


Biometric Recognition: Security and Privacy Concerns

Biometrics offers greater security and convenience than traditional methods of personal recognition. In some applications, biometrics can replace or supplement the existing technology. In others, it is the only viable approach. But how secure is biometrics? And what are the privacy implications?



Is Alice authorized to enter this facility? Is Bob entitled to access this Web site or privileged information? Are we administering our service exclusively to the enrolled users? Does Charlie have a criminal record? Every day, a variety of organizations pose questions such as these about *personal recognition*.

One emerging technology that is becoming more widespread in such organizations is *biometrics*—automatic personal recognition based on physiological or behavioral characteristics.¹ The term comes from the Greek words *bios* (life) and *metrikos* (measure). To make a personal recognition, biometrics relies on who you are or what you do—as opposed to what you know (such as a password) or what you have (such as an ID card).

Biometrics has several advantages compared with traditional recognition. In some applications, it can either replace or supplement existing technologies; in others, it is the only viable approach to personal recognition. With the increasing infrastructure for reliable automatic personal recognition and for associating an identity with other personal behavior, concern is naturally growing over whether this information might be abused to violate individuals' rights to anonymity. We argue here, however, that the accountable, responsible use of biometric systems can in fact protect individual privacy.

Measurement requirements

What biological measurements qualify as biometrics? Any human physiological or behavioral trait can serve as a *biometric characteristic* as long as it satisfies the following requirements:

- *Universality*. Each person should have the characteristic.
- *Distinctiveness*. Any two persons should be different in terms of the characteristic.
- *Permanence*. The characteristic should be sufficiently invariant (with respect to the matching criterion) over a period of time.
- *Collectibility*. The characteristic should be quantitatively measurable.

However, for a practical biometric system, we must also consider issues of performance, acceptability, and circumvention. In other words, a practical system must meet accuracy, speed, and resource requirements, and it must be harmless to the users, accepted by the intended population, and sufficiently robust to various fraudulent methods and attacks.

Biometric systems

A *biometric system* is essentially a pattern-recognition system that recognizes a person based on a feature vector derived from a specific physiological or behavioral characteristic that the person possesses. Depending on the application context, a biometric system typically operates in one of two modes: *verification* or *identification*. (Throughout this article, we use the generic term “recognition” where we do not wish to distinguish between verification and identification.)

In verification mode, the system validates a person's identity by comparing the captured biometric characteristic with the individual's biometric template, which is prestored in the system database. In such a system, an individual who desires to be recognized (for example, Bob)

SALIL
PRABHAKAR
*Digital
Persona*

SHARATH
PANKANTI
*IBM T.J.
Watson
Research
Center*

ANIL K. JAIN
*Michigan
State
University*

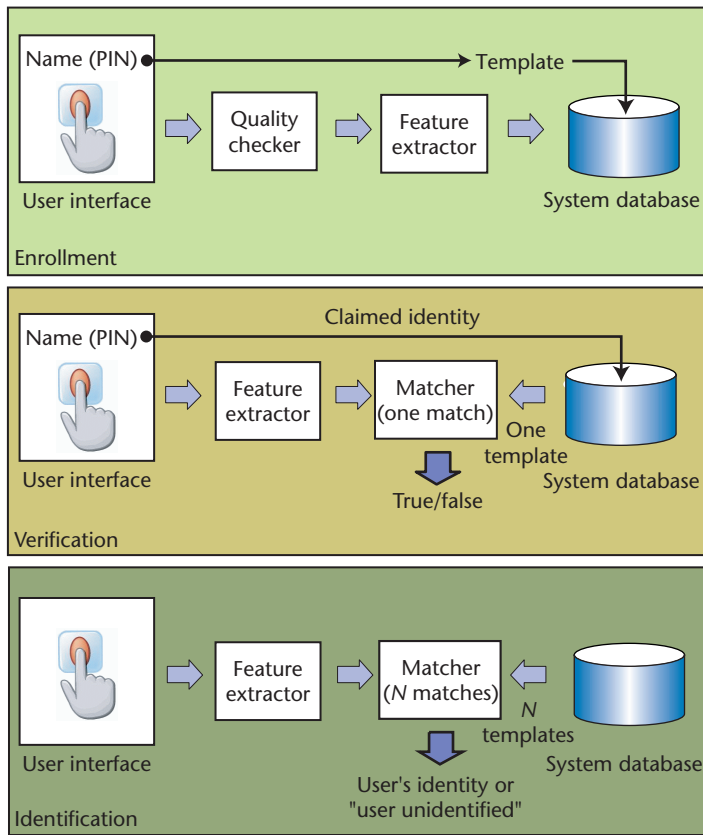


Figure 1. Block diagrams of enrollment, verification, and identification tasks. Enrollment creates an association between an identity and its biometric characteristics. In a verification task, an enrolled user claims an identity and the system verifies the authenticity of the claim based on her biometric feature. An identification system identifies an enrolled user based on her biometric characteristics without the user having to claim an identity.

claims an identity—usually via a personal identification number (PIN), login name, smart card, or the like—and the system conducts a one-to-one comparison to determine whether the claim is true. The question being answered is, “Is this person Bob?” Identity verification is typically used for *positive recognition*, where the aim is to prevent multiple people from using the same identity.

In identification mode, the system recognizes an individual by searching the entire template database for a match. The system conducts a one-to-many comparison to establish an individual’s identity (or fails if the subject is not enrolled in the system database). The question being answered is, “Who is this person?” Identification is a critical component of *negative recognition* applications, in which the system establishes whether the person is who she (implicitly or explicitly) denies being. The purpose of negative recognition is to prevent a single person from using multiple identities. Identifi-

cation can also be used in positive recognition for convenience (because the user is not required to claim an identity). While the traditional methods of personal recognition such as passwords, PINs, keys, and tokens work for positive recognition, only biometrics can be used for negative recognition.

Figure 1 contains block diagrams of a verification system and an identification system, both performing the task of user enrollment. The enrollment module registers individuals into the biometric system database. During the enrollment phase, a biometric reader (such as a fingerprint sensor or CCD camera) first scans the individual’s biometric characteristic to produce its digital representation. The system generally performs a quality check to ensure that successive stages can reliably process the acquired sample. To facilitate matching, a feature extractor processes the input sample to generate a compact but expressive representation, called a template. Depending on the application, the biometric system might store the template in its central database or record it on a smart card issued to the individual.

Biometric system errors

Two samples of the same biometric characteristic from the same person—for example, two impressions of your right index finger—are not exactly the same because of imperfect imaging conditions (such as sensor noise and dry fingers), changes in the user’s physiological or behavioral characteristics (such as cuts and bruises on the finger), ambient conditions (such as temperature and humidity), and the user’s interaction with the sensor (such as finger placement). Therefore, a biometric matching system’s response is typically a matching score s (usually a single number) that quantifies the similarity between the input and the database template representations. (For simplicity, we assume that the system actually communicates the matching score to the user. Some systems might communicate only the final decision based on a predetermined matching criterion or a threshold.) The higher the score, the more certain the system is that the two biometric measurements come from the same person.

A threshold t regulates the system decision. The system infers that pairs of biometric samples generating scores higher than or equal to t are *mate pairs* (that is, they belong to the same person). Consequently, pairs of biometric samples generating scores lower than t are *nonmate pairs* (that is, they belong to different persons). The distribution of scores generated from pairs of samples from different persons is called an *impostor distribution*; the score distribution generated from pairs of samples from the same person is called a *genuine distribution* (see Figure 2).

A biometric verification system can make two types of errors:

- Mistaking biometric measurements from two different persons to be from the same person (called *false match* or *false accept*)
- Mistaking two biometric measurements from the same person to be from two different persons (called *false non-match* or *false reject*)

An operational biometric system makes a trade-off between false match rate (FMR) and false nonmatch rate (FNMR). In fact, both FMR and FNMR are functions of the system threshold t : If the system's designers decrease t to make the system more tolerant to input variations and noise, FMR increases. On the other hand, if they raise t to make the system more secure, then FNMR increases accordingly. We can depict system performance at all operating points (thresholds t) in the form of a *receiver operating characteristic (ROC) curve*. An ROC curve plots FMR against $(1 - \text{FNMR})$ or FNMR for various values of threshold t (see Figure 3).

Besides these two recognition error rates, we can also use the rates of *failure to capture (FTC)* and *failure to enroll (FTE)* to summarize a biometric system's accuracy. The FTC rate, which only applies when the biometric device implements automatic-capture functionality, denotes the percentage of times the biometric device fails to automatically capture a sample when presented with a biometric characteristic. This type of error typically occurs when the device cannot locate a biometric signal of sufficient quality—for example, if it receives an extremely faint fingerprint or an occluded face.

The FTE rate, on the other hand, denotes the percentage of times users cannot enroll in the recognition system. There is a trade-off between the FTE rate and the perceived system accuracy (FMR and FNMR). FTE errors typically occur when the system rejects poor-quality templates during enrollment. Consequently, the database contains only high-quality templates, and the perceived system accuracy improves. Because of the interdependence among the failure rates and error rates, all these rates—FTE, FTC, FNMR, and FMR—constitute important performance metrics of a biometric system.

Given a biometric system's accuracy in verification mode, we can approximately infer its accuracy in identification mode under simplifying assumptions. Let us denote the identification false nonmatch and false match rates as FNMR_N and FMR_N , where N represents the number of identities in the system database. (For simplicity, we assume a single identification attempt per subject and a single biometric template per enrolled user.) Then, it can be shown that $\text{FNMR}_N \cong \text{FNMR}$, and $\text{FMR}_N = 1 - (1 - \text{FMR})^N \cong N \times \text{FMR}$. (This approximation holds only when $N \times \text{FMR} < 0.1$.) A report by the United Kingdom Biometric Working Group includes a detailed discussion of accuracy issues.²

A biometric system's accuracy requirements depend




Figure 2. Biometric system error rates: The curves show false match rate (FMR) and false nonmatch (FNMR) rate for a given threshold t over the genuine and imposter score distributions. FMR is the percentage of nonmate pairs whose matching scores are greater than or equal to t , and FNMR is the percentage of mate pairs whose matching scores are less than t .

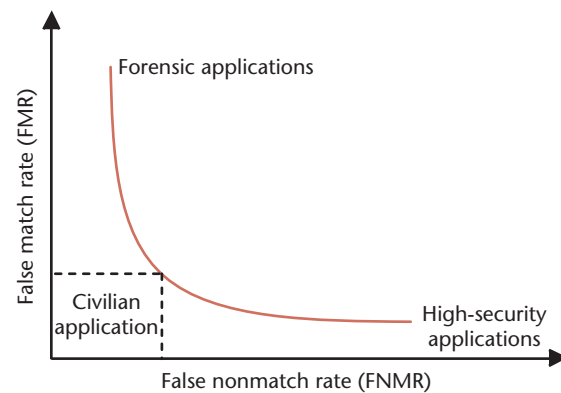


Figure 3. Receiver operating characteristic curve: Different biometric application types make different trade-offs between the false match rate and false nonmatch rate (FMR and FNMR). Lack of understanding of the error rates is a primary source of confusion in assessing system accuracy in vendor and user communities alike.

greatly on the application. For example, in some forensic applications, such as criminal identification, FNMR rate (and not FMR) is the critical design issue: that is, we do not want to miss a criminal even at the risk of manually examining a large number of potentially incorrect matches that the biometric system identifies. On the other extreme, FMR might be one of the most important factors in a highly secure access-control application, where the primary objective is deterring impostors.

Table 1. Comparison of several biometric technologies (assessments based on authors' perceptions).






BIOMETRIC	FINGERPRINT	FACE	HAND GEOMETRY	IRIS	VOICE
					
Barriers to universality	Worn ridges; hand or finger impairment	None	Hand impairment	Visual impairment	Speech impairment
Distinctiveness	High	Low	Medium	High	Low
Permanence	High	Medium	Medium	High	Low
Collectibility	Medium	High	High	Medium	Medium
Performance	High	Low	Medium	High	Low
Acceptability	Medium	High	Medium	Low	High
Potential for circumvention	Low	High	Medium	Low	High



Figure 4. Biometrics application examples. (a) Digital Persona's fingerprint verification system provides personal recognition for computer and network login. (b) Indivios manufactures a fingerprint-based point-of-sale (POS) terminal that verifies customers before charging their credit cards. (c) BioThentica's fingerprint-based door lock restricts access to premises. (d) The Inspass immigration system, developed by Recognition Systems and installed at major airports in the US, uses hand geometry verification technology.

In several civilian applications, performance requirements lie between these two extremes, and we need to consider both FMR and FNMR values. In applications such as bank ATM card verification, for example, a false match would mean the loss of several hundred dollars, while a high FNMR might lead to the loss of a valued customer. Figure 3 depicts the FMR and FNMR trade-offs in different biometric application types.

Comparison of biometrics

Several biometric characteristics are in use in various applications. Each biometric has its strengths and weaknesses, and the choice typically depends on the application. No single biometric can effectively meet the requirements of all applications—none is “optimal.” We match a specific biometric to an application depending

on the application's operational mode and the biometric characteristic's properties. For example, both the fingerprint- and iris-based techniques are more accurate than the voice-based technique. However, in a telebanking application, the voice-based technique might be preferable because the bank could integrate it seamlessly into the existing telephone system. Table 1 briefly compares five biometric techniques along seven factors. (The *Handbook of Fingerprint Recognition* includes a more detailed list.³)

Applications of biometric systems

Biometric applications fall into three main groups:

- *commercial* applications, such as computer network logins, electronic data security, e-commerce, Internet access, ATMs, credit cards, physical access control, cellular phones, PDAs, medical records management, and distance learning;
- *government* applications such as national ID cards, correctional facilities, driver's licenses, social security, border control, passport control, and welfare-disbursement; and
- *forensic* applications such as corpse identification, criminal investigation, terrorist identification, parenthood determination, and missing children.

Figure 4 shows some examples of biometric applications in use. Traditionally, commercial applications have used knowledge-based systems employing PINs and passwords, government applications have utilized systems based on tokens such as ID cards and badges, and forensic applications have relied on human experts to match biometric features.

Now let us examine the advantages and disadvantages of biometrics in various applications. In the commercial

category, applications require positive recognition and may use the biometric system either in verification or identification mode. The government and forensic categories consist of mainly negative-recognition applications that require identification.

Positive recognition: Commercial applications

As mentioned, traditional technologies available for achieving a positive recognition include knowledge-based methods (using, for example, PINs and passwords) and token-based methods (using possessions such as keys and cards).

Most people set their passwords to words or digits they can easily remember—for example, names and birthdays of family members, favorite movie or music stars, and dictionary words. (In 2001, a survey of 1,200 British office workers found that almost half chose their own name, a pet's name, or a family member's name as a password. Others based their passwords on celebrity or movie character names, such as Darth Vader and Homer Simpson.⁴)

Such passwords are easy to crack by guessing or by simple brute-force dictionary attacks. Although it is possible, and even advisable, to keep different passwords for different applications and to change them frequently, most people use the same password across different applications and never change it. Compromising a single password can thus cause a break in security in many applications. For example, a hacker might create a bogus Web site enticing users with free air miles or pornography if they register with a login name and password. The hacker could then have fair success in using the same login name and password to attack users' corporate accounts.

Longer passwords are more secure but harder to remember, which prompts some users to write them down in accessible locations (such as Post-it notes hidden under the keyboard). Strong passwords are difficult to remember and result in more help desk calls for forgotten or expired passwords. Cryptographic techniques can provide very long passwords (encryption keys) that the user need not remember; however, these are in turn protected by simple passwords, which defeats their purpose.

Given that a hacker needs to break only one password among those of all the employees to gain access to a company's intranet, a single weak password compromises the overall security of every system that user has access to. Thus, the entire system's security is only as good as the weakest password.

Finally, when a user shares a password with a colleague, there is no way for the system to know who the actual user is.

Similarly, possession-based personal recognition suffers from many problems as well. For example, keys and tokens can be shared, duplicated, lost, or stolen, or an attacker could make a master key that opens many locks.

It is significantly more difficult to copy, share, or distribute biometrics. Biometrics cannot be lost or forgotten, and online biometrics-based recognition systems require the person being recognized to be present at the point of recognition. Biometrics are difficult for attackers to forge and for users to repudiate. Furthermore, the security level is relatively equal for all users in a system, which means that one account is no easier to break than any other (for example, through social engineering). The main advantage of a biometric system is that it gives users greater convenience (they no longer have to remember multiple, long and complex, frequently changing passwords) while maintaining sufficiently high accuracy and ensuring that the user is present at the point and time of recognition.


Trojan horse attacks against a biometric system's modules and replay attacks against its communication channels (see Figure 1) are similar to those against password-based personal recognition systems. We can secure biometric systems against these attacks using the building blocks of standard cryptographic techniques.

Withstanding brute-force attacks

Now let us consider a brute-force attack on a commercial biometric system operating in verification mode. A brute-force attack's chances of success depend on the biometric verification's matching accuracy. Let us assume that a certain commercial biometric verification system operates at 0.001 percent FMR. At this setting, several biometric systems (such as state-of-the-art fingerprint- and iris-recognition systems) easily deliver less than 1 percent FNMR.⁵ An FMR of 0.001 percent indicates that if a hacker launches a brute-force attack with a large number of different fingerprints, on the average, one out of 100,000 attempts will succeed. We can consider this

Biometrics cannot be lost or forgotten.... They are difficult for attackers to forge and for users to repudiate.

equivalent to the security offered by a randomly chosen five-digit PIN (although a brute-force attack against a five-digit PIN is guaranteed to succeed in 100,000 attempts and requires only 50,000 attempts on the average). To attack a biometric-based system, however, the hacker must generate (or acquire) a large number of samples of the biometric (for example, fingerprints); this is more dif-



COPYRIGHT SPRINGER, USED BY PERMISSION.

Figure 5. Fake fingers made from consenting users. (a) Rubber stamp made from a live-scan fingerprint image. (b) Wafer-thin plastic sheet housing a three-dimensional replication of a fingerprint.

ficult than generating a large number of PINs or passwords. Finally, system administrators can arbitrarily reduce a biometric system's FMR for higher security—at the cost of the increased inconvenience to users resulting from a higher FNMR. (Similarly, longer PINs or passwords also increase security while inconveniencing the users who must remember and type them correctly.)

Certain commercial applications would prefer to operate biometric systems in identification mode because of the added convenience of not requiring each user to claim an identity. Usually, we see speed as the biggest problem in scaling up an identification application, but identification accuracy actually scales even worse than speed. Consider an identification application with 10,000 users. We can certainly find a combination of a fast fingerprint-matching algorithm and special-purpose hardware capable of making an identification in a few seconds. On the other hand, a matching algorithm with a verification FMR of 0.001 percent will have an identification FMR_N of $10,000 \times 0.001$ percent, or 10 percent. In other words, an impostor has a good chance of gaining access to the system simply by using all 10 fingers on her two hands! Therefore, although small- to medium-scale commercial applications (for example, those with a few hundred users) might still use single-biometric identification, the only obvious solution for building accurate identification systems for large-scale applications appears to be *multimodal-biometric systems* (for example, requiring multiple fingerprints, a face and fingerprint, or some other combination, from each user).³

Risks of stolen biometrics

Instead of launching a brute-force attack, a hacker might use a very specific target and present the system with a copy of a known person's biometric sample. InterGov (www.intergov.org) reports that insiders commit about

80 percent of all cybercrimes (an assessment based only on reported security breaches). In such cases, the individual breaching the system's security very likely knows an authorized user personally, can acquire a sample biometric (for example, a latent fingerprint), can make a duplicate (such as a three-dimensional mold of the fingerprint) and present it to the biometric system.

Let us analyze this threat. To lift a latent fingerprint, the hacker must know the legitimate user's whereabouts and the surfaces she has touched. Next, the hacker must lift a latent fingerprint of good quality. This is not easy in practice because most latent fingerprints we leave are incomplete, wrapped around irregular surfaces, or partially canceled by fingers slipping. Then, the hacker has to make an accurate three-dimensional model of the finger as shown in Figure 5 (a simple two-dimensional fingerprint image cannot deceive a typical live-scan fingerprint sensor);³ this requires both expertise and laboratory equipment such as a high-resolution scanner, a three-dimensional printing device (such as a stereo-lithography printer), and so on. Recently, Tsutomu Matsumoto and colleagues documented several detailed methods of creating a fake finger from silicone and gelatin to fool many commercially available fingerprint sensors.⁶ Although producing a gummy clone of an available real finger (from a consenting user) is relatively simple, reconstructing a fake finger from a latent fingerprint remains quite complicated. Additionally, a single fake finger cannot serve in attacks against multiple users. Finally, creating a fake finger is as difficult the second or third time as it was the first time, and thus there are no economies of scale in repeating the attack. In fact, a fake-biometric attack on a biometric-based network access application presents a much smaller risk than an attack on a password-based system. This is because a hacker could launch an attack against a password-based network access application remotely, without knowing any of the users. Also, the hacker could use the same password (for example, a dictionary word) to launch an attack against all the enrolled users at no extra cost.

Validity detection and multimodal-biometrics for increased security

Many commercial applications could improve their personal recognition systems' security by adding required credentials or building blocks—for example, using a token or password together with biometric recognition. However, in high-security applications (such as access control to nuclear energy facilities), it is important that each component of the recognition system is secure in itself and that the many components provide additional layers of security. In many commercial applications, adding more credentials (such as passwords and tokens) can be undesirable because doing so reintroduces the problems associated with knowledge- and possession-

History of biometrics

For thousands of years, humans have used body characteristics such as face, voice, gait, and so on to recognize each other. In the mid 19th century, Alphonse Bertillon, chief of the criminal identification division of the police department in Paris, developed and then practiced the idea of using various body measurements (for example, height, length of arms, feet, and fingers) to identify criminals. In the late 19th century, just as his idea was gaining popularity, it was eclipsed by a far more significant and practical discovery: the distinctiveness of human fingerprints. Soon after this discovery, many major law-enforcement departments embraced the idea of “booking” criminals’ finger-

prints and storing them in databases (initially, card files). Later, police gained the ability to “lift” leftover, typically fragmentary, fingerprints from crime scenes (commonly called *latents*) and match them with fingerprints in the database to determine criminals’ identities.

Biometrics first came into extensive use for law-enforcement and legal purposes—identification of criminals and illegal aliens, security clearances for employees in sensitive jobs, paternity determinations, forensics, positive identifications of convicts and prisoners, and so on. Today, however, many civilian and private-sector applications are increasingly using biometrics to establish personal recognition.

based systems (passwords can be forgotten or guessed, and tokens can be lost or stolen). In these applications, fake biometric attacks remain a serious concern. However, we can address this threat in two ways: first, by building *vitality* detection mechanisms in the biometric recognition system hardware and software; second, by designing multimodal-biometric systems that incorporate several different biometric characteristics (for example, face, fingerprint, and hand geometry).

Fingerprint devices can incorporate vitality detection by measuring optical, electric, or thermal properties of the human skin or other biomedical characteristics such as pulse. Iris-recognition devices can measure the involuntary pupillary *hippus* (constant small constrictions and dilations of the pupil caused by spontaneous movements of the iris rather than external stimulation) to ensure that the eye is alive. The resources required to defeat biometric sensors increase as they incorporate more methods of vitality detection. If the hacker knows the vitality detection method being used, however, she can probably thwart it. For example, if the fingerprint sensor also uses finger pulse detection, the hacker can build a pulse generator into the fake finger.

Therefore, in our opinion, the best method for vitality detection is to use a characteristic distinctive to each individual, and not easily available to an adversary for copying—that is, another biometric. For example, we could build a multimodal-biometric system that combines a strong biometric, such as a fingerprint, with another biometric (possibly a weaker one) that is difficult to acquire covertly, such as a face thermogram. Another approach might use a fingerprint system requiring each user to present several fingers in a specific order; a hacker would have to find latent fingerprints from multiple fingers from both hands and also know the order of presentation. Such a solution comes at a relatively low cost: It requires no other sensor, and an existing fingerprint-verification system could be easily adapted to handle impressions from multiple fingers. Interestingly, such multimodal-biometric

solutions could also significantly improve recognition accuracy—at the cost of longer acquisition and processing times (and possibly extra hardware).

Replacing compromised biometrics

One disadvantage of biometrics is that they cannot be easily revoked.⁷ If a biometric is ever compromised, it is compromised forever. With a credit card, the bank can issue the user a new card with a new number. But a user has only a limited number of biometrics—one face, 10 fingers, and so on—and they are not easy to replace. Also, because different applications might use the same biometric, a thief who acquires a person’s biometric in one application could also use it in others.

However, integrating cryptographic techniques along with biometric matchers can help address this problem.⁸ For example, instead of storing the original biometric signal in the system database during enrollment, the system could store only its noninvertible transformed version (for example, a hash). During recognition, the biometric sensor would transform the signal using the same noninvertible transform and perform matching in the transformed space. Different applications can use different noninvertible transforms (or different parameters of the same transform), so a template would be usable only by the application that created it. In fact, the user herself could provide the transform’s parameters in terms of a password or PIN. If a hacker ever compromises such a biometric template, the system can issue a new one using a different transform or different parameters. The disadvantage of this technique is that invertible and simple noninvertible transforms are not very strong (in the information-theoretic sense), and using strong noninvertible transforms lowers system accuracy because the matcher cannot effectively deal with all the biometric signal variations (of the same person) in the transformed space.⁸ A simple and effective method of creating an easily revocable biometric template is to encrypt the biometric template with the user’s password.

Ultimately, in commercial applications, the decision to add or replace existing personal recognition methods with biometrics-based solutions should be based on a cost-benefit analysis. For example, is the installation and maintenance cost of a biometric-based computer login system less than the currently used password system? According to the Gartner Group, between 20 percent and 50 percent of all help desk calls are for password resets. Forrester Research states that the average help desk labor cost for a single password reset is about US\$38.

Identification is a much harder problem than verification because an identification system must perform a large number of comparisons.

Negative recognition: Government and forensic applications

Negative recognition applications, such as background-checking of employees and preventing terrorists from boarding airplanes, must perform personal recognition in identification mode. As we noted earlier, at a given level of accuracy, identification is a much harder problem than verification because an identification system must perform a large number of comparisons.

To illustrate the difference, let us suppose airport authorities are looking for the FBI's 100 most-wanted criminals (yielding a database size of 100), and that the state-of-the-art fingerprint verification system operates at 1 percent FNMR and 0.001 percent FMR. If we deployed this system in verification mode, it would fail to match the correct users 1 percent of the time and erroneously verify wrong users 0.001 percent of the time. Now consider what these numbers mean for the system deployed in identification mode. While the identification FNMR_N is about 1 percent, the identification FMR_N is about 100×0.001 percent, or 0.1 percent. This means that although the system has a 99 percent chance of catching a criminal, it produces a large number of false alarms. For example, if 200,000 people use a major US airport in one day, the system produces 200 false alarms!

If the system used faces instead of fingerprints for the identification—which might be preferable in the airport application, because cameras could acquire the faces covertly—the number of misses and false alarms would

be considerably higher: face identification systems have rather poor accuracy, especially in environments with cluttered backgrounds and varied lighting conditions. Thus, relying exclusively on automatic biometric systems for negative identification might be infeasible.

Adding traditional personal recognition tools such as passwords and PINs would not be at all useful for negative recognition. So, although biometric systems might not yet be extremely accurate in supporting large-scale identification applications, they are the only choice for negative recognition applications. Furthermore, a biometric system operating in semiautomatic mode, with a human expert examining all the alarms and making the final decision, can be quite effective. For example, if we need 100 airport security agents to manually match every person at an airport against the FBI's 100 most-wanted database, we might need only five agents to take a closer look at the 200 daily alarms the biometric system generates. In such semiautomatic applications, the biometric system only generates an alarm that calls for a closer, manual examination of the individual; an alarm does not directly translate into an arrest. In fact, the trade-off between FMR and FNMR rates in a biometric system is no different from that in any detection system—including the extant metal detectors already in use at all airports.

Other negative recognition applications, such as background checks and forensic criminal identification, can also operate in semiautomatic mode, and their use follows a similar cost-benefit analysis. For example, in attempting to match latent fingerprints, law enforcement agencies typically use an automatic fingerprint identification system (AFIS) only to narrow down the number of fingerprint matches from a few million to a few hundred for a human expert to perform. A forensic expert always makes the final decision.

In our opinion, using biometrics in negative recognition applications does not infringe on civil liberties because unless you are already in the "criminal database," the recognition system has no record of you. However, we do need appropriate legislation to protect the abuse of such systems.

Privacy and biometrics

Privacy is the ability to lead your life free of intrusions, to remain autonomous, and to control access to your personal information. As the incidence and magnitude of identity fraud increase, strong biometrics such as fingerprints will increasingly come into play for positively recognizing people; the conventional technologies—knowledge- or token-based, for example—cannot deliver this functionality.

For instance, US legislation requires strong recognition schemes such as biometrics to limit access to sensitive medical records. In some applications, developers have envisaged using biometrics for anonymous access. These applications could index sensitive individual in-

formation without explicitly specifying the user's name, and the access mechanisms would entail specific biometric-based recognition—for example, allowing access to the records if the person's left index fingerprint matches the fingerprint associated with the record. Furthermore, by requiring automated access mechanisms to go through a secure biometric system, system administrators could track all accesses to the privileged information, improving accountability for transactions within the information systems. In other words, biometrics-based accesses are less repudiable than other types of access control mechanisms. Thus, biometric identifiers—especially strong identifiers such as fingerprints—can clearly enhance the integrity of systems holding personal information.

On the other hand, several privacy concerns surround the use of biometrics for personal recognition. We have traditionally conceived of human recognition as a mutually reciprocated action between two individuals. Automatic methods of individual recognition, especially those based on biometrics, might culturally be perceived as undignified to humans. In addition, some biometrics (fingerprints, faces, and DNA) carry negative connotations because of their prevalent use in criminal investigation.

On other fronts, some object to biometrics on religious grounds, citing biblical references such as Revelation 13:16–17: “He also forced everyone, small and great, rich and poor, free and slave, to receive a mark on his right hand or on his forehead, so that no one could buy or sell unless he had the mark, which is the name of the beast or the number of his name.” Still others have raised concerns about the hygiene of biometric sensors that require contact. Given that we routinely touch many objects touched by strangers—money, for example—this objection seems frivolous.

However, biometrics does raise three systematic privacy concerns:⁹

- *Unintended functional scope.* Because biometric identifiers are biological in origin, collectors might glean additional (possibly statistical) personal information from scanned biometric measurements. For instance, certain malformed fingers might be statistically correlated with certain genetic disorders. With the rapid advances in human genome research, fear of inferring further information from biological measurements might also be on the rise. Such derived medical information could become a basis for systematic discrimination against segments of the population perceived as “risky.”
- *Unintended application scope.* Strong biometric identifiers such as fingerprints allow the possibility of unwanted identifications. For instance, persons legally maintaining aliases (say, for safety reasons) could be identified based on their fingerprints. In addition, bio-

metric identifiers could link bits and pieces of behavioral information about individuals enrolled in widely different applications; detractors often construe this potential as a means for organizations—governmental or corporate—to accumulate power over individuals and their autonomy.

- *Covert recognition.* Biometric characteristics are not secrets. It is often possible to obtain a biometric sample, such as a person's face, without that person's knowledge. This permits covert recognition of previously enrolled people. Consequently, those who desire to remain anonymous in any particular situation could be denied their privacy by biometric recognition.

We can address the possible abuse of biometric information (or its derivatives) and related accountability procedures in several ways:

- legislation by governments and the public—for example, European Union legislation against sharing biometric identifiers and personal information;¹⁰
- assurance of self-regulation—for example, a consortium of biometric vendors could choose to adhere to a set of ethical guidelines in their product design; and
- autonomous enforcement by independent regulatory organizations—for example, a central biometrics authority.

Until we reach consensus on the proper limits to biometrics use, many people will surely hesitate to provide either raw or processed biometric measurements to centralized applications and to untrustworthy applications that might share the data with other applications. As a result, applications delivering highly decentralized recognition capabilities will be the most acceptable for the time being.

Those who desire to remain anonymous in any particular situation could be denied their privacy by biometric recognition.

One way to decentralize a biometric system is to store the biometric information not in a centralized server but in decentralized, encrypted databases, over which the individual has complete control. For instance, a system could issue the user a smart card with her fingerprint template stored on it. In addition, if the smart card were to integrate onboard a small fingerprint sensor and pro-

cessing power to perform feature extraction and matching, it could compare the input fingerprint retrieved from the sensor directly with the resident template and deliver the decision, possibly encrypted, to the outside world. Such a decentralized system would permit all the advantages of biometric-based recognition without many of its stipulated privacy problems.

Reliable personal recognition is critical to many business processes. Because conventional knowledge- and token-based methods rely on surrogate representations of a person's identity to establish personal recognition, any system assuring reliable positive personal recognition must necessarily involve a biometric component. In fact, a sound personal recognition system design must incorporate many biometric and nonbiometric components.

Biometric-based systems also have limitations with adverse implications for a system's security. For example, the accuracy of current biometric systems is not perfect, and elaborate spoofing attacks can defeat a practical biometric system. Although the evolution of biometric technology will surely overcome some of these limitations, it is important to understand that foolproof personal recognition systems simply do not exist—and perhaps never will. Security is a risk-management strategy that identifies, controls, eliminates, or minimizes uncertain events that can adversely affect system resources and information assets. A system's security requirements depend on the application's requirements (the threat model) and the cost-benefit analysis. In our opinion, properly implemented biometric systems are effective deterrents to perpetrators.

Finally, the use of biometrics indeed raises several privacy concerns. A sound trade-off between security and privacy might be necessary; but we can only enforce collective accountability and acceptability standards through common legislation. On the positive side of the privacy issue, biometrics provides tools to enforce accountable logs of system transactions and to protect individuals' right to privacy.

As biometric technology matures, interaction will increase among applications, the market, and the technology. The technology's value, user acceptance, and the service provider's credibility will influence this interaction. It is too early to predict where and how biometric technology will evolve and which applications will ultimately embed it. But it is certain that biometric-based recognition will profoundly influence the way we conduct our daily business. □

References

1. A.K. Jain, R. Bolle, and S. Pankanti, eds., *Biometrics: Personal Identification in a Networked Society*, Kluwer Academic Publishers, 1999.
2. *Best Practices in Testing and Reporting Biometric Device Performance*, version 2.0, tech. report, United Kingdom Biometric Working Group, 2002; www.cesg.gov.uk/technology/biometrics.
3. D. Maltoni et al., *Handbook of Fingerprint Recognition*, Springer, 2003.
4. *Password Clues, The CentralNic Password Survey Report*, CentralNic, 13 July 2001; www.centralnic.com/page.php?pid=73.
5. D. Maio et al., "FVC2002: Second Fingerprint Verification Competition," *Proc. Int'l Conf. Pattern Recognition*, vol. 3, IEEE CS Press, 2002, pp. 811–814.
6. T. Matsumoto et al., "Impact of Artificial Gummy Fingers on Fingerprint Systems," *Proc. SPIE, Optical Security and Counterfeit Deterrence Techniques IV*, vol. 4677, Int'l Soc. for Optical Engineering, 2002, pp. 275–289.
7. B. Schneier, "Inside Risks: The Uses and Abuses of Biometrics," *Comm. ACM*, vol. 42, no. 8, Aug. 1999, p. 136.
8. A. Jules and M. Sudan, "A Fuzzy Vault Scheme," *Proc. IEEE Int'l Symp. Information Theory*, IEEE Press, 2002, p. 408.
9. J.D. Woodward, "Biometrics: Privacy's Foe or Privacy's Friend?" *Proc. IEEE*, vol. 85, no. 9, Sept. 1997, pp. 1480–1492.
10. *European Data Directive 95/46/EC*, Feb. 1995; www.privacy.org/pi/intl_orgs/ec/final_EU_Data_Protection.html.

Salil Prabhakar leads the Algorithms Research Group at Digital Persona. His research interests include pattern recognition, image processing, computer vision, machine learning, data mining, document analysis, biometrics, and multimedia applications. He holds a BTech from the Institute of Technology, Banaras Hindu University, Varanasi, India, and a PhD from Michigan State University, both in computer science and engineering. Contact him at salilp@digitalpersona.com.

Sharath Pankanti is with the Exploratory Computer Vision and Intelligent Robotics Group at the IBM T.J. Watson Research Center. For the past few years he has been working on analysis and interpretation of video depicting human activities. Previously, he worked on the Advanced Identification Solutions Project, dealing with reliable and scalable fingerprint recognition systems. His research interests include biometrics, pattern recognition, computer vision, and human perception. He holds a PhD from Michigan State University in computer science and engineering. Contact him at sharat@us.ibm.com.

Anil K. Jain is a University Distinguished Professor in the Department of Computer Science and Engineering at Michigan State University. He works in the areas of statistical pattern recognition, exploratory pattern analysis, Markov random fields, texture analysis, 3D object recognition, document image analysis, and biometric authentication. Jain is a fellow of both the IEEE and International Association of Pattern Recognition (IAPR). He has received the Fulbright research award, a Guggenheim fellowship, and the Alexander von Humboldt research award, and he delivered the 2002 Pierre Devijver lecture sponsored by the IAPR. Contact him at jain@cse.msu.edu.