

1

Introduction to Behavioral Biometrics

1.1 Introduction

This book fills several roles in the biometrics literature. It is intended to serve as a reference source for case studies in behavioral biometrics. There are a number of very useful texts on the topic of biometrics, but they tend to focus on physiological biometrics, or they focus at a generic level, covering the full spectrum, becoming overly general. There are texts that focus on the algorithmic approaches deployed in biometrics, and as such serve as a source of machine learning algorithms. To date, there is no text that is solely dedicated to the topic of behavioral biometrics. This book serves to provide a number of case studies of major implementations within the field of behavioral biometrics. Though not as informative as the actual published work, the case studies are comprehensive and provide the user with a strong sense of the approaches employed in the various subdomains of behavioral biometrics. The intended audience is students, at the advanced undergraduate and postgraduate levels, and researchers wishing to explore this fascinating research topic. In addition, this text will serve as a reference for system integrators, CIOs, and related professionals who are charged with implementing security features at their organization. The reader will be directed to appropriate sources when detailed implementation issues are concerned, especially those involving specific machine learning algorithms. A single text of this size cannot cover both the domain of behavioral biometrics *and* the machine learning algorithms they employ.

Biometrics in the context presented in this book is concerned with a scientific approach to user verification and/or identification. The focus will be on *behavioral* biometrics – the verification and/or identification of individuals based on the way they provide information to the authentication system. For instance, individuals could be required to provide a signature, enunciate a particular phrase, or enter a secret code through an input device in order to provide evidence of their identity. Note that there is an implicit simplicity to behavioral biometrics in that typically, no special machinery/hardware is required for the authentication/identification process other than the computer (or ATM) device itself. In addition, the approaches prevalent in this domain are very familiar to us – practically everyone has provided a signature to verify their identity, and we have one or more passwords for logging into computer

systems. We are simply used to providing proof of identity in these fashions in certain circumstances. These two factors provide the foundation for the behavioral approach to biometrics. These modes of identification are substantially different from the other classes of biometrics: physiological and token-based biometrics. For instance, what is termed physiological (or biological) biometrics requires that we present some aspect of our physicality in order to be identified. Typical instances of physiological biometrics include iris scans, retina scans, and fingerprints. Lastly, token-based biometric systems require the possession of some object such as a bank or identity card. Each class of biometrics is designed to provide an efficient and accurate method of verifying the identity (authentication) and/or the identification of an individual.

1.2 Types of Behavioral Biometrics

There are a variety of subdivisions within the behavioral biometrics domain. Each subdivision has its own characteristics in terms of ease of use, deployability, user acceptance, and quality of the identification/verification task. In order of presentation in this text, the following subdivisions can be identified as

- **Voice Recognition:**

in which users are requested to enunciate text as a means of identifying themselves. Voice can be employed for either speaker identification or speaker authentication. With respect to speaker identification, a person enunciates text, and the speech patterns are analyzed to determine the identity of the speaker. In the literature, this is referred to as speaker-independent recognition. This mode poses several interesting issues, such as what happens if the speaker is not contained within the database of speakers? As in all major forms of biometrics, any individual wishing to utilize the biometric device must, at some stage, introduce themselves to the system, typically in the form of an enrollment process. One of the principal tasks of the enrollment process is to register the person as a potential user of the biometric system (enrollment will be discussed further later in this chapter). In a speaker-independent system, the user's voice pattern is analyzed and compared to all other voice samples in the user database. There are a number of ways this comparison is made, and specific details are provided via case studies in the appropriate chapters (Chapter 2 for voice recognition). The closest match to the particular voice data presented for identification becomes the presumed identity of the speaker. There are three possible outcomes: i) The speaker is correctly identified; ii) the speaker is incorrectly identified as another speaker; or iii) the speaker is not identified as being a member of the system. Clearly, we would like to avoid the last two possibilities, which reflect the false acceptance rate (FAR) (type II error) and the false rejection rate (FRR) (type I error) as much as possible. When speakers attempt an authentication task, the speakers have provided some evidence of their identity, and the purpose of the voice recognition process is to verify that these persons have a legitimate claim to that identity. The result of this approach is a binary decision: either you are verified as the claimed identity or you are not.

The other major division within voice recognition biometrics is whether the enunciated text is fixed or free, that is, do the users enunciate a specific phrase (text dependent), or are

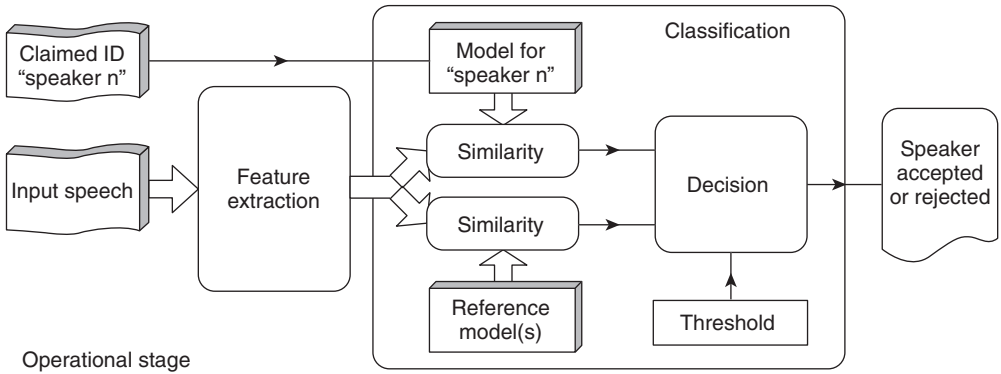


Figure 1.1 An example of a voice recognition processing system (Source: Ganchev, 2005)

they allowed to enunciate any phrase (text-independent)? The speaker-dependent version is easier from a matching perspective, in that the spoken text is directly matched to the information stored in the database. The text-independent approach allows speakers to enunciate any speech they wish to. This approach requires a model of each speaker, which is certainly more computationally expensive than the text-dependent approach. These and other related issues will be discussed further in the next chapter (Figure 1.1).

• Signature Verification:

where users are required to present handwritten text for authentication. This is probably the most familiar of all biometrics – though currently not the most prevalent – due to the advent of computer-based passwords. There are two essentially distinct forms of signature-based biometrics: online and off-line. With an online signature verification system, the signature characteristics are extracted as the user writes, and these features are used to immediately authenticate the user. Typically, specialized hardware is required, such as a pressure-sensitive pen (a digital pen) and/or a special writing tablet. These hardware elements are designed to capture the dynamical aspects of writing, such the pen pressure, pen angle, and related information (see Chapter 3 for details). In a remote access approach, where specialized hardware may not be feasible, the online approach is most suitable from a small portable device such as a PDA, where the stylus can be used for writing. The off-line approach utilizes the static features of the signature, such as the length and height of the text, and certain specialized features such as loops (not unlike a fingerprint approach). Typically, the data are acquired through an image of the signature, which may be photocopied or scanned into a computer for subsequent analysis. As in all behavioral biometric approaches, a writing sample must be stored in the authentication database, and the writing sample is compared to the appropriate reference sample before the acceptance/rejection decision to be made. Again, there is the possibility of having text-dependent or text-independent signature verification. The same caveats that apply to voice also apply here – and voice and signature are really very similar technologies – only the mode of communication has changed, which results in a different set of features that can be extracted. An example of an online signature setup is presented in Figure 1.2.



Figure 1.2 An online signature verification system (Source: Interlink Electronics ePad (www.primidi.com/2003/05/31.html))

♦ **Keystroke Dynamics:**

is a behavioral biometric that relies on the *way we type* on a typical keyboard/keypad type device. As a person types, certain attributes are extracted and used to authenticate or identify the typist. Again, we have two principal options: text-dependent and text-independent versions. The most common form of text-dependent systems requires users to enter their login ID and password (or commonly just their password). In the text-independent version, users are allowed to enter any text string they wish. In some implementations, a third option is used where a user is requested to enter a long text string on the order of 500–1500 characters. Users enroll into the system by entering their text either multiple times if the short text-independent system (i.e. password) is employed, or typically once if the system employs a long text string. From this enrollment process, the user's typing style is acquired and stored for subsequent authentication purposes. This approach is well suited for remote access scenarios: no specialized hardware is required and users are used to providing their login credentials. As discussed in more detail in Chapter 4, some of the attributes that are extracted when a person types are the duration of a key press (dwell time) and the time between striking successive keys (digraph if the time is recorded between successive keys). These features, along with several others, are used to build a model of the way a person types. The security enhancement provided by this technology becomes evident if you leave your password written on a sticky notepad tucked inside your desk, which someone happens to find. Without this level of protection, possession of the password is that that is required for a user to access your account. With the addition of a keystroke dynamics-based biometric, it is not sufficient that the password is acquired: the password has to be entered exactly (or at least within certain tolerance limits) the way the enrolled user entered it during enrollment. If not, the login attempt is rejected. An example of the notion of a digraph is depicted in Figure 1.3.

♦ **Graphical Authentication Systems:**

are employed as an alternative to textual-based password systems. There are issues with textual-based passwords regarding the strength, which refers to how easy it would be to guess




Figure 1.3 The combinations of digraphs that can be generated from the character sequence “N” followed by “O”. Note the subscripts “p” and “r” indicate press and release, respectively.

someone’s password, given free access to the computer system on which they are stored (an off-line attack). Studies have indicated that most people make their passwords easy to remember – such as their names, certain memorable places, etc. Generally speaking, the full password space is not utilized when people are allowed to select their own passwords. On a typical PC-type keyboard, there are 95 printable characters, and for a password of eight characters, there are 95^8 (or 6×10^{15}) possible passwords that can be generated. This is a relatively large search space to exhaustively explore, though not impossible in a realistic time frame with today’s modern computing power (and the deployment of a grid of computers). But typically, most users explore a small fraction of this possible password space, and the possibility of a successful off-line attack is very real (see Chapter 4 for some examples). As indicated, the principal reason for the lack of a thorough exploration of password space is the issue of memorability. Here is where graphical passwords take over.

Graphical passwords are composed of a collection of images, each representing an element of the user’s password. The images are presented to the user – who must select the password elements – possibly in a predefined order, but more often than not, order is removed from the equation, depending on the implementation. A key difference between textual- and graphical-based passwords is that in the former, recall is required, and in the latter, recognition is involved. The psychological literature has provided ample evidence that recognition is a much easier task than recall. In addition, it appears that we have an innate ability to remember pictures better than text. These two factors combined provide the rationale for the graphical password-based approach. There are a variety of graphical-based password systems that have been developed, and this interesting approach is discussed in some detail in Chapter 6. An example of a classical approach, dubbed Passfaces™, is presented in Figure 1.4. In this system, the user’s password is a collection of faces (typically four to six), which must be selected in order from a series of decoy face images.

• Mouse Dynamics:

is a biometric approach designed to capture the static and dynamic aspects of using the mouse as a tool for interacting with a user interface, which contains the elements of their password,



Figure 1.4 An example of the Passfaces™ graphical password authentication scheme. Note that on each page of faces, the user is required to select the correct face image; note that in this system, there is an implied order to the selection process (Source: Passfaces website – www.passfaces.com)

typically presented in a graphical fashion. Mouse movement information such as the change in the mouse pointer position over time and space is recorded, providing the basis for determining trajectories and velocity, which can be used to build a reference model of the user. Therefore, mouse dynamics is used in conjunction within a graphical password scenario, though the password may not consist of a collection of images to be identified. Instead, this approach is based on human computer interaction (HCI) features – how one interacts with an application is used to authenticate a user. Provided there is enough entropy in the game – enough possibilities for interacting with it, then one may be able to differentiate users based on this information. Some examples of this approach, which are rather sparsely represented in the literature, are presented in detail in Chapter 6, and an example of a system developed by Ahmed & Traore, (2003) is presented in Figure 1.5.

♦ **Gait as a Biometric:**

relies on the walking pattern of a person. Even the great Shakespeare himself stated that “For that John Mortimer . . . in face, in gait in speech he doth resemble” (Shakespeare, W., King




Figure 1.5 A graph presenting the user profile (solid top line versus a series of imposters based on average speed of mouse movements (Source: Awad et al., 2005)

Henry the Sixth, part 2, ca. 1590-1591). As Shakespeare himself intimated, there are subtle differences in the way a person ambulates. The results of a number of gait-based biometrics indicate that these differences are statistically significant – leading equal error rate (EER) values on the order of 5% or less. There are two principal approaches to gait biometrics: machine-vision and sensor-based approaches. The former is the more traditional approach and is suitable for scenarios where the authentication process must be mass produced, such as at airports. In this scenario, a user can be scanned from a distance relative to the authentication point. Provided the data acquisition can occur quickly, this type of approach may be very attractive. The sensor-based approach (see Figure 1.6 for an example of an accelerometer, the typical sensor used in gait analysis) acquires dynamic data, measuring the acceleration in three orthogonal directions. Sensor-based systems are quite suitable for user authentication – as they are obviously attached to the individual accessing the biometric device. Machine-vision based approaches are more general, and are typically employed for user identification.

The feature space of gait biometrics is not as rich as other technologies. This probably reflects the conditions under which the data are acquired – either a machine-vision approach with issues regarding lighting and other factors that typically degrade the performance of related biometrics such as face recognition. Even under the best of conditions (the gold-standard condition – see Appendix A for details), there are really only three degrees of freedom from which to draw features from. The current trend is to focus on dynamic aspects of walking, and the results tend to be somewhat better than static features when comparing EER values. When deployed in a multimodal approach, gait data, in conjunction with speech biometrics, for instance, tend to produce very low EER values (see Appendix A for details).



Figure 1.6 A photograph of a subject wearing a sensor-based gait device termed an accelerometer. Note that it is capable of measuring acceleration in three different orthogonal directions (Source: Gafurov et al., 2006)

Research continues to find ways to enhance the feature space of gait biometrics, but considering what is currently available, an EER of 3–5% is quite respectable.

♦ **Smile Recognition:**

is a technique that uses high-speed photography, and a zoom lens generates a series of smile maps. These maps are composed of the underlying deformation of the relevant muscles and tiny wrinkles, which move in a characteristic fashion when a person smiles. A collection of directional vectors is produced which form the contours describing the dynamical aspects of smiling. This approach requires further analysis to determine how effective it will be as a behavioral biometrics, as current results are produced from a small study cohort.

♦ **Lip Movement Recognition:**

For the purpose of recognizing individuals, we suggest a lip recognition method using shape similarity when vowels are uttered. In the method, we apply mathematical morphology, in which three kinds of structuring elements such as square, vertical line, and horizontal line are used for deriving pattern spectrum. The shapeness vector is compared to the reference vector to recognize the individual from lip shape. According to experimental results with eight lips that uttered five vowels, it is found that the method successfully recognizes lips with 100% accuracy.

♦ **Odor as a Biometric:**

is an often-overlooked class of behavioral biometrics based on our sense of smell – olfaction-based biometrics. The human olfactory system is capable of detecting a wide range of odorants using a relatively sparse receptor system (see Freeman, 1991 for an excellent review). There are two principal processes involved in olfaction: stimulus reception and identification. There are questions regarding the specificity and sensitivity of the sense of




Figure 1.7 The olfactory biometric scheme, highlighting the sensor array and pattern recognition components (Source: Korotkaya, 2003)

smell. There are a number of professions that rely on a keen sense of smell – wine tasters, perfume experts, and human body recovery are a few examples (Yamazaki et al., 2001, Teo et al., 2002). It would therefore seem reasonable to assume that olfaction does have sufficient capacity to accurately identify a wide range of odors with high sensitivity. The question then shifts to whether or not humans exude sufficiently distinct odors such that we can be discriminated by them. Does the use of deodorant, colognes, and perfumes obfuscate our body odor beyond recognition? Lastly, how do we get a computer to perform olfaction?

The answer to the last question relies on the development of an artificial nose – the ENose (Keller, 1999, Korotkaya, 2003) – depicted in Figure 1.7. It is composed of two modules: a sensor array and a pattern recognition system. The sensor array consists of a collection of sensors (typically 10–20) each designed to react with a particular odorant. The pattern recognition system is used to map the activation pattern of the sensor array to a particular odorant pattern. The sensor array can be designed from a variety of materials, conductor sensors:

- made from metal oxide and polymers;
- piezoelectric sensors;
- metal-oxide-silicon field-effect transistors;
- optical fiber sensors.

Each of these technologies can be deployed as the basis for the sensor aspect of an ENose system (for details, please consult Gardner, 1991, Korotkaya, 2003).

There are a number of pattern recognition systems that can be employed – cluster analysis, neural networks, and related classification algorithms can be employed with success. The current operation of the ENose system is essentially a 1 : 1 correspondence between sensor array number and odorants. Though the human olfactory system contains a great number of receptors (on the order of 1×10^6), they are used in a combinatorial fashion, that is, there is not a 1 : 1 correspondence between an odorant and the activation of a particular receptor. It is a distributed system – and ENose, if it is to succeed at all, must adopt a similar approach. To date, there is not a clear direction in this area; it is really up to the neuroengineers to develop the required technology before it can be adapted to the biometrics domain. Though interesting, this approach will have to therefore wait for further parallel advancements in engineering before it can be considered a truly viable behavioral biometric – especially in a remote access context.

• **Biological Signals as a Behavioral Biometric:**

is a novel approach that relies on the measurement of a variety of biological signals. These include the electrocardiogram (ECG), the electroencephalogram (EEG), and the electrooculogram (EOG) to name a few potential candidates. In the late 1970s, Forsen published a report

that evaluated the largest collection of potential biometric technologies known at the time (Forsen et al., 1977). Included in this impressive list was the deployment of the ECG and EEG – quite prescient for 1977! The basic approach is to extract the signals from the user during the enrollment period, to extract features, and to generate a classifier. When the user then attempts to log in, the particular class of signal is recorded, and a matching score is computed, which determines the decision outcome. This is really no different than any other behavioral biometric (and physiological for that matter) – the novelty here is the data that are acquired. In order to acquire biological signal data, specialized hardware is required. One of the tenets (or at least selling points) of behavioral biometrics is that no specialized hardware is required. It is anticipated that with the current rate of technological advancement, the amount of hardware required will be reduced to acceptable levels.

• ECG as a Behavioral Biometric:

The ECG is simply a recording of the electrical activity associated with the beating of the heart. A series of leads is positioned appropriately over the heart – which picks up the small electrical signals produced by various regions of the heart that generate electricity (i.e. the pacemaker or the sinoatrial node). The recording of the human heartbeat generates a characteristic profile (depicted in Figure 8.3). The question to be addressed is whether there is sufficient variability between individuals such that this signal can form a reliable marker for any particular individual. The data presented in Chapter 8 of this volume indicates that there is plenty of evidence to suggest that *it is* sufficiently discriminating to produce a high degree of classification accuracy (near 100% in some studies). Figure 1.8 presents a typical authentication scheme employing ECG data (taken from Mehta & Lingayat, 2007).




Figure 1.8 A time series recording of electrocardiogram data and some preprocessing results. The x -axis is time and the y -axis represents the signals acquired from each of the 12 leads. The bottom row represents the SVM detection results (Source: Mehta and Lingayat, 2007)

♦ **EEG as a Behavioral Biometric:**

The EEG is a recording from the scalp surface of the electrical activity of a collection of synchronously firing, parallel-oriented neurons. The EEG records the electrical activity of the brain and, as such, is continuously active (even for patients in the locked-in-state condition, resulting from a stroke). Embedded within the ongoing EEG activity are changes that occur in a correlated fashion with particular types of cognitive activities. The activities are typical cognitive functions such as thinking of a name, reading aloud, and listening to music. These signals can be isolated from the underlying background activity through a series of filtering and related techniques, which are discussed in some detail in Chapter 8 of this volume (see the references therein for more details). The goal in this approach is to associate particular electrical signatures that occur within the brain with particular cognitive tasks, such as entering a password to playing a video game.

The data obtained from EEG is sufficiently robust to generate a significant amount of intersubject variability, and many studies have produced statistically significant classification results using “raw” EEG data. In addition, through the process of biofeedback, a type of operant conditioning, people can control, to some degree, the activity of the brain in response to particular tasks (Miller, 1969). This is the essence of the brain–computer interface (BCI) (Figure 1.9) and forms the basis of an exciting area of research that is being applied to biometrics. For instance, users can control the movement of a cursor, type on a virtual keyboard, and related activities.

That this technology can be used as an authentication system is receiving serious research efforts, and the results appear to be quite promising, even at this early stage in the evolution of this technology. Again, there are the issues of the requisite hardware, which, as in the case for ECG technologies, can be expected to diminish with time. An example of a typical BCI protocol stack is presented in Figure 1.10.



Figure 1.9 A subject interacting with a virtual keyboard while wearing a standard 10–20 electroencephalogram skullcap (Source: Internet image – www.lce.hut.fi/research/css/bci *Copyright requested)

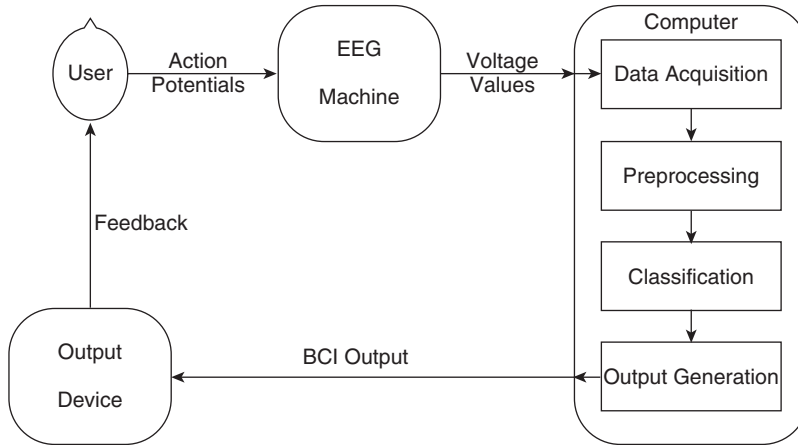


Figure 1.10 An example of a typical BCI protocol stack, displaying the principal features and the feedback loop (Source: Felzer, 2001)

1.3 The Biometric Process

Virtually all biometric-based authentication systems operate in a standard triage fashion: enrollment, model building, and decision logic. This set of processes is depicted in Figure 1.11. The purpose of enrollment is to acquire data from which the other two modules can be generated. In addition, it serves to incorporate a user into the pool of valid users – which is essential if one wishes to authenticate at a later date. The enrollment process varies little across biometrics modalities with respect to the user’s participation: to provide samples of data. How much data are required depends on how the biometric operates. Typically, the inherent variability of a biometric modality will have a significant impact on the quality of the data obtained during enrollment. Issues of user acceptability, in terms of the effort to enroll, are a significant constraint and must be taken into account when developing the particular biometric. It is of no use if the system generates 100% classification accuracy if it is too invasive or labor intensive. This is an issue that distinguishes physiological from behavioral biometrics. Physiological biometrics is based on the notion of anatomical constancy and individual variation. One would expect that in this situation, enrollment would be minimal. For instance, in a fingerprint-based system, once all fingers are recorded, there would be no need to repeat the process 10 times for instance. A fingerprint is a fingerprint? The same may not hold true for behavioral biometrics, where there is an inherent variability in the way the process is repeated. Signatures are rarely identical, and the irony of it all is that the technology scrutinizes our behavior at such a low level that it is bound to find some variation even when we as humans, examining two versions of a signature produced by the same individual, find no clear differences.

There are two principal classes of features that can be acquired during enrollment, which can be categorized into static and dynamic features. Typically, static features capture the global aspects of the enrollment data. For instance, in signature verification, the static features capture the width/height ratios, and the overall time interval during which the signature is

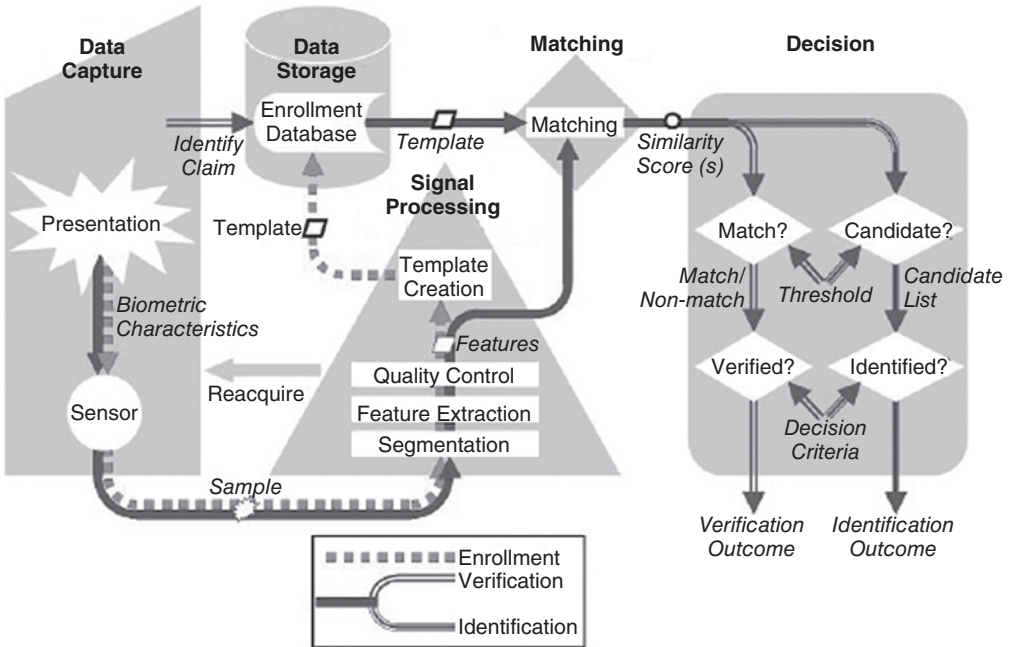


Figure 1.11 The elements that comprise a complete biometric-based system, suitable for both verification (authentication) and identification (Source: ISO/IEC JTC 1/SC 37 N1272, Figure 1 2005-08-30)

entered. Dynamic features include how the enrollment data change while they are being entered, such as the acceleration or the change in typing speed over time. One could envision that the static data are used as a gross approximation, with the dynamical features added in the event of a borderline decision. This presupposes that the static data are less informative than the dynamical data. But at the same time, the issue of constancy might weigh static data more heavily than dynamical data, which tends to be more variable. Finding this balance is a difficult task as it is not known in advance of the study. Typically, the results of the study are used to weigh the features – and different studies produce varying results – as the conditions are rarely identical between studies. There are also issues of data fusion – how does one incorporate a variety of features, which may operate on different timescales and differing magnitudes? These are important issues that will be discussed in Chapter 7, where multimodal biometrics is addressed.

Once these issues have been resolved, the ultimate result of the enrollment process is the generation of a biometric information record (BIR) for each user of the system. How do we transform the data that are collected during enrollment into a useful model? In part, this is a loaded question. On the one hand, one would assume that a model was available prior to collecting the data. But in reality, a lot of exploratory analysis is performed, where one collects all the data that appear possible to collect, and generates a collection of models, trying each to find out which provides the best classification accuracy. But the question is where did the model come from in the first place? This is the way science progresses, so we proceed as normal barring any other indication.

There are a vast number of models that have been employed in behavioral biometrics. It is beyond the scope of this text to explore this area, as it would fill a number of volumes. The case studies that occupy the majority of this text provide some examples of a variety of approaches that have been successfully applied in this domain. Assuming that a BIR is created for each successfully enrolled person, a database is created with the BIR data. There are issues here as well. Should the data be encrypted to help reduce the success of an off-line attack? Generally, the answer is yes, and many systems do employ online encryption technology.

The decision logic is designed to provide an automated mechanism for deciding whether or not to accept or reject a user's attempt to authenticate. When users make a request to authenticate, their details are extracted and compared in some way to the stored BIR. In order to decide whether to accept or reject the request, a decision process must be invoked in order to decide whether or not to accept the request. Typically, this entails comparing the features extracted from the authentication attempt with the stored BIR. There are a number of similarity metrics that have been employed in this domain. A factor that significantly impacts the matching/scoring process is whether or not the system utilizes a static or dynamic approach. For instance, in keystroke dynamics, one can employ a fixed text or a variable text approach to authentication. For a fixed text approach, a specific set of characters are typed – which can be directly compared to the BIR. This is a much easier decision to make than one based on a more dynamic approach, where the characters entered are contained within a much larger search space of possible characters. Of course, the ease with which the decision can be made is contingent upon the model building component but nonetheless has a significant impact on the decision logic. Given that a decision has been rendered regarding an authentication attempt, how do we categorize the accuracy of the system? What metrics are available to rate various decision models?

In part, this depends on the exact task at hand: is it a verification or identification task? Clearly an authentication task (also known as identification), the goal is to confirm the identity of the individual. This can simplify the match and scoring processes considerably as it reduces the search task to a 1 : 1 mapping between the presumed identity and that stored in the database. The verification task is depicted in Figure 1.12.

The task of identification is considerably more difficult than authentication in most cases. The entire database must be examined as there is no information that could narrow down the

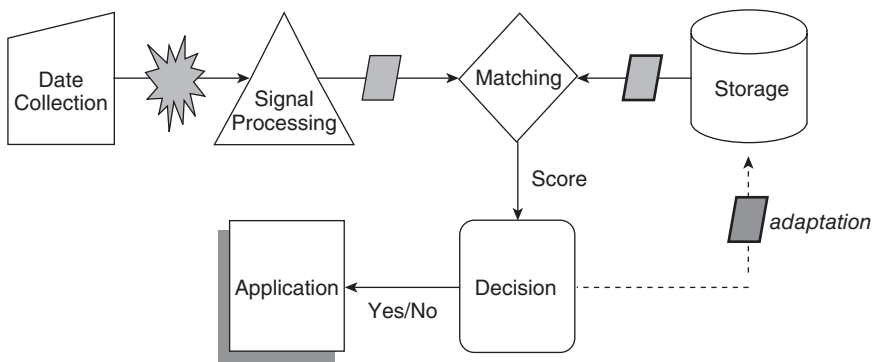


Figure 1.12 A graphical depiction of the verification process model indicating the principal elements and their potential interactions

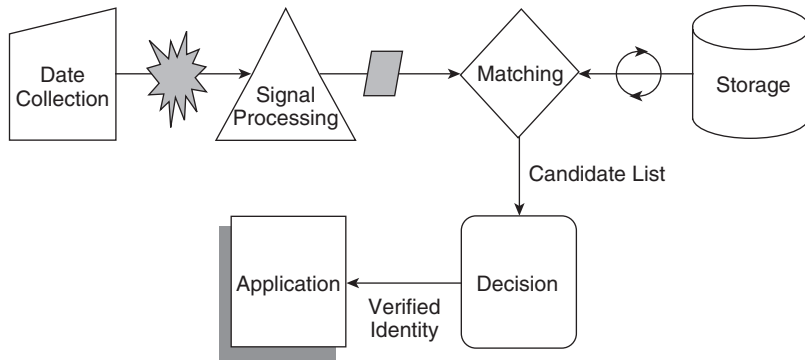


Figure 1.13 The identification process model depicting the principal difference between verification and identification, the candidate list element (see the text for details)

search. As depicted in Figure 1.13, the two process models are similar – barring the candidate list component, found only in the identification model. Another subtle distinction between these two approaches is depicted by the “adaptation” component present in the verification process model (Figure 1.12). Adaptation of the BIR is a vitally important feature of a mature and viable biometrics. Take for instance a keystroke dynamics-based authentication system. After the users complete enrollment and continue entering their password, the typing style might change slightly due to a practice effect or for other reasons. If the user is continuously matched against the enrollment data, the system may begin to falsely reject the genuine user. To prevent such an occurrence, the user’s BIR must be updated. How the user’s BIR evolves over time is an implementation issue. We tend to keep a rolling tally of the latest 10 successful login attempts, updating any statistical metrics every time the user is successfully authenticated. This is possible in a verification task – or at least it is easier to implement. In an identification task, the issue is how does the system actually confirm that the identification process has been successful? The system must only update the BIR once it has been successfully accessed – and this cannot be known without some ancillary mechanism in place to identify the user – sort of a catch-22 scenario. Therefore, adaptation most easily fits into the authentication/verification scheme, as depicted in Figure 1.13.

1.4 Validation Issues

In order to compare different implementations of any biometric, a measure of success and failure must be available in order to benchmark different implementations. Traditionally, within the biometrics literature, type I (FRR) and type II (FAR) errors are used as a measure of success. Figure 1.14 illustrates the relationship between FAR, FRR, and the EER, which is the intersection of FAR and FRR when co-plotted. Note that some authors prefer to use the term crossover error rate (CER) as opposed to the EER, but they refer to identical concepts. When reading the literature, one will often find that instead of FAR/FRR, researchers report FAR and the imposter pass rate (IPR). The confusion is that this version of FAR is what most authors’ term FRR, and the IPR is the common FRR. Another common metric prevalent in the physiological literature is the false matching rate (FMR) and false

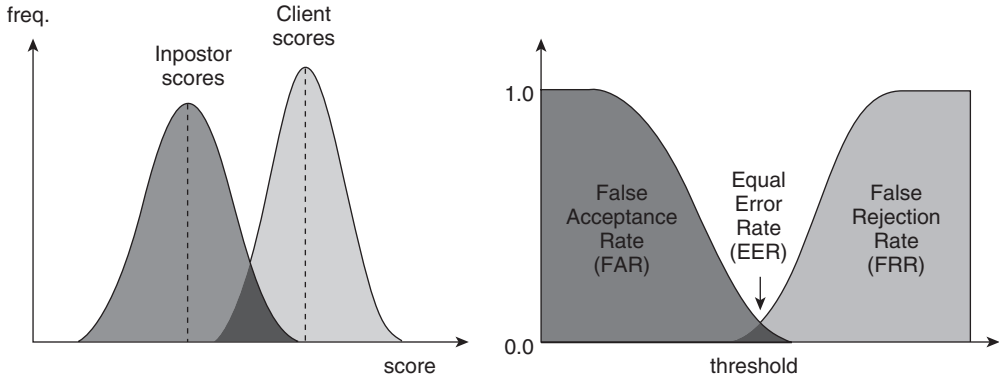


Figure 1.14 When the FAR and FRR are plotted on the same graph, as a function of a classification parameter, the intersection of the two functions is termed the EER or CER (Source: Google image: www.bioid.com/sdk/docs/images/EER_all.gif)

non-matching-rate (FNMR). The FMR is used as an alternative to FAR (FRR). Its use is intended to avoid confusion in applications that reject the claimants (i.e. an impostor) if their biometric data match that of an enrollee. The same caveat applies to FNMR as well.

A common result reported in the literature is the interdependence between the FAR and FRR. Most studies report that one cannot manipulate one of the metrics without producing an inverse effect on the other. Some systems can produce a very low FAR – but this generally means that the system is extremely sensitive and the legitimate user will fail to authenticate (FRR) at an unacceptable level. From a user perspective, this is very undesirable, and from the corporate perspective, this can be quite expensive. If users fail to authenticate, then their account is usually changed, and hence the user will have to reenroll into the system. In addition, the help desk support staff will be impacted negatively in proportion to the user support required to reset the users' account details. On the other hand, when the FRR is reduced to acceptable levels, then the FAR rises, which tends to increase the level of security breaches to unacceptable levels. Currently, there is no direct solution to this problem. One possible approach is to use a multimodal biometric system, employing several technologies. This approach doesn't solve the FAR/FRR interdependency but compensates for the effect by relaxing the stringency of each component biometric such that both FAR and FRR are reduced to acceptable levels without placing an undue burden on the user. The use of a multimodal approach is a very active research area and will be discussed in some detail in Chapter 7.

In addition to FAR/FRR and their variants, it is surprising that the concepts of positive predictive value (*PPV*) and negative predictive value (*NPV*), along with the concepts of sensitivity and specificity, often reported in the classification literature. *PPV* is the positive predictive value and the *NPV* negative predictive value of a classification result. The *PPV* provides the probability that a positive result is actually a true positive (that is a measure of correct classification). The predictive negative value (PNV) provides the probability that a negative result will reflect a true negative result. From a confusion matrix (sometimes referred to as a contingency matrix), one can calculate the *PPV*, *NPV*, sensitivity, specificity, and classification accuracy in a straightforward fashion, as displayed in Table 1.1.

The values for *PPV*, *NPV*, sensitivity, specificity, and overall accuracy can be calculated according to the following formulas (using the data in the confusion matrix):

Table 1.1 A sample confusion matrix for a two-class decision system

	Negative	Positive	
Negative	190 (<i>TN</i>)	10 (<i>FP</i>)	Specificity
Positive	10 (<i>FN</i>)	190 (<i>TP</i>)	Sensitivity
	<i>NPV</i>	<i>PPV</i>	Accuracy

TN = true negative, *FP* = false positive, *FN* = false negative, *TP* = true positive.

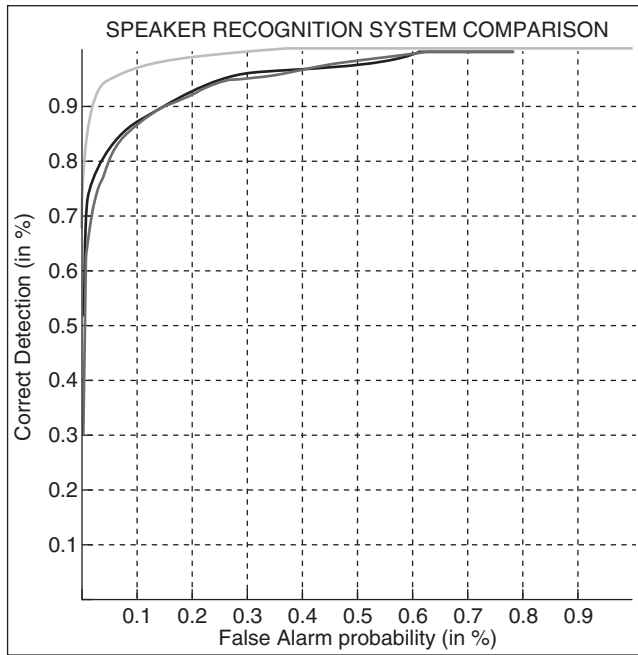


Figure 1.15 An example of an ROC curve, which displays the relationship between specificity and sensitivity (the *x*-axis is 1 specificity), and the *y*-axis is the sensitivity. The closer the curve approaches the *y*-axis, the better the result. Typically, one calculates the area under the curve to generate a scalar measure of the classification accuracy (Source: Martin et al., 2004)

$$\begin{aligned}
 \text{Sensitivity} &= TP / (FN + TP) \\
 \text{Specificity} &= TN / (TN + FP) \\
 PPV &= TP / (TP + FP) \\
 NPV &= TN / (TN + FN) \\
 \text{Accuracy} &= (TN + TP) / (TN + FP + FN + TP)
 \end{aligned}$$

Furthermore, the use of specificity and sensitivity can be used to produce an receiver operator characteristic (ROC) curve (see Figure 1.15). The ROC curve displays the interplay between

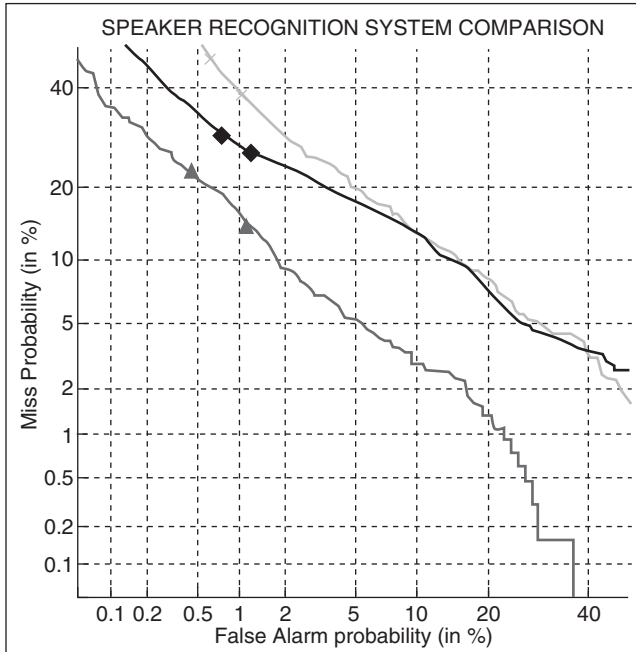


Figure 1.16 An example of a DET curve (using the same data used to plot the ROC curve in Figure 1.15). (Source: Martin et al., 2004)

sensitivity and specificity; it quantifies the relationship between FAR/FRR, in a form that is more quantitative than a simple EER/CER plot. In addition, the likelihood ratio (LR) can be obtained simply by measuring the slope of the tangent line at some cutoff value. These measurements are very useful in assessing the quality of a classification result. They are used quite frequently in the data mining literature and related fields, but for some reason have not found a place in the biometrics literature.

In addition to the above metrics, the detection error trade-off (DET) curve may be reported (see Figure 1.16 for an example of a DET curve). To generate a DET curve, one plots the FAR or equivalent on the x-axis and the FAR or equivalent on the y-axis. Typically, this plot yields a reasonably straight line, and provides uniform treatment to both classes of errors. In addition, by the selection of judicious scaling factors, one can examine the behavior of the errors more closely than the ROC. Of course, the FAR/FRR values are obtained as a function of some threshold parameter. With a complete system in hand, we can now address the important issue of biometric databases – valuable sources of information that can be used to test our particular biometric implementation.

1.5 Relevant Databases

One of the key issues with developing a novel biometric or developing a new classification strategy is to test it on an appropriate dataset. The case studies presented in this book employ a variety of techniques to acquire the data required for biometric analysis. But generally

Table 1.2 FVC2000 summary of the four databases employed in the competition. Note that w is the width and d is the depth (the dimensions of the image)

	Sensor type	Image size	Set A ($w \times d$)	Set B ($w \times d$)	Resolution
DB1	Low-cost optical sensor	300×300	100×8	10×8	500 dpi
DB2	Low-cost capacitive sensor	256×364	100×8	10×8	500 dpi
DB3	Optical sensor	448×478	100×8	10×8	500 dpi
DB4	Synthetic generator	240×320	100×8	10×8	About 500 dpi

Each is distinguished based on the type of sensor that was used to acquire the fingerprints (DB1-3), and DB4 contained synthetic signatures (Source: <http://bias.csr.unibo.it/fvc2000> – free access website).

speaking, most people use local data collected in their own particular style. In this section, we will review some examples of databases that have been made publicly available for research purposes (and in Section 1.6, a discussion of ontologies and standards is discussed – both are intimately related).

The majority of biometric databases contain information on fingerprint, voice, and face data (Ortega-Garcia, 2003, Ortega & Bousono-Crespo, 2005). The fingerprint verification competition (FVC2000) databases were started in 2000 as an international competition to test classification algorithms (FVC2000, FVC2002, FVC2004). The FVC2000 competition was the first such event, bringing researchers from academia and industry in to compete. The data consisted of four separate databases (see Table 1.2 for details), which included different types of fingerprint scanners (optical, capacitive, etc.) along with synthetic data. The primary purpose of this competition was to determine how accurately we could identify a fingerprint based on automated techniques (cf. automated fingerprint identification system). The competition was advertised to anyone wishing to enter – with the express purpose of producing a classifier with the lowest EER. The results from this first competition (FVC2000) are summarized in Table 1.3. As can be observed from the results for the EER was approximately 1.7%. Note that this value was the average across all four databases. According to Maio and colleagues (2004), the purpose of this competition can be summarized by this quote from the authors: “The goal of a technology evaluation is to compare competing algorithms from a single technology. Testing of all algorithms is done on a standardized database collected by a ‘universal’ sensor. Nonetheless, performance against this database will depend upon both the environment and the population in which it was collected. Consequently the ‘three bears’ rule might be applied, attempting to create a database that is neither too difficult nor too easy for the algorithms to be tested. Although sample or example data may be distributed for developmental or tuning purposes prior to the test, the actual testing must be done on data that has not been previously seen by the algorithm developers. Testing is done using ‘off-line’ processing of the data. Because the database is fixed, results of technology tests are repeatable” (Table 1.3).

The competitors were judged on several criteria, but the average EER across all four databases was considered the de facto benchmark. As can be seen, the average EER was approximately 1.7%, and the adjusted EER (Avg EER* in Table 1.3) represents an adjustment based on whether there were rejections during the enrollment (see the fourth column in Table 1.3). These results are impressive, and it is interesting to note that the latest competition, FVC2004,

Table 1.3 Summary of the classification results from the first fingerprint verification competition (FVC2000)

Algorithm	Avg EER (%)	Avg EER* (%)	Avg REJ _{ENROLL} (%)	Avg REJ _{MATCH} (%)	Avg enroll time (sec)	Avg match time (sec)
<i>Sag1</i>	1.73	1.73	0.00	0.00	3.18	1.22
<i>Sag2</i>	2.28	2.28	0.00	0.00	1.11	1.11
<i>Cspn</i>	5.19	5.18	0.14	0.31	0.20	0.20
<i>Cetp</i>	6.32	6.29	0.00	0.02	0.95	1.06
<i>Cwai</i>	7.08	4.66	4.46	3.14	0.27	0.35
<i>Krdl</i>	10.94	7.59	6.86	6.52	1.08	1.58
<i>Utwe</i>	15.24	15.24	0.00	0.00	10.42	2.67
<i>Fpin</i>	15.94	15.94	0.00	0.00	1.22	1.27
<i>Uinh</i>	19.33	17.31	3.75	5.23	0.71	0.76
<i>Diti</i>	20.97	20.97	0.00	0.00	1.24	1.32
<i>Ncmi</i>	47.84	47.88	0.00	0.09	1.44	1.71

Please note that for a correct interpretation of the results, Avg EER alone is not an exhaustive metric, but Avg REJ_{ENROLL} should be also taken into account (Source: <http://bias.csr.unibo.it/fvc2000> – free access website).

yielded a slightly higher average EER, just over 2%. Presumably, the technology – both from a signal acquisition and classification perspective, had increased during the 4 years between these competitions. This interesting fact alludes to the caution that should be applied when considering the classification results obtained from such studies. These were large-scale datasets – one should be cautious when examining much smaller datasets – how well do they cover the possible spectrum of events possible within the domain of interest? Have the classification algorithms been tailored to the data? A common issue of over-fitting may result if one is not careful. Ideally, after one has developed a classification algorithm that works well with a local database – in essence treating it as the training case – then the algorithm(s) should then be applied to a non-training database to see how well the results extrapolate. For more details on these datasets, please consult Maio and colleagues (2003).

The next dataset to be examined is from the behavioral biometrics literature. The signature verification competition (SVC2004) premiered in 2004, in conjunction with the FVC2004 and the FAC2004 (the latter being a face verification competition using the BANCA dataset, sponsored by the International Conference on Biometric Authentication (<http://www.cse.ust.hk/svc2004/>)). Two datasets were used in this competition: the first (DB1) contained static information regarding the coordinate position of the pen, and the second (DB2) contained coordinate information plus pen orientation and pen pressure. The signatures contained controls and forgeries – the latter consisted of skilled forgeries and causal forgeries (see Chapter 3 for details on different types of forgeries). Generally, the skilled forgeries were obtained from participants who could see the actual signature being entered and had some amount of time to practice. The results from this study are summarized in Table 1.4. It is interesting to note that the average EER for signature was not very different from that of the fingerprint competition (for FVC2004, the best average EER was 2.07% versus 2.84% for signature verification; see Yeung et al., 2004 for more details). Note also that there was a very considerable range of EER values obtained (see Table 1.4) in the signature verification competition.

Table 1.4 Summary of the classification results from the first signature verification competition (SVC2004) sponsored by the International Conference on Biometrics consortium (Source: <http://www.cse.ust.hk/svc2004/#introduction> – free access website)

Test set (60 users):

Team ID	10 genuine signatures + 20 skilled forgeries				10 genuine signatures + 20 random forgeries			
	Avg EER (%)	SD EER (%)	Max EER (%)	Min EER (%)	Avg EER (%)	SD EER (%)	Max EER (%)	Min EER (%)
106	2.84	5.64	30.00	0.00	2.79	5.89	50.00	0.00
124	4.37	6.52	25.00	0.00	1.85	2.97	15.00	0.00
126	5.79	10.30	52.63	0.00	5.11	9.06	50.00	0.00
119b	5.88	9.21	50.00	0.00	2.12	3.29	15.00	0.00
119c	6.05	9.39	50.00	0.00	2.13	3.29	15.00	0.00
115	6.22	9.38	50.00	0.00	2.04	3.16	15.00	0.00
119a	6.88	9.54	50.00	0.00	2.18	3.54	22.50	0.00
114	8.77	12.24	57.14	0.00	2.93	5.91	40.00	0.00
118	11.81	12.90	50.00	0.00	4.39	6.08	40.00	0.00
117	11.85	12.07	70.00	0.00	3.83	5.66	40.00	0.00
116	13.53	12.99	70.00	0.00	3.47	6.90	52.63	0.00
104	16.22	13.49	66.67	0.00	6.89	9.20	48.57	0.00
112	28.89	15.95	80.00	0.00	12.47	10.29	55.00	0.00

This variability in the results must be reported – and the use of an average EER goes some way toward presenting the variability in the results. One will also notice that the details of the collection of the datasets is generally underdetermined – in that even for SVC2004, there are significant differences in the description of the datasets between DB1 and DB2 – making it difficult at best to produce these databases. These issues will be discussed next in the context of international standards and ontologies.

1.6 International Standards

The biometrics industry has undergone a renaissance with respect to the development and deployment of a variety of physiological and behavioral biometrics. Physiological biometrics such as fingerprints and iris and retinal scanners were developed first, followed by behavioral-based biometric technologies such as gait, signature, and computer interaction dynamics. These developments were driven for the most part by the needs of e-commerce and homeland security issues. Both driving forces have become borderless and hence must be compatible with a variety of customs and technological practices in our global society. Thus, the need arose to impose a standardization practice in order to facilitate interoperability between different instantiations of biometric-based security. As of 1996, the only standard available was the forensic fingerprint standards. Standards bodies such as the National Institute of Standards (NIST) and the International Standards Organization (ISO) have become directly involved in creating a set of standards to align most of the major biometric methodologies into a common




Figure 1.17 The three blocks contained within the CBEFF standard template for biometric data storage. The “SBH” block is the standard biometric header; the “BSMB” is the biometric specific memory block, and the “SB” block is an optional signature block (Source: Podio et al., 2001 (Figure2))

framework for interoperability purposes (<http://www.iso.org/>). The first major standardization effort was initiated in 1999 by the NIST (<http://www.nist.gov/>). Through a meeting with the major biometric industry players, a decision as to whether a standard template could be generated that would suit all of the industry leaders was examined. In the end, no agreement was met, but within a year, the Common Biometrics Exchange File Format (CBEFF) format was proposed. The CBEFF 1.0 was finalized as a standard in 2001 under the auspices of the NIST/Biometric Consortium (BC) and was made publicly available under an NIST publication NISTIR 6529 (January 2001). In 2005, CBEFF 1.1 was released under ANSI/INCITS 398-2005, and CBEFF 2.0 was released under the auspices of ISO/IEC JTC1 (ISO/IEC 19785-1) in 2006 (<http://www.incits.org/>). The overall structure of the CBEFF is depicted in Figure 1.17. It consists of three blocks onto which the required information is mapped onto. The purpose of the CBEFF was to provide biometric vendors a standard format for storing data for the sole purpose of interoperability. The basic format of the CBEFF is depicted in Figure 1.17. It consists of three elements – a header block, the data block, and an optional signature block (SB). Each block consists of a number fields that are either mandatory or optional. The essential features of the CBEFF template are

- facilitating biometric data interchange between different system components or systems;
- promoting interoperability of biometric-based application programs and systems;
- providing forward compatibility for technology improvements;
- simplifying the software and hardware integration process.

In summary, the standard biometric header (SBH) is used to identify the source and the type of biometric data – the format owner, the format type, and security options – these fields are mandatory. There are, in addition, several optional fields that are used by the BioAPI (discussed later in this chapter). The biometric specific memory block (BSMB) contains details on the format and the actual data associated with the particular biometric, and its specific format is not specified. Lastly, the optional SB is an optional signature that can be used for source/destination verification purposes. Table 1.5 lists the fields contained within the CBEFF blocks. For more details, please consult Reynolds (2005).

In addition to the development of the CBEFF standardized template, several variations and/or enhancements have been added to facilitate application development and to enhance security. In particular, the International Biometrics Industry Association (IBIA) is the body responsible for ensuring that all biometric patrons are properly registered and provided with a unique ID number (<http://www.bioapi.org/>). Clients can then register their biometric solutions with an appropriate patron. This patron/client relationship is depicted in Figure 1.18.

The CBEFF template does not specify at any level how the applications that acquire and utilize biometric information should be developed. To enhance the software development

Table 1.5 A depiction of the fields within the standard CBEFF template. (Source: International Standards Organization, ANSI/INCITS 398-2005)

Ident	Cond Code	Field Number	Field Name	Char Type	Field size per occurrence		Occur count		Max byte count
					min	max	min	max	
					LEN	M	18.001	LOGICAL RECORD LENGTH	
IDC	M	18.002	IMAGE DESIGNATION CHARACTER	N	2	5	1	1	12
RSV	–	18.003	RESERVED FOR FUTURE	–	–	–	–	–	–
HDV	M	18.099	INCLUSION						
	M	18.100	CBEFF HEADER VERSION	N	5	5	1	1	12
BTY	M	18.101	BIMOETRIC TYPE	N	9	9	1	1	16
BDQ	M	18.102	BIOMETRIC DATA QUALITY	AN	2	4	1	1	11
BFO	M	18.103	BDB FORMAT OWNER	AN	5	5	1	1	12
BFT	M	18.104	BDB FORMAT TYPE	AN	5	5	1	1	12
RSV	–	18.105	RESERVED FOR FUTURE	–	–	–	–	–	–
		18.199	INCLUSION						
UDF	0	18.200	USER-DEFINED	–	–	–	–	–	–
		18.998	FIELDS						
BDB	M	18.999	BIOMETRIC DATA	B	2	–	1	1	–

cycle, the BioAPI was developed – and indeed was part of the driving force for the development of the CBEFF. The BioAPI has its own version of the CBEFF – defined as a biometric identification record (BIR) (In later versions, BIR is used more generically and stands for biometric information record). A BIR refers to any biometric data that is returned to the application, including raw data, intermediate data, processed sample(s) ready for verification or identification, as well as enrollment data. The BIR inherits the standard structure of CBEFF and inserts detailed information into the SBH which makes it possible to be interpreted by BioAPI devices. The BioAPI has extended the original CBEFF by developing a suite of software development tools. By subsuming the CBEFF (via inheritance), it provides a complete program development environment for creating a variety of biometric applications. For more details, please consult BioAPI (<http://www.nationalbiometric.org/>) (Figure 1.19).

The last issue that has been addressed with regards to biometric standards is that of enhanced security – which was not part of the original CBEFF model. To enhance the security features of this model, the X9.84 specification was created. It was originally designed to integrate biometrics within the financial industry. Subsequently, the security features can be used in biometric applications regardless of the nature of the end user. In 2000, ANSI X9.84-

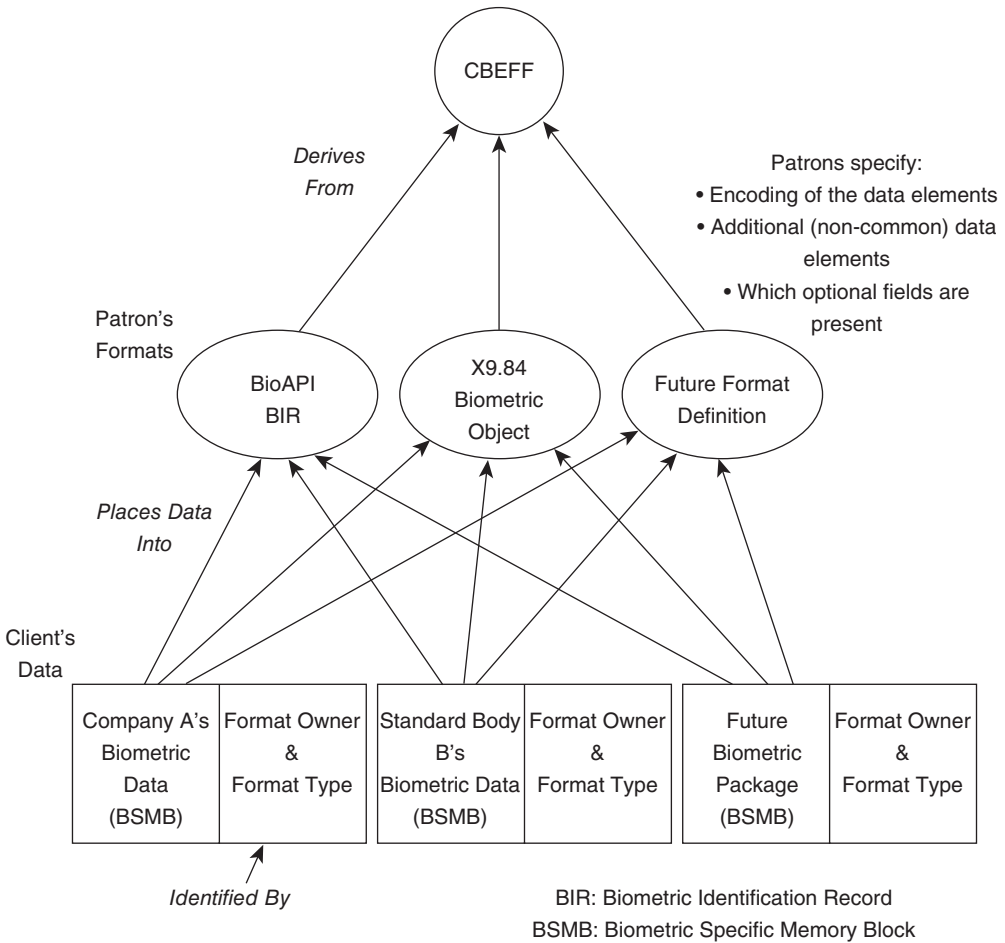


Figure 1.18 Patron/client architecture of the current working model. A client must register with a patron, who has the responsibility to ensure that the standards are adhered to and that any new technologies are properly defined and subsequently registered appropriately (Source: Tilton, 2003)

2002, Biometric Information Management and Security for the Financial Services Industry, was published (http://www.incits.org/tc_home/m1htm/docs/m1050246.htm). This standard provides guidelines for the secure implementation of biometric systems, applicable not only to financial environments and transactions but far beyond. The scope of X9.84-2002 covers security and management of biometric data across its life cycle, usage of biometric technology for *verification* and *identification* for banking customers and employees, application of biometric technology for physical and logical access controls, encapsulation of biometric data, techniques for securely transmitting and storing biometric data, and security of the physical hardware (Tilton, 2003). X9.84 begins by defining a biometric data processing framework. This framework defines common processing components and transmission paths within a biometrically enabled system that must be secured. Figure 1.19 summarizes the X9.84 specification.

Basic Functions	Primitive Functions
Module Management BioAPI_ModuleLoad BioAPI_ModuleAttach	BioAPI_Capture Captures raw/intermediate data from sensor
Data Handling BioAPI_GetBIRFromHandle BioAPI_GetHeaderFromHandle	BioAPI_Process Converts raw sample into processed template for matching
Callback & Event Operations BioAPI_SetStreamCallback	BioAPI_CreateTemplate Converts raw sample(s) into processed template for enrollment
Biometric Operations <u>BioAPI_Enroll</u> – Captures biometric data and creates template <u>BioAPI_Verify</u> – Captures live biometric data and matches it against one enrolled template <u>BioAPI_Identify</u> – Captures live biometric data and matches it against a set of enrolled template	BioAPI_VerifyMatch Performs a 1:1 match BioAPI_IdentifyMatch Performs a 1:N match BioAPI_Import Imports non-real-time data for processing

Figure 1.19 Summary of some of the major modules within the BioAPI version 1.1 framework. See BioAPI (<http://www.nationalbiometric.org/>) for details. This list contains many (but not all) of the primitive and basic functions required for the Win32 reference implementation of the BioAPI framework (Source: International Standards Organization, ANSI/INCITS 398-2005)

Note that recently, the BioAPI has been updated to version 2.0, which extends the previous version (ISO/IEC-19794-1, 2005). The principal change is the expansion of the patron/client model – which now includes devices, allowing for a proper multimodal approach. This should help facilitate interoperability – as it has moved the emphasis from the business collaboration perspective down to particular implementations. We will have to wait and see if this enhancement facilitates.

To summarize what is available in terms of a standard for biometric data interchange, we essentially have an available application programming interface Application Programming Interface (BioAPI), a security layer (X9.84), and a standardized template (BIR and CBEFF). The API is used to integrate the client (biometric applications) via a common template to other biometric clients implementing implements the interoperability requirement set forth by the standards organization. If you examine the patron list (<http://www.ibia.org/>), you will notice that there are no behavioral biometric patrons. This could be explained by a paucity of biometric clients, but if you look at the literature, there are a number of behavioral-based biometrics in the marketplace. Consider BioPassword[®], a leader in keystroke dynamics-based biometrics. They claim to be driving forward via an initiative with INCITS, a keystroke dynamics-based interchange format (CEBFF compliant) BSMB. Yet they have not yet registered as a patron/client with IBIA – it simply might be a matter of time. In addition to BioPassword[®], there are a number of other vendors with behavioral-based biometrics – employing gait analysis, signature verification, and voice detection as viable biometric solutions. Why no behavioral biometric solution has registered is an interesting question (although BioPassword[®] is spearheading the registration of their keystroke dynamics product, BioPassword[®]).

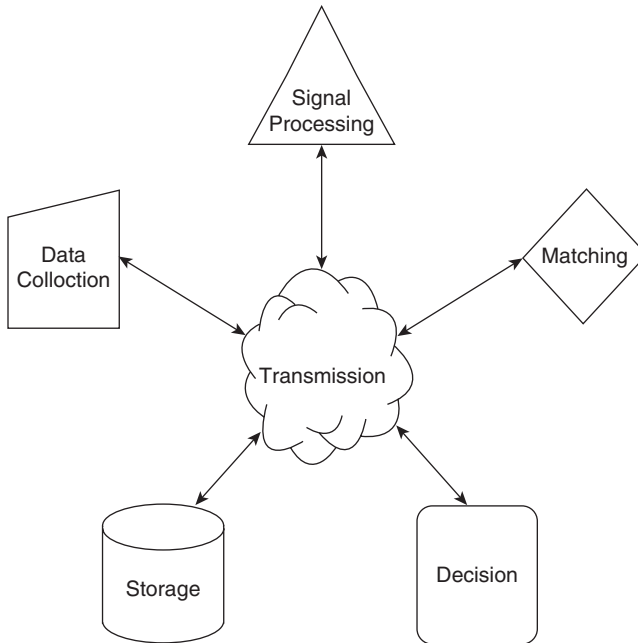


Figure 1.20 Depiction of the X9.84 biometric framework for secure transmission of biometric data over data channel that requires security

This could be the result of the difficulty in establishing a new patron or inherent differences between physiological versus behavioral biometrics.

As is displayed in Figure 1.18, there are only a few patrons – BioAPI and X9.84. These patrons are the result of a large organizational structure that is an amalgamation of international standard bodies and industry leaders. In order to augment the list of patrons, the industry must be willing to cooperate and work with these standards bodies in order to assert their standards with the constraint of being CBEFF compliant. At the client level, organizations (standards or industry) can produce a CBEFF-compliant BSMB for instance, but any additional changes are made at the API level and hence are proprietary in a sense.

Another possibility for the lack of behavioral biometric patrons is the inherent difference(s) between physiological and behavioral biometrics. For instance – though both classes of biometrics require an enrollment process – enrollment in behavioral biometrics may be significantly different than the method employed in physiological biometrics. For instance, enrolling in a fingerprint or retinal scanner may be a more straightforward task than enrolling in a keystroke or mouse dynamics-based biometric. In addition, there is a temporal factor in some behavioral-based biometrics. With keystroke dynamics-based systems, typing styles may change over time or a users' typing style may adapt as they learn their login details through a practice effect. The temporal changes must be captured if the authentication module is to perform at optimal levels. The same considerations apply to mouse dynamics, which are similar to keystroke dynamics except that they are applied to a graphical authentication system. Signature-based authentication systems tend to be more stable than keystroke/mouse

dynamics – and so may be more similar to physiological biometrics in this respect. Are these considerations worthy of addressing? If so, how can the existing standards address these issues?

The primary consideration in this chapter is that the BioAPI/X9.84 standards are not robust enough to allow a complete integration of all extant biometric modalities. Though all must conform to the general framework (depicted in Figure 1.19) if they are to be considered for inclusion. The common denominator is that all clients must conform to the CBEFF template format. If you examine the format, you will notice that the vast majority of the fields are optional (17/20 for the SBH alone). Of those that are required, there is a limited vocabulary that is available to choose from – principally codes for format type, etc. An erroneous entry for a field value is handled by the client software and is not included directly as part of the standard. The values for optional fields do not conform to any standard and hence are, from an ontology perspective, ineffectual. Lastly, the “ontology” engendered by either patron has only minimal support for behavioral biometrics such as keystroke/mouse dynamics. We propose that an existing or a new patron be developed that can address these issues at the level of the standard itself – not at the level of the implementation. We propose that a proper ontology must be created in order to ensure that standards are developed properly and encompass all the extant biometric tools that are currently available, from an interoperability and research perspective (Revett, 2007a). An analogy with existing ontologies may be useful in this regard.

One useful ontology that has been very successful is the microarray gene expression data (MGED) ontology (<http://www.mged.org/>). This ontology describes how gene expression data should be annotated in order to facilitate sharing data across various laboratories around the globe. The ontology is actual in that it has a data model that incorporates named fields and values for each field. It has separate modules that relate to the acquisition of the experimental material, a model for how an experiment was performed, and lastly, a module for storing the data in a Web-accessible format. This ontology has been very successful – as many research laboratories around the world are using them – allowing seamless sharing of data. We propose a similar sort of structure for a behavioral biometric-based ontology, which includes first and foremost a true ontology where data fields are required and values for these fields are from a controlled list. A data structure similar to the CBEFF can be used, but it is not *the* single point of commonality between different biometric systems. Rather, the CBEFF is simply a data storage module that can be used by any biometric system. The fields contained within the data storage module must be more comprehensive and must be generated in the form of some type of object model, similar to the MGED standard.

This discussion has described the need for a comprehensive ontology for behavioral biometrics. The need for such an ontology is premised on the examples of how the attribute selection process and testing protocol can influence the results at all stages of the software development cycle of biometric software. Poor attribute selection will invariably produce a product that is inferior with respect to generation of adequate FAR/FRR results. Even if the attributes are selected reasonably, how they are utilized in the authentication algorithm is highly instance dependent and will clearly vary from one implementation to another. Skewed testing phase results will generally produce a negative impact on the quality of the resultant biometric – possibly increasing the duration of the testing phase – and certainly will increase the cost of product development. In addition, without knowledge of how attribute extraction and the testing protocol, it will be impossible to compare the results of different

authentication algorithms even on the same dataset. The differences might result from variations in the protocol more so than on the authentication algorithm per se.

What we are proposing in this chapter is a comprehensive ontology – not just a data template as the CBEFF standard provides. The CBEFF has too many optional fields and does not include sufficient data regarding the biometric implementation to allow comparisons between different methodologies. Though this may not have been the original intention, issues highlighted in this chapter suggest that this is a critical aspect of such an effort. Interoperability between various biometrics is a noble goal. But the current standards appear to be biased toward physiological-based biometrics. Granted these are fairly stable technologies – and the attribute extraction process is well posed – they are incomplete with respect to the inclusion of behavioral-based biometrics. In addition, the ability of various researchers (both in academia and industry) to explore the same data – for corroboration and analysis purposes – is greatly hindered, resulting in duplication of effort.

This is a critical feature of a standard. The standards essentially neglect behavioral biometrics, yielding a divide between the two technologies. A proper ontology may be the answer. The MGED standards has proven extremely effective with regards to a very complicated domain – DNA microarray experiments. Something akin to the MGED ontology may be what is required to achieve interoperability between the two classes of biometrics. Having such an ontology in hand would not impede the production of new biometric solutions; in contrast, it would streamline the development process in most cases. In terms of proprietary product development – in terms of proprietary product development – an ontology does not imply that data and algorithms will be shared across the community, divulging trade secrets. Trademark work can still maintain its anonymity – there is no need to disclose secrets during the development process. When a product has reached the marketplace, intellectual property rights will have been acquired, and this protection will be incorporated into the ontology by definition. What will be made available is how the process was performed: details regarding study conditions, the attribute selection process, data preprocessing, classification algorithms, and data analysis protocols are the principal foci of the proposed ontology. The details of the classification algorithms do not have to be disclosed.

To date, there is a single proposal for an ontology/standard that encompasses a behavioral biometric authentication scheme, propounded by BioPassword (termed the Keystroke Dynamics Format for Data Interchange), published by Samantha Reynolds, from BioPassword (Reynolds, 2005). A summary of the keystroke dynamics format is presented in Table 4.13. It is unfortunate that this data format summary (or mini ontology) has so many incomplete fields, especially within the “valid values” column. One of the key features of an ontology is that it serves as a named vocabulary. All field values must be selected from a finite set of values. Still, this is a very solid start toward the development of an ontology for behavioral biometrics, and hopefully will be completed in the near future. But during this evolutionary process, it is hoped that it will be able to incorporate other types of behavioral biometrics as well. The ultimate aim would be to unite physiological and behavioral biometrics into a common universally encompassing standard.

1.7 Conclusions

The chapter has highlighted some of the major issues involved in behavioral biometrics. A summary of the principal behavioral biometrics was presented (though the coverage was not

exhaustive) highlighting the principal techniques. The focus of this book is on a remote access approach to biometrics, and as such, there is an implicit constraint that a minimal amount of hardware is required to deploy the system. One will note that in the list of behavioral biometrics, ECG, EEG, and gait were added. These approaches require some additional hardware over and above what is typically supplied with a standard PC. Their inclusion is to set the background for Chapter 8, which discusses the future of behavioral biometrics. Therefore, it should be noted that these technologies may not fall under our current working definition of a remote access approach, which can be defined as “a technique for authenticating an individual who is requesting authentication on a machine which is distinct from the server which performs the authentication process.” But if behavioral biometrics is to expand its horizons, we may have to consider other options from traditional ones such as voice, signature, and keystroke interactions. Who knows what the future of technology will bring to us – which might make these possibilities and others a feasible option in the near future.

It is hoped that this text will highlight some of the advances of behavioral biometrics into the foreground by highlighting some of the success stories (through case study analysis) that warrant a second look at this approach to biometrics. There are a variety of techniques that have been attempted, each very creative and imaginative, and based on solid computational approaches. Unfortunately, the machine learning approaches cannot be addressed in a book of this length, so the reader is directed as appropriate to a selection of sources that can be consulted as the need arises. In the final analysis, this author believes that behavioral biometrics – either alone or in conjunction with physiological biometrics – either is standard reality or in virtual reality – can provide the required security to enable users to feel confident that their space on a computer system is fully trustworthy. This applies to a standard PC, ATM, PDA, or mobile phone.

1.8 Research Topics

1. Odor has been claimed to be a useful behavioral biometric – how would one explore the individuality of this biometric?
2. Is it theoretically possible to make FAR and FRR independent of one another?
3. Do lip movements suffer the same degree of light dependence as face recognition in general?
4. Can DNA be practically implemented as a biometric, and if so, would it be best utilized as a physiological or behavioral biometric tool?
5. What factors are important in the development of a biometric ontology? How do the current standards need to be enhanced to produce a unified biometric ontology (incorporating physiological and behavioral biometrics)?
6. What new behavioral biometric lie on the horizon? Have we exhausted the possibilities?

