CHAPTER

13

Biometrics

And the Gileadites took the passages of Jordan before the Ephraimites: and it was so, that when those Ephraimites which were escaped said. Let me go over; that the men of Gilead said unto him, Art thou an Ephraimite? If he said, Nay; Then said they unto him, Say now Shibboleth: and he said Sibboleth: for he could not frame to pronounce it right. Then they took him, and slew him at the passages of the Jordan: and there fell at that time of the Ephraimites forty and two thousand.

—JUDGES 12:5–6

13.1 Introduction

The above quotation may be the first recorded military use of a security protocol in which the authentication relies on a property of the human being—in this case his accent. (There had been less formal uses before this, as when Isaac tried to identify Esau by his bodily hair, but got deceived by Jacob; or indeed when people recognized each other by their faces, which I'll discuss later.)

Biometrics identify people by measuring some aspect of individual anatomy or physiology (such as your hand geometry or fingerprint), some deeply ingrained skill, or other behavioral characteristic (such as your handwritten signature), or something that is a combination of the two (such as your voice).

Over the last quarter century or so, people have developed a large number of biometric devices; this rapidly growing market is now worth about $50 million a year [414]. Earlier I mentioned the use of hand geometry to identify staff at a nuclear reactor in the late 1970s. But the best established biometric techniques predate the computer age altogether—namely the use of handwritten signatures, facial features, and fingerprints. We will look at these first, then go on to the fancier, more high-tech techniques.

13.2 Handwritten Signatures

Handwritten signatures had been used in classical China, but carved personal seals were considered to be higher status, and are still used for serious transactions in China, Japan, and Korea to this day. Europe was the other way around: seals had been used in medieval times, but as writing spread after the Renaissance, people increasingly just wrote their names to signify assent to business and other documents. Over time, the signature became accepted as the standard way of doing this in the West. Every day, billions of dollars' worth of contracts are concluded by handwritten signatures on documents, and how these can be replaced by electronic signatures is a hot policy and technology issue.

How secure are handwritten signatures?

The probability that a forged signature will be accepted as genuine mainly depends on the amount of care taken when examining it. Many bank card transactions in stores are accepted without even a glance at the specimen signature on the card—so much so that many Americans do not even bother to sign their credit cards. (This can cause problems when traveling in more punctilious countries such as Germany or Switzerland.) But even diligent signature checking doesn't reduce the risk of fraud to zero. An experiment showed that 105 professional document examiners, who each did 144 pairwise comparisons, misattributed 6.5% of documents. Meanwhile, a control group of 34 untrained people of the same educational level got it wrong 38.3% of the time [431], and the nonprofessionals' performance couldn't be improved by giving them monetary incentives [432]. Errors made by professionals are a subject of continuing discussion in the industry, but are thought to reflect the examiner's assumptions and preconceptions [81]. As the participants in these tests were given reasonable handwriting samples rather than just a signature, it seems fair to assume that the results for verifying signatures on checks or credit card vouchers would be significantly worse.

So handwritten signatures are surrounded by a number of conventions and special rules which vary from one country to another. For example, to buy a house in England using money borrowed from a bank of which you're not an established customer, the procedure is to go to a lawyer's office with a document such as a passport, sign the property transfer and loan contract, and get the contract countersigned by the lawyer. The requirement for government-issued photo-ID is imposed by the mortgage lender to keep its insurers happy, while the requirement that a purchase of real estate be in writing was imposed by the government some centuries ago in order to collect stamp duty on property transactions. Other types of document (such as expert testimony) may have to be notarized in particular ways. Many curious anomalies go back to the nineteenth century, and the invention of the typewriter. Some countries require that machine-written contracts be initialed on each page, while some don't; and these differences have sometimes persisted for over a century. Clashes in conventions still cause serious problems. In one case, a real estate transaction in Spain was held to be invalid because the deal had been concluded by fax, and a U.K. company went bust as a result.

In most of the English-speaking world, however, most documents do not need to be authenticated by special measures. The essence of a signature is the intent of the signer, so an illiterate's "*X*" on a document is just as valid as a monarch's flourish. In fact, a plaintext name at the bottom of an email message also has just as much legal force [810], except where there are specific regulations requiring the transaction to be in writing. There may be thousands of such in each jurisdiction. Meanwhile, it's actu-

ally very rare for signatures to be disputed in court cases, as the context generally makes it clear who did what. So we have a very weak biometric mechanism that works quite well in practice—except that it's choked by procedural rules that vary by country and by application.

Sorting out this mess, and imposing reasonably uniform rules for electronic documents, is a subject of much international activity. A summary of the issues can be found in [811], with an analysis by country in [68]; and I'll discuss some of the issues further in Part 3. For now, note that the form of a signature, the ease with which it can be forged, and whether it has legal validity in a given context, are largely independent questions.

There is one application, though, where effective automatic recognition of handwritten signatures could be very valuable. This is check clearing.

In a bank's check processing center, it is typical practice that you only verify signatures on checks over a certain amount—perhaps $1,000, perhaps $10,000, perhaps a percentage of the last three months' movement on the account. The signature verification is done by an operator who sees, simultaneously presented on-screen, the check image and the customer's reference signature.

Verifying checks for small amounts is not economic unless it can be automated, so a number of researchers have worked on systems to compare handwritten signatures automatically. This turns out to be a very difficult image-processing task because of the variability between one genuine signature and another. A much easier option is to use a *signature tablet.* This is a sensor surface on which the user does a signature; it records not just the shape of the curve but also its dynamics (the velocity of the hand, where the pen was lifted off the paper, and so on). Tablets are used to identify users in some high-value applications, including securities dealing.

Like alarm systems, most biometric systems have a trade-off between false accept and false reject rates, often referred to in the banking industry as the *fraud* and *insult* rates, and in the biometric literature as *type 1* and *type 2* errors. Many systems can be tuned to favor one over the other. The *equal error rate* is when the system is tuned so that the probabilities of false accept and false reject are equal. For common signature recognition systems, the equal error rate is about 1%. This is not fatal in an operation such as a bank dealing room. If one of the dealers tries to log on one morning and his PC rejects his signature, he can just try again. If there is a persistent failure, he can call the system administrator and have the machine reset. However, it is a show-stopper in a retail store. If one transaction in a hundred fails, the aggravation to customers would be unacceptable. So U.K. banks set a target for biometrics of a fraud rate of 1% and an insult rate of 0.01%, which is beyond the current state of the art in signature verification [317].

What can be done to bridge the gap? An interesting experiment was conducted by the University of Kent, England, to cut fraud by welfare claimants who were drawing their benefits at a post office near Southampton. The novel feature of this system is that it was used to screen signatures and to support human decisions, rather than to take decisions itself. So instead of being tuned for a low insult rate, with a correspondingly high fraud rate, it had fraud and insult rates approximately equal. When a signature was rejected, this merely told the staff to look more closely, and to ask for a driver's license or other photo ID. With 8,500 samples taken from 343 customers, 98.2% were verified correctly at the first attempt, rising to 99.15% after three attempts. The experiment was judged to be a success [282]. However, this rate was achieved by

excluding *goats*—a term used by the biometric community for people whose templates don't classify well. With them included, the false reject rate was 6.9% [283].

In general, biometric mechanisms tend to be much more robust in attended operations, where they assist a guard rather than replacing him. The false alarm rate may then actually help by keeping the guard alert.

13.3 Face Recognition

Recognizing people by their facial features is the oldest identification mechanism of all, going back at least to our early primate ancestors. Biologists believe that a significant part of our cognitive function evolved to provide efficient ways of recognizing other people's facial features and expressions [646]. For example, we are extremely good at detecting whether another person is looking at us or not. In theory, humans' ability to identify people by their faces appears to be very much better than any automatic system produced to date.

The human ability to recognize faces is also important to the security engineer because of the widespread reliance placed on photo IDs. Drivers' licenses, passports, and other kinds of identity card are not only used directly to control entry to computer rooms, but also bootstrap most other systems. The issue of a password, or a smartcard, or the registration of a user for a biometric system using some other technique such as iris recognition, is often the end point of a process which was started by that person presenting photo ID when applying for a job, opening a bank account, or whatever.

But even if people are good at recognizing friends in the flesh, how good are they at identifying strangers by photo ID?

The simple answer is that they're not. Psychologists at the University of Westminster conducted a fascinating experiment with the help of a supermarket chain and a bank [450]. They recruited 44 students and issued each of them with four credit cards each with a different photograph on it, as follows.

- One of the photos was a "good, good" one. It was genuine and recent.
- The second was a "bad, good one." It was genuine but a bit old; the student now had different clothing, hairstyle, or whatever. In other words, it was typical of the photo that most people have on their photo ID.
- The third was a "good, bad one." From a pile of a hundred or so random photographs of different people, investigators chose the one that most looked like the subject. In other words, it was typical of the match that criminals could get if they had a stack of stolen cards.
- The fourth was a "bad, bad" one. It was chosen at random except that it had the same sex and race as the subject. In other words, it was typical of the match that really lazy, careless criminals would get.

The experiment was conducted in a supermarket after normal business hours, but with experienced cashiers on duty who were aware of the purpose of the experiment. Each student made several trips past the checkout using different cards. It transpired that none of the checkout staff could tell the difference between "good, bad" photos and "bad, good" photos. In fact, some of them could not even tell the difference between "good, good" and "bad, bad." As this experiment was done under optimum conditions—with experienced staff, plenty of time, and no threat of embarrassment or

violence if a card was rejected—real-life performance can be expected to be worse. (In fact, many stores do not pass on to their checkout staff the reward offered by credit card companies for capturing stolen cards, so even the basic motivation may be absent.)

The response of the banking industry to this experiment was ambivalent. At least two banks that had experimented with photos on credit cards had experienced a substantial drop in fraud—to less than one percent of the expected amount in the case of one Scottish bank [67]. The overall conclusion was that the benefit to be had from photo ID is essentially its deterrent effect [293].

The extreme difficulty of getting people to use their facial recognition skills effectively is one of the reasons for trying to automate the process. Attempts go back to the nineteenth century, when Galton devised a series of spring-loaded "mechanical selectors for facial measurements [328]. But automated face recognition actually subsumes a number of separate problems. In identity verification, the subject looks straight at the camera under controlled lighting conditions, and their face is compared with the one on file. A related but harder problem is found in forensics, where we may be trying to establish whether a suspect's face fits a low-quality recording on a security video. The hardest of all is surveillance, where the goal may be to scan a moving crowd of people at an airport and try to pick out anyone who is on a list of perhaps a few hundred known suspects.

Even picking out faces from an image of a crowd is a nontrivial computational task [502]. A recent empirical study of the robustness of different facial feature extraction methods found that, given reasonable variations in lighting, viewpoint, and expression, no method was sufficient by itself, and error rates were up to 20% [10]. Systems that use a combination of techniques can get the error rate down, but not to the 1% or less which is possible with many other biometrics [556, 818].

In short, the technology still does not work very well, when viewed solely in terms of error rates. However, from the system viewpoint, it can work very well indeed. In 1998, the London borough of Newham placed video cameras prominently in the high street and ran a PR campaign about how their new computer system constantly scanned the faces in the crowd for several hundred known local criminals. They managed to get a significant reduction in burglary, shoplifting, and street crime. The system even worries civil libertarians—despite the fact that it appears to work primarily by deterrence [739]. Of course, as time passes and technology improves, both the potential and the worries may increase.

13.4 Fingerprints

Fingerprints are important. By 1998, fingerprint recognition products accounted for 78% of the total sales of biometric technology. These products look at the friction ridges that cover the fingertips and classify patterns of *minutiae*, such as branches and end points of the ridges. Some also look at the pores in the skin of the ridges. A technical description of the leading automatic fingerprint identification systems can be found in [496].

The use of fingerprints to identify people was discovered independently a number of times. Mark Twain mentioned thumbprints in 1883, in *Life on the Mississippi*, where he claims to have learned about them from an old Frenchman who had been a prison-keeper. Long before that, they were accepted in a seventh-century Chinese legal code as an alternative to a seal or a signature; and they were required by an eighth-century Japanese code when an illiterate man wished to divorce his wife. They were mentioned in work by Malpighi in Italy in the seventeenth century; and used in 1691 by 225 citizens of Londonderry in Ireland to sign a petition asking for reparations following the siege of the city by King William.

The first modern systematic use appears to have been in India during the midnineteenth century, when William Herschel (grandson of the astronomer) was a colonial official in Hooghly. He used fingerprints to stop impersonation of pensioners who had died, and to prevent rich criminals paying poor people to serve their jail sentences for them. Henry Faulds, a medical missionary in Japan, discovered them independently in the 1870s and brought them to the attention of Darwin, who in turn motivated Galton to work out a scheme for classifying their patterns. His classification, of *loops, whorls, arches*, and *tents*, is still in use today.

According to the English-language version of history, fingerprints passed into mainstream police use in 1900, when a former police chief from Bengal, Edward Henry, became Commissioner of the Metropolitan Police in London.[1] Henry's contribution was to develop Galton's classification into an indexing system known as *binning*. By assigning one bit to whether or not each of a suspect's 10 fingers had a whorl—a type of circular pattern—he divided the fingerprint files into 1,024 bins. In this way, it was possible to reduce the number of records that have to be searched by orders of magnitude.

Fingerprints are now used by the world's police forces for essentially two different purposes. In the United States, their main use is in identification. FBI files are used to check out arrested suspects to determine whether they're currently wanted by other law enforcement agencies. They are also used to screen job applicants; for example, anyone wanting a U.S. government clearance at Secret or above must have an FBI fingerprint check. They are also used in crime scene forensics. In Europe, where people carry identity cards and identity is thus more readily established, forensics provide the main application.

Fingerprints found at a crime scene are matched against database records. Prints that match to more than a certain level are taken as hard evidence that a suspect visited the crime scene, and are often enough to secure a conviction on their own. In some countries, fingerprints are required from all citizens and all resident foreigners.

To cut the costs of manual fingerprint matching, a number of automated systems have been developed. Algorithms suitable for the image-processing step are surveyed in [522], and there is a tutorial plus a description of an IBM system in [415]. While some of these systems simply replace the previous manual classification and matching

[1] In the Spanish version, they were first used in Argentina where they secured a murder conviction in 1892; while Cuba, which set up its fingerprint bureau in 1907, beat the United States, whose first conviction was in Illinois in 1911. The Croation version notes that the Argentinian system was developed by one Juan Vucetich, who had emigrated from Dalmatia. The German version refers to Professor Purkinje of Breslau, who wrote about fingerprints in 1828. Breslau is now Wroclaw in Poland, so the Poles have a story too. Indians point to the bureau established in Calcutta in 1898. Success truly has many fathers!

process, or aim to improve on it [779], others use fingerprint reading devices to authenticate people in real time for applications such as building entry control and benefit payment [258]. They are also used in banking systems in countries such as India and Saudi Arabia, where the use of ink fingerprints was already common thanks to the large proportion of people who are without formal education.

They have not really taken off in banking systems in North America or Europe because of the association with crime, though a few U.S. banks do ask for fingerprints if you cash a check there and are not a customer. They find this cuts check fraud by about a half. Some have gone as far as fingerprinting new customers, and found that customer resistance is less than expected, especially if they use scanners rather than ink and paper [314]. Again, the effect is largely deterrent: matching a single print against the whole FBI database is much harder than typical crime scene work, where the suspects are the hundred or so locally active burglars. Nonetheless, there have been moves to ban the use of fingerprints in U.S. banking as a violation of privacy.

So how good is fingerprint recognition? The error rate in forensic applications can be very low, the limitation being the size and quality of the image taken from the crime scene. It varies from one country to another, depending on police procedures. Britain traditionally required that fingerprints match in 16 *points* (corresponding minutiae), and a U.K. police expert estimated that this will only happen by chance somewhere between one in four billion and one in ten billion matches [485]. Greece accepts 10 matching minutiae, Turkey 8; the United States has no set limit (it certifies examiners instead). This means that in the United States, matches can be found with poorer quality prints, but they can be open to doubt. In Britain, fingerprint evidence went for almost a century without a successful challenge; in the United States, challenges do succeed from time to time, and disputes between rival experts are not unknown.

A recent case has upset the traditional U.K. complacency [538]. Shirley McKie, a Scottish policewoman, was prosecuted on the basis of a fingerprint match on the required 16 points, verified by four examiners of the Scottish Criminal Records Office. The defense called two American examiners who presented testimony that it is not an identification.

McKie was acquitted and, as no indication was made as to whether the jury concurred with the foreign experts or merely considered their testimony as negating the Scottish experts, the Scottish Criminal Records Office asserted for over a year that this was a valid identification. But by June 2000, the matter had gone as far as the Scottish Parliament, and the justice minister himself had to climb down. The problem appears to have been that if they accepted that the fingerprint was not Shirley's, they might also have to release one David Asbury who had been convicted of murder in that case. His fingerprint identification is now also being questioned by experts and an appeal on his behalf is underway [334].

Four comments are in order here.

- Even if the probability of a false match on 16 points is one in ten billion (10^{-10}) as claimed by the police, once many prints are compared against each other, probability theory starts to bite. A system that worked well in the old days, whereby a crime scene print would be compared manually with the records of 57 known local burglars, breaks down once thousands of prints are compared every year with an online database of millions. It was inevitable that, sooner or later, enough matches would have been done to find a 16-point mismatch. Indeed, as most people on the fingerprint database are petty crimi-

nals who will not be able to muster the kind of resolute defense that McKie did, I wouldn't be surprised if there had already been other wrongful convictions.

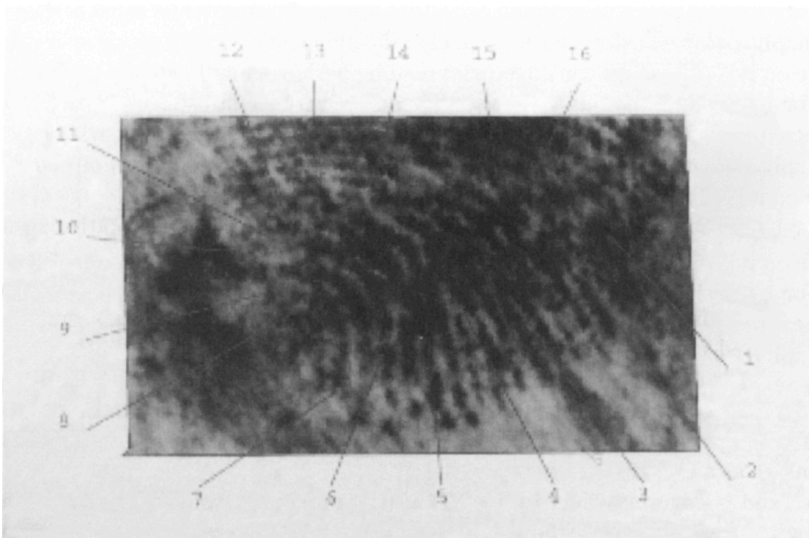

Figure 13.1 Crime scene print.

- As Figure 13.1 should make clear, fingerprint impressions are often very "noisy," being obscured by dirt, so mistakes are quite possible. The skill (and prejudices) of the examiner enter into the equation in a much more significant way than a naive jury might think. The errors caused by noise can manifest themselves at more than one level. For example, binning error rates are believed to cause a false reject rate of several percent [154].
- The belief that any security mechanism is infallible generates the complacency and carelessness needed to undermine its proper use. No consideration appears to have been given to increasing the number of points required from 16 to, say, 20, with the introduction of computer matching. Sixteen was tradition, the system was infallible, and there was certainly no reason to make public funds available for defendants to hire their own experts. In fact, as all the U.K. experts are policemen or former policemen, there are no independent experts available for hire.
- A belief of infallibility ensures that the consequences of the eventual failure will be severe. As with the Munden case described in Section 9.4.3, which helped torpedo claims about cash machine security, an assumption that a security mechanism is infallible causes procedures, cultural assumptions, and even laws to spring up which ensure that its eventual failure will be denied for as long as possible, and may have disastrous effects for the individuals involved.

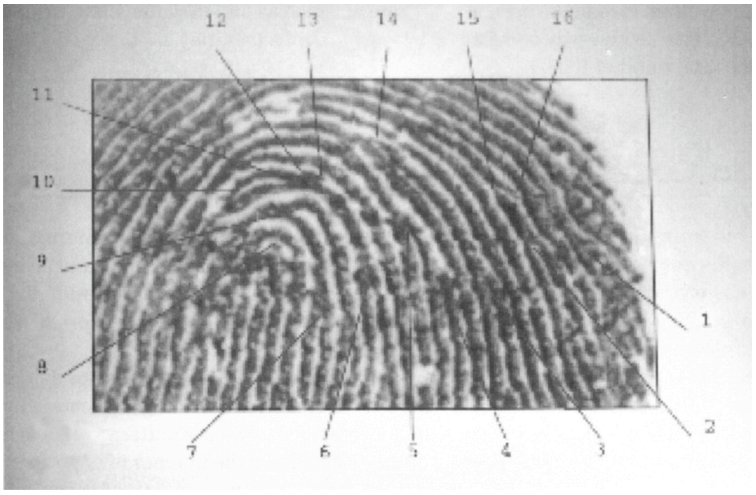

Figure 13.2 Inked print.

However, even when we do have a correct match (with 20, or 24, or however many points), its implications are not entirely obvious. It is possible for fingerprints to be transferred using adhesive tape, or for molds to be made—even without the knowledge of the target—using techniques originally devised for police use. So it is possible that the suspect whose print is found at the crime scene was framed by another criminal (or by the police—most fingerprint fabrication cases involve law enforcement personnel rather than other suspects [110]). Of course, even if the villain wasn't framed, he can always claim that he was and the jury might believe him.

Moving now to automated identification, the better systems have an equal error rate which seems to be somewhat below 1%. Although in theory the false accept probability can be made arbitrarily small, in practice false accepts happen because of features incorporated to reduce the false reject rate—such as allowance for distortion and flexibility in feature selection [650].

Fingerprint damage can also impair recognition. When I was a kid, I slashed my finger while cutting an apple, and this left a scar about half an inch long on my left middle finger. When I presented this finger to the system used in 1989 by the FBI for building entry control, my scar crashed the scanner. (It was registered and worked OK with the successor system from the same company when I tried again 10 years later.) But even where scars don't cause gross system malfunctions, they still increase the error rate. A number of people, such as manual workers and pipe smokers, damage their fingerprints frequently; and both the young and the old have faint prints [171]. Automated systems also have problems with amputees, people with birth defects such as extra fingers, and the (rare) people born without conventional fingerprint patterns at all [485].

Perhaps the most important aspect of fingerprint systems is not their error rate, as measured under laboratory conditions, but their deterrent effect.

This is particularly pronounced in welfare payment systems. Even though the fingerprint readers used to authenticate welfare claimants have an error rate as much as 5% [163], they have turned out to be such an effective way of reducing the welfare rolls that they are being adopted in one place after another [553].

13.5 Iris Codes

We turn now from the very traditional ways of identifying people to the modern and innovative. Recognizing people by the patterns in the irises of their eyes is far and away the technique with the best error rates of automated systems when measured under lab conditions. It appears to be the most secure possible way of controlling entry to premises such as plutonium stores.

As far as is known, every human iris is measurably unique. It is fairly easy to detect in a video picture, does not wear out, and is isolated from the external environment by the cornea (which in turn has its own cleaning mechanism). The iris pattern contains a large amount of randomness, and appears to have many times the number of degrees of freedom of a fingerprint. It is formed between the third and eighth month of gestation, and (like the fingerprint pattern) is *phenotypic* in that there appears to be limited genetic influence; the mechanisms that form it appear to be chaotic. So the patterns are different even for identical twins (and for the two eyes of a single individual), and they appear to be stable throughout life.

A signal processing technique (Gabor filters) has been found which extracts the information from an image of the iris into a 256-byte *iris code*. This involves a circular wavelet transform taken at a number of concentric rings between the pupil and the outside of the iris (Figure 13.3), and has the beautiful property that two codes computed from the same iris will typically match in 90% of their bits [218]. This is much simpler than in fingerprint scanners where orienting and classifying the minutiae is a hard task. The speed and accuracy of iris coding has led to a number of commercial iris recognition products [794]. Iris codes provide the lowest false accept rates of any known verification system—zero, in tests conducted by the U.S. Department of Energy. The equal error rate has been shown to be better than one in a million, and if one is prepared to tolerate a false reject rate of one in ten thousand, then the theoretical false accept rate would be less than one in a trillion.

The main practical problem facing deployment of iris scanning in the field is getting the picture without being too intrusive. The iris is small (less than half an inch) and an image including several hundred pixels of iris is needed. A cooperative subject can place his eye within a few inches of a video camera, and the best standard equipment will work up to a distance of two or three feet. Cooperation can be assumed with entry control to computer rooms, but it is less acceptable in general retail applications, as some people find being so close to a camera uncomfortable. There's no technical reason why a camera could not acquire the iris from a distance of several feet given automatic facial feature recognition, pan and zoom—it would just cost a bit more—but that brings Orwellian overtones of automatic recognition of individuals passing in a crowd. (In Europe, data protection law would be a potential show-stopper.) Secondary problems include blinking, eyelashes obscuring the eye, and sunglasses.

Possible attacks on iris recognition systems include—in unattended operation at least—a simple photograph of the target's iris. This may not be a problem in entry

control to supervised premises, but if everyone starts to use iris codes to authenticate bank card transactions, then your code will become known to many organizations. As iris codes can be compared rapidly (just exclusive-or them together and count the number of zero bits), they may start to assume the properties of names, rather than being passwords (as in current systems). So it might be possible to use your iris code to link together your dealings with different organizations.

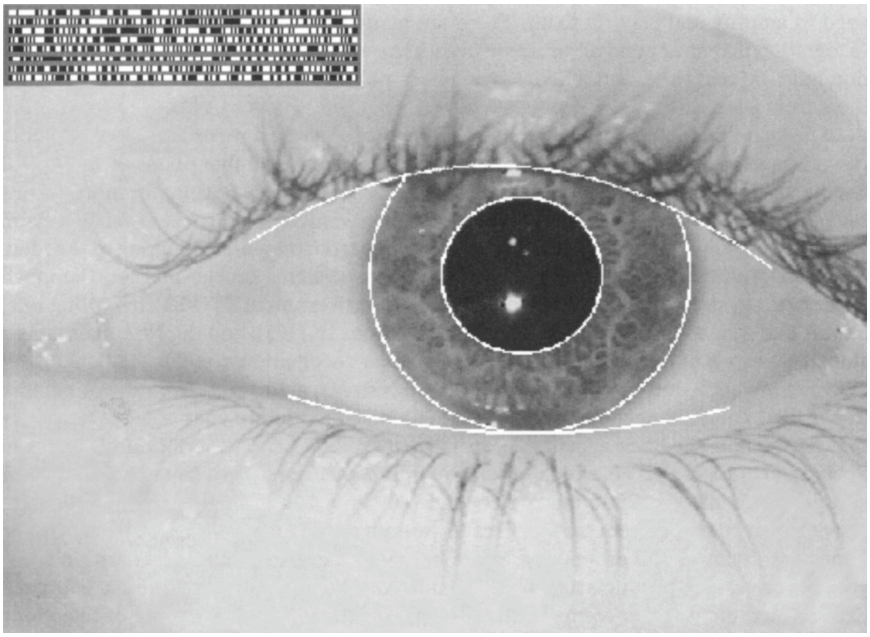

Figure 13.3 An iris with iris code (courtesy John Daugman).

A possible solution to the impersonation problem is to design terminals that measure *hippus*—a natural fluctuation in the diameter of the pupil which happens at about 0.5 Hz. But even this isn't infallible. One might try, for example, to print the target's iris patterns on contact lenses (though existing vanity contact lens printing techniques are so coarse-grained that they are detectable).

Despite the difficulties, iris codes remain a very strong contender as they can, in the correct circumstances, provide much greater certainty than any other method that the individual in question is the same as the one who was initially registered on the system. They can meet the goal of automatic recognition with zero false acceptances.

13.6 Voice Recognition

Voice recognition—also known as *speaker recognition*—is the problem of identifying a speaker from a short utterance. While *speech recognition* systems are concerned with transcribing speech and need to ignore speech idiosyncrasies, voice recognition systems need to amplify and classify them. There are many subproblems, such as whether the recognition is text-dependent or not, whether the environment is noisy, whether operation must be real time, and whether one needs only to verify speakers or to recognize them from a large set.

In *forensic phonology*, the objective is, usually, to match a recorded telephone conversation, such as a bomb threat, to speech samples from a number of suspects. Typical techniques involve filtering and extracting features from the spectrum; for more details see [461]. A more straightforward biometric authentication objective is to verify a claim to identity in some telephone systems. These range from telephone banking to the identification of military personnel, with over a dozen systems on the market. Campbell describes a system that can be used with the U.S. government STU-III encrypting telephone, and that achieves an equal error rate of about 1% [161]; and the NSA maintains a standard corpus of test data for evaluating speaker recognition systems [414].

There are some interesting attacks on these systems, quite apart from the possibility that a villain might somehow manage to train himself to imitate your voice in a manner that the equipment finds acceptable. In [324] there is a brief description of a system fielded in US EP-3 aircraft which breaks up intercepted messages from enemy aircraft and ground controllers into quarter second segments that are then cut and pasted to provide new, deceptive messages. This is primitive compared with what can now be done with digital signal processing. Some informed observers expect that within a few years, there will be products available which support real-time voice and image forgery. Crude voice morphing systems already exist, and enable female victims of telephone sex pests to answer the phone with a male sounding voice. Better ones will enable call centers to have the same 'person' always greet you when you phone. With that sort of commercial pressure driving the technology, it's only a matter of time before remote biometrics become very much harder.

13.7 Other Systems

A number of other biometric technologies have been proposed. For a survey of the market, see [553]. Some, such as those based on *facial thermograms* (maps of the surface temperature of the face, derived from infrared images), the shape of the ear, gait, lip prints, and the patterns of veins in the hand, don't seem to have been marketed as products. Other technologies may provide interesting biometrics in the future. For example, the huge investment in developing digital noses for quality control in the food and drink industries may lead to a "digital doggie," which recognizes its master by scent.

Others biometric techniques, such as typing patterns, were used in products in the 1980s but don't appear to have been successful (typing patterns, also known as keystroke dynamics, had a famous precursor in the wartime technique of identifying wireless telegraphy operators by their *fist*, the way in which they used a Morse key).

Still others, such as hand geometry, have useful niche markets. In addition to its use since the 1970s in nuclear premises entry control, hand geometry is now used at airports by the U.S. Immigration and Naturalization Service to provide a "fast track" for frequent flyers. It is fairly robust, with an equal error rate under lab conditions of 0.1–0.2%. (In fact, hand geometry derives from *anthropometrics*, a system of identifying criminals by skeletal measurements, which was introduced in Paris in 1882 by Alphonse Bertillon, but replaced by fingerprints a generation later.)

One other biometric deserves passing mention—the use of DNA typing. This has become a valuable tool for crime-scene forensics and for determining parenthood in child support cases, but is too slow for applications such as building entry control. Being genotypic rather than phenotypic, its accuracy is also limited by the incidence of monozygotic twins—about one white person in 120 has an identical twin. There's also a privacy problem, in that it should soon be possible to reconstruct a large amount of information about an individual from their DNA sample. For a survey of forensic DNA analysis techniques, and suggestions of how to make national DNA databases consistent with European data protection law, see [680].

13.8 What Goes Wrong

As with other aspects of security, we find the usual crop of failures due to bugs, blunders, and complacency. The main problem faced by DNA typing, for example, was an initially high rate of false positives, due to careless laboratory procedure. This not only scared off some police forces, which had sent in samples from different volunteers and got back false matches, but also led to disputed court cases and alleged miscarriages of justice.

Biometrics are like many other protection mechanisms (alarms, seals, tamper-sensing enclosures, ...) in that environmental conditions can cause havoc. Noise, dirt, vibration, and unreliable lighting conditions all take their toll. Some systems, like speaker recognition, are vulnerable to alcohol intake and stress. Changes in environmental assumptions, such as from closed to open systems, from small systems to large ones, from attended to standalone, from cooperative to recalcitrant subjects, and from verification to identification—can all undermine a system's viability.

There are a number of more specific and interesting attacks on various biometric systems.

- There have been some attacks on the methods used to index biometric data. The classic one is the helpful villain who gives an inexperienced policeman his fingerprints in the wrong order, so that instead of the hand being indexed under the Henry system as '01101' it becomes perhaps '01011', so his record isn't found and he gets the lighter sentence due a first offender [485].
- Forensic biometrics often don't tell as much as one might assume. Apart from the possibility that a fingerprint or DNA sample might have been planted by the police, it may just be old. The age of a fingerprint can't be determined directly, and prints on areas with public access say little. A print on a bank door says much less than a print in a robbed vault. So in premises vulnerable to robbery, cleaning procedures may be critical for evidence. If a suspect's prints are found on a bank counter, and she claims to have gone there three days previously, she may be convicted by evidence that the branch counter is polished every evening. Putting this in system terms, freshness is often a critical issue, and some quite unexpected things can find themselves inside the "trusted computing base."

- Another aspect of freshness is that most biometric systems can, at least in theory, be attacked using suitable recordings. We mentioned direct attacks on voice recognition, attacks on iris scanners by photos on a contact lens, and molds of fingerprints. Even simpler still, in countries where fingerprints are used to pay pensions, there are persistent tales of "Granny's finger in the pickle jar" being the most valuable property she bequeathed to her family. This reinforces the lesson that unattended operation of biometric authentication devices is tricky.
- Certain systems—notably handwriting systems—are vulnerable to collusion. Villains can voluntarily degrade handwriting ability. By giving several slightly different childish sample signatures, they can force the machine to accept a lower threshold than usual. The kind of attack to expect is that Alice opens a bank account and her accomplice Betty withdraws money from it; Alice then complains of theft and produces a watertight alibi. As with alarm and shared control systems, commercial users have to worry about colluding employees or customers, while the military threat model is usually just the single disloyal soldier.
- Commercial system builders must also worry about false repudiation—such as whether a user who practices enough can generate two signatures that pass for identical on the signature tablet, even if they are visually quite different.
- The statistics are often not understood by system designers, and the birthday theorem is particularly poorly appreciated. With 10,000 biometrics in a database, for example, there are about 50,000,000 pairs. So even with a false accept rate of only one in a million, the likelihood of there being at least one false match will rise above one-half as soon as there are somewhat over a thousand people (in fact, 1,609 people) enrolled. So identification is a tougher task than verification [219]. The practical consequence is that a system designed for authentication may fail when you try to rely on it for evidence. A good way to explain to judges, and other non-technical people, why the system error rate differs from the single sample error rate is that there is "one chance to get it right, but N chances to get it wrong." For a good discussion of error rates see [154].
- Another aspect of statistics comes into play when designers assume that by combining biometrics they can get a lower error rate. The curious and perhaps counter-intuitive result is that a combination will typically result in improving either the false accept or the false reject rate, while making the other worse. One way to look at this is that if you install two different burglar alarm systems at your home, then the probability that they will be simultaneously defeated goes down while the number of false alarms goes up. In some cases, such as when a very good biometric is combined with a very imprecise one, the effect can be worse overall [219].
- Most biometrics are not as accurate for all people, and some of the population can't be identified as reliably as the rest (or even at all). The elderly, and manual workers, often have damaged or abraded fingerprints. People with dark-colored eyes and large pupils give poorer iris codes. Disabled people, with no fingers or no eyes, risk exclusion if such systems become widespread. Illiterates who make an "X" are more at risk from signature forgery.

Biometric engineers sometimes refer to such subjects dismissively as goats, but this is blind to political reality. A biometric system that is (or is seen to be) socially regressive—in that it puts the disabled, the poor, the old, and ethnic minorities at greater risk of impersonation—may meet with principled resistance. In fact, a biometric system might be defeated by legal challenges on a number of grounds [626]. It may also be vulnerable to villains who are (or pretend to be) disabled. Fallback modes of operation will have to be provided; if these are less secure, then forcing their use may yield an attack, and if they are at least as secure, then why use biometrics at all?

- Finally, Christian fundamentalists are uneasy about biometric technology. They find written of the Antichrist in Revelation 13:16-17: "And he causes all, both small and great, rich and poor, free and slave, to receive a mark on their right hand or on their foreheads, and that no one may buy or sell except one who has the mark or the name of the beast, or the number of his name." So biometrics can arouse political opposition on the right as well as the left.

So there are some non-trivial problems to be overcome before biometrics will be ready for mass-market use, in the way that magnetic strip cards are used at present. But despite the cost and the error rates, they have proved their worth in a number of applications, most notably where their deterrent effect is useful.

13.9 Summary

Biometric measures of one kind or another have been used to identify people since ancient times, with handwritten signatures, facial features, and fingerprints being the traditional methods. Systems have been built that automate the task of recognition, using these methods and newer ones, such as hand geometry, voiceprints, and iris patterns. These systems have different strengths and weaknesses. In automatic operation, most have error rates of the order of 1% (though iris recognition is better, hand geometry slightly better, and face recognition worse). There is always a trade-off between the false accept rate (the fraud rate) and the false reject rate (the insult rate). The statistics of error rates are deceptively difficult.

If any biometric becomes very widely used, there is increased risk of forgery in unattended operation: voice synthesizers, photographs of irises, fingerprint molds, and even good old-fashioned forged signatures must all be thought of in system design. These do not rule out the use of biometrics, as traditional methods such as handwritten signatures are usable in practice despite very high error rates. Biometrics are usually more powerful in attended operation, where, with good system design, the relative strengths and weaknesses of the human guard and the machine recognition system may complement one another. Finally, many biometric systems achieve most or all of their result by deterring criminals rather than being effective at identifying them.

Research Problems

Potentially profitable research problems relate to the design, or improvement, of biometric systems. Is it possible to build a system—other than iris scanning—that will meet the banks' goal of a 1% fraud rate and a 0.01% insult rate? Is it possible to build a static signature verification system that has a good enough error rate (say 1%) for it to be used for screening images of checks? Are there any completely new biometrics that might be useful in some circumstances?

One I thought up while writing this chapter, in a conversation with William Clocksin and Alan Blackwell, was instrumenting a car so as to identify a driver by the way in which he or she operated the gears and the clutch. This might be hooked in to a high-end car alarm system of the kind that, if your car appears to be stolen, phones a GPS fix to a control center which then calls you to check. We haven't patented this; if you can make it work, all we ask is an acknowledgment—and some thought about how to prevent insurance companies (and governments) demanding access to the data!

Further Reading

The history of fingerprints is good reading. The standard reference is Lambourne [485], while Block has a good collection of U.S. case histories [120]. In addition to the references cited for facial and handwriting recognition in the text, there's an IBM experimental system described at [433] and a survey of the literature at [181]. The standard work on iris codes is Daugman [218]. For voice recognition, there is a tutorial in [161] which focuses on speaker identification while for the forensic aspects, see Klevans and Rodman [461]. A special issue of the *Proceedings of the IEEE* on biometric systems—volume 85 no 9 (September 1997) provides a very useful snapshot of the state of the technical art. Finally, for technical detail on a range of systems, there is a book by Anil Jain, Ruud Bolle, and Sharath Pankanti which contains chapters on a number of biometric system written by their designers [414].