Israeli citizens and frequent international travelers to rapidly negotiate the passport inspection process; some financial institutions in Japan have installed palm-vein authentication systems in their ATMs to validate the identity of a customer conducting a transaction; customers phoning in to schedule product shipments through Union Pacific's railcar system are authenticated by a speaker recognition software. With the proliferation of biometric-based solutions in civilian and law enforcement applications, it is important that the vulnerabilities and limitations of these systems are clearly understood. Some of the challenges commonly encountered by biometric systems are listed below.

1. Noise in sensed data: The biometric data being presented to the system may be contaminated by noise due to imperfect acquisition conditions or subtle variations in the biometric itself. For example, a scar can change a

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

subject's fingerprint while the common cold can alter the voice characteristics of a speaker. Similarly, unfavorable illumination conditions may significantly affect the face and iris attributes of a person. Noisy data can result in an individual being incorrectly labeled as an impostor thereby increasing the False Reject Rate (FRR) of the system.

2. Non-universality: The biometric system may not be able to acquire meaningful biometric data from a subset of individuals resulting in a failure-to-enroll (FTE) error. For example, a fingerprint system may fail to image the friction ridge structure of some individuals due to the poor quality of their fingerprints. Similarly, an iris recognition system may be unable to obtain the iris information of a subject with long eyelashes, drooping eyelids or certain pathological conditions of the eye[5]. Exception processing will be necessary in order to accommodate such users into the authentication system.
3. Upper bound on identification accuracy: The matching performance of a unibiometric system cannot be continuously improved by tuning the feature extraction and matching modules. There is an implicit upper bound on the number of distinguishable patterns (i.e., the number of distinct biometric feature sets) that can be represented using a template. The capacity of a template is constrained by the variations observed in the feature set of each subject (i.e., *intra*-class variations) and the variations between feature sets of different subjects (i.e., *inter*-class variations). Golfarelli et al. [22] state that the number of distinguishable patterns in two of the most commonly used representations of hand geometry and face are only of the order of 10^5 and 10^3, respectively. Table 1 lists the error rates associated with four biometric modalities - fingerprints, face, voice, iris - as suggested by recent public tests. These statistics suggest that there is a tremendous scope for performance improvement especially in the context of large-scale authentication systems.
4. Spoof attacks: Behavioral traits such as voice [19] and signature [25] are vulnerable to spoof attacks by an impostor attempting to mimic the traits corresponding to legitimately enrolled subjects. Physical traits such as fingerprints can also be spoofed by inscribing ridge-like structures on synthetic material such as gelatine and play-doh [49, 63]. Targeted spoof attacks can undermine the security afforded by the biometric system and, consequently, mitigate its benefits [64].

Some of the limitations of a unibiometric system can be addressed by designing a system that consolidates *multiple* sources of biometric information. This can be accomplished by fusing, for example, multiple biometric systems utilizing different traits, or multiple feature extraction and matching algorithms operating on the same biometric. Such systems, known as multibiometric systems [68, 32, 27], can improve the matching accuracy of a biometric system while increasing population coverage and deterring spoof attacks. In this chapter, the various sources of biometric information that can be fused as well as the different levels of fusion that are possible are discussed.

Table 1. The false accept and false reject error rates (FAR and FRR) associated with the fingerprint, face, voice and iris modalities. The accuracy estimates of biometric systems depend on a number of test conditions including the sensor employed, acquisition protocol used, subject disposition, number of subjects, number of biometric samples per subject, demographic profile of test subjects, subject habituation, time lapse between data acquisition, etc.

Biometric Trait	**Test**	**Test Conditions**	**False Reject Rate**	**False Accept Rate**
Fingerprint	FVC 2004 [45]	Exaggerated skin distortion, rotation	2%	2%
Fingerprint	FpVTE 2003 [76]	US Government operational data	0.1%	1%
Face	FRVT 2002 [60]	Varied lighting, outdoor/indoor, time	10%	1%
Voice	NIST 2004 [62]	Text independent, multi-lingual	5-10%	2-5%
Iris	ITIRT 2005 [28]	Indoor environment, multiple visits	0.99%	0.94%

[5] http://news.bbc.co.uk/2/hi/uk_news/politics/3693375.stm

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

5 Multibiometric Systems

Evidence accumulation and information fusion is an active area of research in several different fields including weather forecasting [57], UAV coordination [74], object tracking [3], robot navigation [1], land-mine detection [24], etc. The notion of problem solving by combining the decisions rendered by multiple "experts" (or algorithms) in a cooperative framework has received substantial attention in the literature. Indeed, information fusion has a long history and the theory of multiple classifier systems (MCS) has been rigorously studied [40, 37, 77, 21].

In the realm of biometrics, the consolidation of evidence presented by multiple biometric sources is an effective way of enhancing the recognition accuracy of an authentication system. For example, the Integrated Automated Fingerprint Identification System (IAFIS) maintained by the FBI integrates the information presented by multiple fingers to determine a match in the master file. Some of the earliest *multimodal*[6] biometric systems reported in the literature combined the face (image/video) and voice (audio) traits of individuals [10, 6]. Besides enhancing matching accuracy, the other advantages of multibiometric systems over traditional unibiometric systems are enumerated below [68].

1. Multibiometric systems address the issue of non-universality (i.e., limited population coverage) encountered by unibiometric systems. If a subject's dry finger prevents her from successfully enrolling into a fingerprint system, then the availability of another biometric trait, say iris, can aid in the inclusion of the individual in the biometric system. A certain degree of flexibility is achieved when a user enrolls into the system using several different traits (e.g., face, voice, fingerprint, iris, hand) while only a subset of these traits (e.g., face and voice) is requested during authentication based on the nature of the application under consideration and the convenience of the user.
2. Multibiometric systems can facilitate the filtering or indexing of large-scale biometric databases. For example, in a bimodal system consisting of face and fingerprint, the face feature set may be used to compute an index value for extracting a candidate list of potential identities from a large database of subjects. The fingerprint modality can then determine the final identity from this limited candidate list.
3. It becomes increasingly difficult (if not impossible) for an impostor to spoof multiple biometric traits of a legitimately enrolled individual. If each sub-system indicates the probability that a particular trait is a 'spoof', then appropriate fusion schemes can be employed to determine if the user, in fact, is an impostor. Furthermore, by asking the user to present a random subset of traits at the point of acquisition, a multibiometric system facilitates a challenge-response type of mechanism, thereby ensuring that the system is interacting with a *live* user. Note that a challenge-response mechanism can be initiated in unibiometric systems also (e.g., system prompts "Please say 1-2-5-7", "Blink twice and move your eyes to the right", "Change your facial expression by smiling", etc.).
4. Multibiometric systems also effectively address the problem of noisy data. When the biometric signal acquired from a single trait is corrupted with noise, the availability of other (less noisy) traits may aid in the reliable determination of identity. Some systems take into account the *quality* of the individual biometric signals during the fusion process. This is especially important when recognition has to take place in adverse conditions where certain biometric traits cannot be reliably extracted. For example, in the presence of ambient acoustic noise, when an individual's voice characteristics cannot be accurately measured, the facial characteristics may be used by the multibiometric system to perform authentication. Estimating the quality of the acquired data is in itself a challenging problem but, when appropriately done, can reap significant benefits in a multibiometric system.
5. These systems also help in the *continuous* monitoring or tracking of an individual in situations when a single trait is not sufficient. Consider a biometric system that uses a 2D camera to procure the face and gait information of a person walking down a crowded aisle. Depending upon the distance and pose of the subject with respect to the camera, both these characteristics may or may not be simultaneously available. Therefore, either (or both) of these traits can be used depending upon the location of the individual with respect to the acquisition system thereby permitting the continuous monitoring of the individual.
6. A multibiometric system may also be viewed as a fault tolerant system which continues to operate even when certain biometric sources become unreliable due to sensor or software malfunction, or deliberate user manipulation. The notion of fault tolerance is especially useful in large-scale authentication systems involving a large number of subjects (such as a border control application).

[6] See Section 6 for a description of the terminology.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

The design of a multibiometric system is defined by several different factors including (a) the Human Computer Interface (HCI) used to acquire biometric information from an individual[7]; (b) the tradeoff between the additional cost incurred in introducing multiple biometric sources and the perceived improvement in matching accuracy; (c) the sources of biometric information used to provide evidence; (d) the level of fusion, i.e, the type of information to be fused; and (e) the fusion methodology adopted.

6 Taxonomy of Multibiometric Systems

A multibiometric system relies on the evidence presented by multiple sources of biometric information. Based on the nature of these sources, a multibiometric system can be classified into one of the following six categories [68]: multi-sensor, multi-algorithm, multi-instance, multi-sample, multimodal and hybrid.

1. Multi-sensor systems:
   Multi-sensor systems employ multiple sensors to capture a single biometric trait of an individual. For example, a face recognition system may deploy multiple 2D cameras to acquire the face image of a subject [41]; an infrared sensor may be used in conjunction with a visible-light sensor to acquire the subsurface information of a person's face [38, 9, 71]; a multispectral camera may be used to acquire images of the iris, face or finger [69, 58]; or an optical as well as a capacitive sensor may be used to image the fingerprint of a subject [47]. The use of multiple sensors, in some instances, can result in the acquisition of complementary information that can enhance the recognition ability of the system. For example, based on the nature of illumination due to ambient lighting, the infrared and visible-light images of a person's face can present different levels of information resulting in enhanced matching accuracy. Similarly, the performance of a 2D face matching system can be improved by utilizing the shape information presented by 3D range images.
2. Multi-algorithm systems:
   In some cases, invoking multiple feature extraction and/or matching algorithms on the same biometric data can result in improved matching performance. Multi-algorithm systems consolidate the output of multiple feature extraction algorithms, or that of multiple matchers operating on the same feature set. These systems do not necessitate the deployment of new sensors and, hence, are cost-effective compared to other types of multibiometric systems. But on the other hand, the introduction of new feature extraction and matching modules can increase the computational complexity of these systems. Ross et al. [67] describe a fingerprint recognition system that utilizes minutiae as well as texture information to represent and match fingerprint images. The inclusion of the texture-based algorithm introduces additional processing time associated with the application of Gabor filters on the input fingerprint image. However, the performance of the hybrid matcher is shown to exceed that of the individual matchers. Lu et al. [44] discuss a face recognition system that combines three different feature extraction schemes (Principal Component Analysis (PCA), Independent Component Analysis (ICA) and Linear Discriminant Analysis (LDA)). The authors postulate that the use of different feature sets makes the system robust to a variety of intra-class variations normally associated with the face biometric. Experimental results indicate that combining multiple face classifiers can enhance the identification rate of the biometric system.
3. Multi-instance systems:
   These systems use multiple instances of the same body trait and have also been referred to as multi-unit systems in the literature. For example, the left and right index fingers, or the left and right irises of an individual, may be used to verify an individual's identity [61, 34]. The US-VISIT border security program presently uses the left- and right-index fingers of visitors to validate their travel documents at the port of entry. FBI's IAFIS combines the evidence of all ten fingers to determine a matching identity in the database. These systems can be cost-effective if a single sensor is used to acquire the multi-unit data in a sequential fashion (e.g., US-VISIT). However, in some instances, it may be desirable to obtain the multi-unit data simultaneously (e.g., IAFIS) thereby demanding the design of an effective (and possibly more expensive) acquisition device.
4. Multi-sample systems:
   A single sensor may be used to acquire multiple samples of the same biometric trait in order to account for the variations that can occur in the trait, or to obtain a more complete representation of the underlying trait. A face system, for example, may capture (and store) the frontal profile of a person's face along with the left and right profiles in order to account for variations in the facial pose. Similarly, a fingerprint system equipped with

[7] This is an important consideration for multi-modal and multi-unit systems (see Section 6).

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

a small size sensor may acquire multiple dab prints of an individual's finger in order to obtain images of various regions of the fingerprint. A mosaicing scheme may then be used to stitch the multiple impressions and create a composite image. One of the key issues in a multi-sample system is determining the *number* of samples that have to be acquired from an individual. It is important that the procured samples represent the *variability* as well as the *typicality* of the individual's biometric data. To this end, the desired relationship between the samples has to be established before-hand in order to optimize the benefits of the integration strategy. For example, a face recognition system utilizing both the frontal- and side-profile images of an individual may stipulate that the side-profile image should be a three-quarter view of the face [26, 56]. Alternately, given a set of biometric samples, the system should be able to automatically select the "optimal" subset that would best represent the individual's variability. Uludag et al. [73] discuss two such schemes in the context of fingerprint recognition. The first method, called `DEND`, employs a clustering strategy to choose a template set that best represents the intra-class variations, while the second method, called `MDIST`, selects templates that exhibit maximum similarity with the rest of the impressions.

5. Multimodal systems:
   Multimodal systems establish identity based on the evidence of multiple biometric traits. For example, some of the earliest multimodal biometric systems utilized face and voice features to establish the identity of an individual [6, 11, 2]. Physically uncorrelated traits (e.g., fingerprint and iris) are expected to result in better *improvement* in performance than correlated traits (e.g., voice and lip movement). The cost of deploying these systems is substantially more due to the requirement of new sensors and, consequently, the development of appropriate user interfaces. The identification accuracy can be significantly improved by utilizing an increasing number of traits although the *curse-of-dimensionality* phenomenon would impose a bound on this number. The curse-of-dimensionality limits the number of attributes (or features) used in a pattern classification system when only a small number of training samples is available [16]. The number of traits used in a specific application will also be restricted by practical considerations such as the cost of deployment, enrollment time, throughput time, expected error rate, user habituation issues, etc.
6. Hybrid systems:
   Chang et al. [8] use the term *hybrid* to describe systems that integrate a subset of the five scenarios discussed above. For example, Brunelli et al. [6] discuss an arrangement in which two speaker recognition algorithms are combined with three face recognition algorithms at the match score and rank levels via a HyperBF network. Thus, the system is multi-algorithmic as well as multimodal in its design. Similarly, the NIST BSSR1 dataset [53] has match scores pertaining to two different face matchers operating on the frontal face image of an individual (multi-algorithm), and a fingerprint matcher operating on the left- and right-index fingers of the same individual (multi-instance).

It is also possible to combine biometric information with non-biometric entities such as tokens in order to enhance the matching performance. For example, [35] discuss a dual factor authenticator that combines a pseudo random number (present in a token) with a facial feature set in order to produce a set of user-specific compact codes known as BioCode. The pseudo random number and the facial feature sets are fixed in length and an iterated inner product is used to generate the BioCode. When an individual's biometric information is suspected to be compromised, then the token containing the random data is replaced, thereby revoking the previous authenticator. The use of biometric and non-biometric authenticators in tandem is a powerful way of enhancing security. However, some of the inconveniences associated with traditional authenticators remain (such as "Where did I leave my token?").

Another category of multibiometric systems combine primary biometric identifiers (such as face and fingerprint) with soft biometric attributes (such as gender, height, weight, eye color, etc.). Soft biometric traits cannot be used to distinguish individuals reliably since the same attribute is likely to be shared by several different people in the target population. However, when used in conjunction with primary biometric traits, the performance of the authentication system can be significantly enhanced [30]. Soft biometric attributes also help in filtering (or indexing) large biometric databases by limiting the number of entries to be searched in the database. For example, if it is determined (automatically or manually) that the subject is an "Asian Male", then the system can constrain its search to only those identities in the database labeled with these attributes. Alternately, soft biometric traits can be used in surveillance applications to decide if at all primary biometric information has to be acquired from a certain individual. Automated techniques to estimate soft biometric characteristics is an ongoing area of research and is likely to benefit law enforcement and border control biometric applications.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

7 Levels of fusion

In a biometric system, the amount of available information gets compressed as one progresses along the various modules of the system. Based on the type of information available in a certain module, different levels of fusion can be defined. Sanderson and Paliwal [70] categorize the various levels of fusion into two broad categories: pre-classification or fusion *before* matching and post-classification or fusion *after* matching (see Figure 4). Such a categorization is necessary since the amount of information available for fusion reduces drastically once the matcher has been invoked. Pre-classification fusion schemes typically require the development of new matching techniques (since the matchers used by the individual sources may no longer be relevant) thereby introducing additional challenges. Pre-classification schemes include fusion at the sensor (or raw data) and the feature levels while post-classification schemes include fusion at the match score, rank and decision levels.

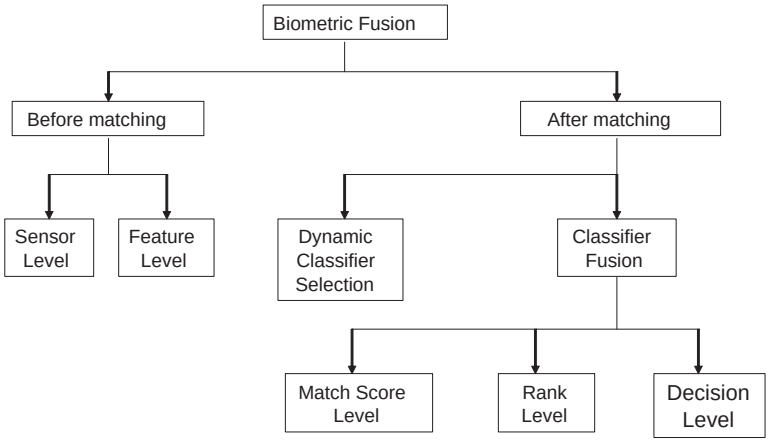

Fig. 4. Fusion can be accomplished at various levels in a biometric system. Most multibiometric systems fuse information at the match score level or the decision level. More recently researchers have begun to fuse information at the sensor and feature levels. In biometric systems operating in the identification mode, fusion can be done at the rank level.

8 Summary

Rapid advancements in the field of communications, computer networking and transportation, coupled with heightened concerns about identity fraud and national security, has resulted in a pronounced need for reliable and efficient identity management schemes in a myriad of applications. The process of identity management in the context of a specific application involves the creation, maintenance and obliteration of identities while ensuring that an impostor does not fraudulently gain privileges associated with a legitimately enrolled individual. Traditional authentication techniques based on passwords and tokens are limited in their ability to address issues such as negative recognition and non-repudiation. The advent of biometrics has served to address some of the shortcomings of traditional authentication methods. Biometric systems use the physical and behavioral characteristics of an individual such as fingerprint, face, hand geometry, iris, gait and voice to establish identity. A broad spectrum of establishments can engage the services of a biometric system including travel and transportation, financial institutions, health care, law enforcement agencies and various government sectors.

The deployment of biometrics in civilian and government applications has raised questions related to the privacy accorded to an enrolled individual [14]. Specifically, questions such as (i) "Will biometric data be used to track people covertly thereby violating their right to privacy?", (ii) "Can the medical condition of a person be surreptitiously elicited from the raw biometric data?", (iii) "Will the acquired biometric data be used only for the intended purpose, or will it be used for previously unexpressed functions, hence resulting in *functionality creep*?", (iv) "Will various biometric databases be linked in order to deduce an individual's social and financial profile?", and (v) "What are the consequences of compromising a user's biometric data?", have advocated societal concerns about the use of biometric

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

solutions in large-scale applications. The promotion of Privacy-Enhancing Technologies (PETs) can assuage some of the legitimate concerns associated with biometric-enabled technology [65, 36]. For example, the use of personal smart cards to store and process the biometric template of an individual can mitigate public concerns related to placing biometric information in a centralized database. Apart from technological solutions to address privacy concerns, government regulations are also required in order to prevent the inappropriate transmission, exchange and processing of biometric data.

The matching performance of a biometric system is affected by several factors including noisy data, large intra-class variations, and improper user interaction [4]. There is an implicit upper bound on the matching accuracy of any biometric system. Jain et al. [31] suggest three primary reasons for this inherent constraint:

1. **Information limitation**: The magnitude of discriminatory information available in a biometric trait is naturally restricted. For example, hand geometry measurements can distinguish fewer identities than, say, fingerprints [59].
2. **Representation limitation**: The ideal representation scheme for a particular biometric trait should be designed to retain all invariant and discriminatory information in the sensed measurements. Practical feature extraction systems, typically based on simplistic models of biometric data, fail to capture the richness of information in a realistic biometric input resulting in the inclusion of redundant or spurious features, and the exclusion of salient features. Consequently, a significant fraction of legitimate feature space cannot be handled by the biometric system resulting in authentication errors (FAR and FRR).
3. **Matcher limitation**: Given a particular representation scheme, the design of an ideal matcher should perfectly model the invariant relationship between different patterns (i.e, biometric samples) originating from the same class (i.e, identity). In practice, however, a matcher may not correctly model the invariance (e.g., due to non-availability of sufficient number of training samples) resulting in poor matcher accuracy.

Multibiometrics is expected to alleviate some of the limitations of unibiometric systems by consolidating the evidence presented by multiple biometric sources. This integration of evidence is known as information fusion and, if appropriately done, can enhance the matching accuracy of a recognition system. Thus, a properly designed multi-biometric system can improve matching accuracy, increase population coverage and deter spoofing activities. With biometrics already being chosen to deliver crucial societal functions, it is only a matter of time before multibiometric systems begin to impact the way we perform identity management in the 21st century.

References

1. M. A. Abidi and R. C. Gonzalez. *Data Fusion in Robotics and Machine Intelligence*. Academic Press, New York, USA, 1992.
2. E. S. Bigun, J. Bigun, B. Duc, and S. Fischer. Expert Conciliation for Multimodal Person Authentication Systems using Bayesian Statistics. In *First International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA)*, pages 291–300, Crans-Montana, Switzerland, March 1997.
3. R. S. Blum and Z. Liu, editors. *Multi-Sensor Image Fusion and Its Applications*. CRC Press, Taylor and Francis Group, Florida, USA, 2006.
4. R. Bolle, J. Connell, S. Pankanti, N. Ratha, and A. Senior. *Guide to Biometrics*. Springer, 2003.
5. R. Bolle, J. Connell, S. Pankanti, N. Ratha, and A. Senior. The Relationship Between the ROC Curve and the CMC. In *Proceedings of Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID)*, pages 15–20, Buffalo, USA, October 2005.
6. R. Brunelli and D. Falavigna. Person Identification Using Multiple Cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(10):955–966, October 1995.
7. J. P. Campbell. Speaker Recognition: a Tutorial. *Proceedings of the IEEE*, 85(9):1437–1462, September 1997.
8. K. I. Chang, K. W. Bowyer, and P. J. Flynn. An Evaluation of Multimodal 2D+3D Face Biometrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):619–624, April 2005.
9. X. Chen, P. J. Flynn, and K. W. Bowyer. IR and Visible Light Face Recognition. *Computer Vision and Image Understanding*, 99(3):332–358, September 2005.
10. C. C. Chibelushi, F. Deravi, and J. S. Mason. Voice and Facial Image Integration for Speaker Recognition. In R. I. Damper, W. Hall, and J. W. Richards, editors, *Multimedia Technologies and Future Applications*, pages 155–161. Pentech Press, London, 1994.
11. C. C. Chibelushi, J. S. D. Mason, and F. Deravi. Feature-level Data Fusion for Bimodal Person Recognition. In *Proceedings of the Sixth International Conference on Image Processing and Its Applications*, volume 1, pages 399–403, Dublin, Ireland, July 1997.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

12. J. Daugman. Recognizing Persons by their Iris Patterns. In A. K. Jain, R. Bolle, and S. Pankanti, editors, *Biometrics: Personal Identification in Networked Society*, pages 103–122. Kluwer Academic Publishers, London, UK, 1999.
13. J. Daugman. How Iris Recognition Works? *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, 2004.
14. S. Davies. Touching Big Brother: How Biometric Technology Will Fuse Flesh and Machine. *Information Technology and People*, 7(4), 1994.
15. G. Doddington, W. Liggett, A. Martin, M. Przybocki, and D. Reynolds. Sheep, Goats, Lambs and Wolves: A Statistical Analysis of Speaker Performance in the NIST 1998 Speaker Recognition Evaluation. In *CD-ROM Proceedings of the Fifth International Conference on Spoken Language Processing (ICSLP)*, Sydney, Australia, November/December 1998.
16. R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. John Wiley & Sons, 2001.
17. N. Duta, A. K. Jain, and K. V. Mardia. Matching of Palmprints. *Pattern Recognition*, 23(4):477–485, February 2002.
18. J. Egan. *Signal Detection Theory and ROC Analysis*. Academic Press, New York, 1975.
19. A. Eriksson and P. Wretling. How Flexible is the Human Voice? A Case Study of Mimicry. In *Proceedings of the European Conference on Speech Technology*, pages 1043–1046, Rhodes, 1997.
20. C. L. Fancourt, L. Bogoni, K. J. Hanna, Y. Guo, R. P. Wildes, N. Takahashi, and U. Jain. Iris Recognition at a Distance. In *Fifth International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA)*, pages 1–13, Rye Brook, USA, July 2005.
21. J. Ghosh. Multiclassifier Systems: Back to the Future. In *Proceedings of Third International Workshop on Multiple Classifier Systems*, pages 1–15, Cagliari, Italy, June 2002.
22. M. Golfarelli, D. Maio, and D. Maltoni. On the Error-Reject Tradeoff in Biometric Verification Systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):786–796, July 1997.
23. P. Grother and P. J. Phillips. Models of Large Population Recognition Performance. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 68–75, Washington D.C., USA, June/July 2004.
24. A. H. Gunatilaka and B. A. Baertlein. Feature-level and Decision-level Fusion of Noncoincidentally Sampled Sensors for Land Mine Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):577–589, June 2001.
25. W. R. Harrison. *Suspect Documents, their Scientific Examination*. Nelson-Hall Publishers, 1981.
26. H. Hill, P. G. Schyns, and S. Akamatsu. Information and Viewpoint Dependence in Face Recognition. *Cognition*, 62(2):201–222, February 1997.
27. L. Hong, A. K. Jain, and S. Pankanti. Can Multibiometrics Improve Performance? In *Proceedings of IEEE Workshop on Automatic Identification Advanced Technologies (AutoID)*, pages 59–64, New Jersey, USA, October 1999.
28. International Biometric Group. Independent Testing of Iris Recognition Technology: Final Report. Available at `http://www.biometricgroup.com/reports/public/ITIRT.html`, May 2005.
29. A. K. Jain, R. Bolle, and S. Pankanti, editors. *Biometrics: Personal Identification in Networked Society*. Kluwer Academic Publishers, 1999.
30. A. K. Jain, K. Nandakumar, X. Lu, and U. Park. Integrating Faces, Fingerprints and Soft Biometric Traits for User Recognition. In *Proceedings of ECCV International Workshop on Biometric Authentication (BioAW)*, volume LNCS 3087, pages 259–269, Prague, Czech Republic, May 2004. Springer.
31. A. K. Jain, S. Pankanti, S. Prabhakar, L. Hong, and A. Ross. Biometrics: A Grand Challenge. In *Proceedings of International Conference on Pattern Recognition (ICPR)*, volume 2, pages 935–942, Cambridge, UK, August 2004.
32. A. K. Jain and A. Ross. Multibiometric Systems. *Communications of the ACM, Special Issue on Multimodal Interfaces*, 47(1):34–40, January 2004.
33. A. K. Jain, A. Ross, and S. Prabhakar. An Introduction to Biometric Recognition. *IEEE Transactions on Circuits and Systems for Video Technology, Special Issue on Image- and Video-Based Biometrics*, 14(1):4–20, January 2004.
34. J. Jang, K. R. Park, J. Son, and Y. Lee. Multi-unit Iris Recognition System by Image Check Algorithm. In *Proceedings of International Conference on Biometric Authentication (ICBA)*, pages 450–457, Hong Kong, July 2004.
35. A. T. B. Jin, D. N. C. Ling, and A. Goh. An Integrated Dual Factor Authenticator Based On The Face Data And Tokenised Random Number. In *First International Conference on Biometric Authentication*, pages 117–123, Hong Kong, China, July 2004.
36. S. Kenny and J. J. Borking. The Value of Privacy Engineering. *The Journal of Information, Law and Technology (JILT)*, 7(1), 2002.
37. J. Kittler, M. Hatef, R. P. Duin, and J. G. Matas. On Combining Classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):226–239, March 1998.
38. A. Kong, J. Heo, B. Abidi, J. Paik, and M. Abidi. Recent Advances in Visual and Infrared Face Recognition - A Review. *Computer Vision and Image Understanding*, 97(1):103–135, January 2005.
39. A. Kumar, D. C. M. Wong, H. C. Shen, and A. K. Jain. Personal Verification Using Palmprint and Hand Geometry Biometric. In *Fourth International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA)*, pages 668–678, Guildford, UK, June 2003.
40. L. I. Kuncheva. *Combining Pattern Classifiers - Methods and Algorithms*. Wiley, 2004.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

41. J. Lee, B. Moghaddam, H. Pfister, and R. Machiraju. Finding Optimal Views for 3D Face Shape Modeling. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 31–36, Seoul, Korea, May 2004.
42. L. Lee, T. Berger, and E. Aviczer. Reliable On-Line Human Signature Verification Systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(6):643–647, June 1996.
43. S. Z. Li and Anil K. Jain, editors. *Handbook of Face Recognition*. Springer-Verlag, 2005.
44. X. Lu, Y. Wang, and A. K. Jain. Combining Classifiers for Face Recognition. In *IEEE International Conference on Multimedia and Expo (ICME)*, volume 3, pages 13–16, Baltimore, USA, July 2003.
45. D. Maio, D. Maltoni, R. Cappelli, J. L. Wayman, and A. K. Jain. FVC2004: Third Fingerprint Verification Competition. In *Proceedings of International Conference on Biometric Authentication (ICBA)*, pages 1–7, Hong Kong, China, July 2004.
46. D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar. *Handbook of Fingerprint Recognition*. Springer-Verlag, 2003.
47. G. L. Marcialis and F. Roli. Fingerprint Verification by Fusion of Optical and Capacitive Sensors. *Pattern Recognition Letters*, 25(11):1315–1322, August 2004.
48. A. Martin, G. Doddington, T. Kam, M. Ordowski, and M. Przybocki. The DET Curve in Assessment of Detection Task Performance. In *Proceedings of the Fifth European Conference on Speech Communication and Technology*, volume 4, pages 1895–1898, Rhodes, Greece, September 1997.
49. T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino. Impact of Artificial Gummy Fingers on Fingerprint Systems. In *Optical Security and Counterfeit Deterrence Techniques IV, Proceedings of SPIE*, volume 4677, pages 275–289, San Jose, USA, January 2002.
50. F. Monrose and A. Rubin. Authentication Via Keystroke Dynamics. In *Proceedings of Fourth ACM Conference on Computer and Communications Security*, pages 48–56, Zurich, Switzerland, April 1997.
51. H. Moon and P. J. Phillips. Computational and Performance Aspects of PCA-based Face Recognition Algorithms. *Perception*, 30(5):303–321, 2001.
52. V. S. Nalwa. Automatic On-Line Signature Verification. *Proceedings of the IEEE*, 85(2):215–239, February 1997.
53. National Institute of Standards and Technology. NIST Biometric Scores Set. Available at `http://http://www.itl.nist.gov/iad/894.03/biometricscores`, 2004.
54. M. Negin, T. A. Chmielewski, M. Salganicoff, T. A. Camus, U. M. C. von Seelan, P. L. Venetianer, and G. G. Zhang. An Iris Biometric System for Public and Personal Use. *IEEE Computer*, 33(2):70–75, February 2000.
55. M. S. Nixon, J. N. Carter, D. Cunado, P. S. Huang, and S. V. Stevenage. Automatic Gait Recognition. In A. K. Jain, R. Bolle, and S. Pankanti, editors, *Biometrics: Personal Identification in Networked Society*, pages 231–249. Kluwer Academic Publishers, London, UK, 1999.
56. A. O'Toole, H. Bulthoff, N. Troje, and T. Vetter. Face Recognition across Large Viewpoint Changes. In *Proceedings of the International Workshop on Automatic Face- and Gesture-Recognition (IWAFGR)*, pages 326–331, Zurich, Switzerland, June 1995.
57. T. N. Palmer. Predicting Uncertainty in Forecasts of Weather and Climate. *Reports on Progress in Physics*, 63:71–116, 2000.
58. Z. Pan, G. Healey, M. Prasad, and B. Tromberg. Face Recognition in Hyperspectral Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1552–1560, December 2003.
59. S. Pankanti, S. Prabhakar, and A. K. Jain. On the Individuality of Fingerprints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1010–1025, 2002.
60. P. J. Phillips, P. Grother, R. J. Micheals, D. M. Blackburn, E. Tabassi, and J. M. Bone. FRVT2002: Overview and Summary. Available at `http://www.frvt.org/FRVT2002`, March 2003.
61. S. Prabhakar and A. K. Jain. Decision-level Fusion in Fingerprint Verification. Technical Report MSU-CSE-00-24, Michigan State University, October 2000.
62. M. Przybocki and A. Martin. NIST Speaker Recognition Evaluation Chronicles. In *Odyssey: The Speaker and Language Recognition Workshop*, pages 12–22, Toledo, Spain, May 2004.
63. T. Putte and J. Keuning. Biometrical Fingerprint Recognition: Don't Get Your Fingers Burned. In *Proceedings of IFIP TC8/WG8.8 Fourth Working Conference on Smart Card Research and Advanced Applications*, pages 289–303, 2000.
64. N. K. Ratha, J. H. Connell, and R. M. Bolle. An Analysis of Minutiae Matching Strength. In *Proceedings of Third International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA)*, pages 223–228, Halmstad, Sweden, June 2001.
65. M. Rejman-Greene. Privacy Issues in the Application of Biometrics: A European Perspective. In J. L. Wayman, A. K. Jain, D. Maltoni, and D. Maio, editors, *Biometric Systems: Technology, Design and Performance Evaluation*, pages 335–359. Springer, 2005.
66. E. Rood and A. K. Jain. Biometric research agenda: Report of the NSF workshop. In *Workshop for a Biometric Research Agenda*, Morgantown, WV, July 2003.
67. A. Ross, A. K. Jain, and J. Reisman. A Hybrid Fingerprint Matcher. *Pattern Recognition*, 36(7):1661–1673, July 2003.
68. A. Ross, K. Nandakumar, and A. K. Jain. *Handbook of Multibiometrics*. Springer, New York, USA, 1st edition, 2006.
69. R. K. Rowe and K. A. Nixon. Fingerprint Enhancement Using a Multispectral Sensor. In *Proceedings of SPIE Conference on Biometric Technology for Human Identification II*, volume 5779, pages 81–93, March 2005.

Appeared in Annals of Telecommunications, Vol. 62, No. 1/2, pp. 11-35, Jan/Feb 2007

70. C. Sanderson and K. K. Paliwal. Information Fusion and Person Verification Using Speech and Face Information. Research Paper IDIAP-RR 02-33, IDIAP, September 2002.
71. D. A. Socolinsky, A. Selinger, and J. D. Neuheisel. Face Recognition with Visible and Thermal Infrared Imagery. *Computer Vision and Image Understanding*, 91(1-2):72–114, July-August 2003.
72. J. A. Swets, W. P. Tanner, and T. G. Birdsall. Decision Processes in Perception. *Psychological Review*, 68(5):301–340, 1961.
73. U. Uludag, A. Ross, and A. K. Jain. Biometric Template Selection and Update: A Case Study in Fingerprints. *Pattern Recognition*, 37(7):1533–1542, July 2004.
74. G. Vachtsevanos, L. Tang, and J. Reimann. An Intelligent Approach to Coordinated Control of Multiple Unmanned Aerial Vehicles. In *60th Annual Forum of the American Helicopter Society*, Baltimore, USA, June 2004.
75. H. P. Wasserman. *Ethnic Pigmentation*. Elsevier, New York, USA, 1974.
76. C. Wilson, A. R. Hicklin, M. Bone, H. Korves, P. Grother, B. Ulery, R. Micheals, M. Zoepfl, S. Otto, and C. Watson. Fingerprint Vendor Technology Evaluation 2003: Summary of Results and Analysis Report. NIST Technical Report NISTIR 7123, National Institute of Standards and Technology, June 2004.
77. L. Xu, A. Krzyzak, and C. Y. Suen. Methods for Combining Multiple Classifiers and their Applications to Handwriting Recognition. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(3):418–435, 1992.
78. D. Zhang, A. W.-K. Kong, J. You, and M. Wong. Online Palmprint Identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1041–1050, 2003.
79. R. Zunkel. Hand Geometry Based Authentication. In A. K. Jain, R. Bolle, and S. Pankanti, editors, *Biometrics: Personal Identification in Networked Society*, pages 87–102. Kluwer Academic Publishers, London, UK, 1999.