CHAPTER

15

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

15.1 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 behavior (such as your handwritten signature), or some combination of the two (such as your voice).

Over the last quarter century or so, people have developed a large number of biometric devices. Since 9/11 the market has really taken off, with a number of large-scale programs including the international standards for biometric travel documents, the US-VISIT program which fingerprints visitors to the USA, Europe's Schengen visa, assorted ID card initiatives, and various registered traveler programs. Some large systems already existed, such as the FBI's fingerprint database, which is now being expanded to contain a

range of biometric data for both identification and forensic purposes. The Biometric systems market was reportedly worth over $1.5bn in 2005 [675], a massive increase from $50 m in 1998 [655]. I already 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. I will look at these first, then go on to the fancier, more 'high-tech' techniques.

15.2 Handwritten Signatures

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

Handwritten signatures are a very weak authentication mechanism by themselves (in that they're easy to forge) but have worked well for centuries because of the context of their use. An important factor is the liability for forgery. UK law provides that a forged handwritten signature is completely null and void, and this has survived in the laws of many countries that were part of the British Empire at the time. It means that the risk from a forged signature falls on the party who relies on it, and it's not possible for a bank to use its standard terms and conditions to dump the risk on the customer. So manuscript signatures are better for the customer, while the PINs and electronic tokens that are now replacing them can be better for the bank. This is not the case everywhere; some Swiss banks make customers liable for forged cheques. In the USA, Regulation E makes banks liable for the electronic systems they deploy, so the introduction of electronics doesn't change the game much. Needless to say, European banks have moved much further than U.S. banks in moving customers away from handwritten signatures.

Now 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[1]. But even diligent signature checking

[1]Indeed it's not in the cardholder's interest to give a specimen signature to a thief — if the thief makes a random signature on a voucher, it's easier for the real cardholder to disown it. Signing the card is in the bank's interest but not the customer's.

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 [682], and the nonprofessionals' performance couldn't be improved by giving them monetary incentives [683]. Errors made by professionals are a subject of continuing discussion in the industry but are thought to reflect the examiner's preconceptions [137] and context [403]. 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 even worse.

So handwritten signatures are surrounded by a number of conventions and special rules that vary from one country to another, and these extend well beyond banking. 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 tax 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 initialled 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 UK 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 the flourish of an educated man. In fact, a plaintext name at the bottom of an email message also has just as much legal force [1358], except where there are specific regulations to the contrary. There may be many obscure signature regulations scattered through each country's laws.

It's actually very rare for signatures to be disputed in court cases, as the context mostly makes it clear who did what. So we have a very weak biometric mechanism that works fairly well in practice — except that it's choked by procedural rules and liability traps 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 [1359], with an analysis by country in [109]. I'll discuss some

of the issues further in Part III. Meanwhile, 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 where better automatic recognition of handwritten signatures could be valuable. This is check clearing.

A bank's check processing center will typically 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 is simultaneously presented on screen with the check image and the customer's reference signature. Verifying checks for small amounts is not economic unless it could 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 by delivery drivers to collect receipts for goods; there have been products since the early 1990s that will compare captured signatures against specimens enrolled previously.

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 trade-off is known as the *receiver operating characteristic*, a term first used by radar operators; if you turn up the gain on your radar set too high, you can't see the target for clutter, while if it's too low you can't see it at all. It's up to the operator to select a suitable point on this curve. The *equal error rate* is when the system is tuned so that the probabilities of false accept and false reject are equal. For tablet-based signature recognition systems, the equal error rate is at best 1%; for purely optical comparison it's several percent. This is not fatal in an operation such as a check processing center, as the automated comparison is used as a filter to preselect dubious checks for scrutiny by a human operator. However, it is a show-stopper in a customer-facing application such as a retail store. If one transaction in a hundred fails, the aggravation to customers would be unacceptable. So UK 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 and indeed fingerprint scanning [500].

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, just as in a check processing center, it was used to screen signatures and 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 is rejected, this merely tells the staff to look more closely, and ask for a driver's license or other photo-ID. With 8500 samples taken from 343 customers, 98.2% were verified correctly at the first attempt, rising to 99.15% after three attempts [452]. But 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% [453]. Because of this disappointing performance, sales of signature recognition technology are only 1.7% of the total biometric market; automation has cost it its leadership of the biometric market.

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.

15.3 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 [1076]. For example, we are extremely good at detecting whether another person is looking at us or not. In normal social applications, humans' ability to identify people by their faces appears to be very much better than any automatic facial-recognition system produced to date.

The human ability to recognize faces is important to the security engineer because of the widespread reliance placed on photo ID. Drivers' licenses, passports and other kinds of identity card are not only used to control entry to computer rooms directly, they are also used to 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 we are good at recognising friends in the flesh, how good are we at identifying strangers by photo ID?

The simple answer is that we're not. Psychologists at the University of Westminster conducted a fascinating experiment with the help of a supermarket chain and a bank [705]. They recruited 44 students and issued each of them with four credit cards each with a different photograph on it:

- 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, and 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 which 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, and 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'. Now 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 most basic incentive is absent.

The response of the banking industry to this experiment was ambivalent. At least two banks who 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 [107]. The overall conclusion was that the benefit to be had from photo ID is essentially its deterrent effect [471].

So maybe people won't use their facial recognition skills effectively in identification contexts, or maybe the information we use to identify people in social contexts is stored differently in our brains from information we get by looking at a single photo. (Recognising passing strangers is in any case much harder than recognising people you know. It's reckoned that misidentifications are the main cause of false imprisonment, with 20% of witnesses making mistakes in identity parades [1360] — not as bad as the near-random outcomes when comparing faces with photos, but still not good.)

But in any case, photo-ID doesn't seem to work, and this is one of the reasons for trying to automate the process. Attempts go back to the nineteenth century, when Francis Galton devised a series of spring-loaded 'mechanical selectors' for facial measurements [510]. But automated face recognition actually subsumes a number of separate problems, and in most of them we don't have the luxury of taking careful 3-d measurements of the subject. In a typical identity verification application, the subject looks straight at the camera under controlled lighting conditions, and his 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 thousands of known suspects. Yet automatic face recognition was one of the technologies most hyped by the security-industrial complex after 9/11 [1084].

Even picking out faces from an image of a crowd is a non-trivial computational task [798]. An academic 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% [13]. Systems that use a combination of techniques can get the error rate down but not to the levels possible with many other biometrics [898, 1370]. Field trials by the U.S. Department of Defense in 2002 found that a leading face-recognition product correctly recognized one individual out of 270 only 51% of the time, and identified one person correctly to within a range of 10 participants 81% of the time [852]. (The vendor in question had put out a press release on the afternoon of September 11th and seen a huge rise in its stock price in the week after trading resumed [782].) By 2003, the technology had improved somewhat, with one vendor recognising 64% of subjects against a database of over 30,000, although performance outdoors was poorer. Tests done in 2001 by the UK National Physical Laboratory (NPL) of a number of biometric technologies found that face recognition was almost the worst, outperforming only vein patterns; its single-attempt equal-error rate was almost ten percent [834]. A UK Passport Office trial in 2005, that was a better approximation to field conditions, found it recognised only 69% of users (though this fell to 48% for disabled participants) [1274].

So the technology still does not work very well in engineering terms. But there are applications where it can have an effect. For example, the Illinois Department of Motor Vehicles uses it to detect people who apply for extra drivers' licenses in false names [454]. Where wrongdoers can be punished, it may be worthwhile to try to detect them even if you only catch a quarter of them (that's still better than the 8% or so of house burglars we catch).

Face recognition has also been used as what Bruce Schneier calls 'security theater'. 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 — but it worked entirely by the placebo effect [1227]. The police have since admitted that they only ever had 20 or 25 villains' faces on the system, and it never recognised any of them [871]. In Tampa, Florida, a similar system was abandoned after an ACLU freedom of information request

discovered that it had recognised no villains [1072]. The ACLU welcomed its demise, remarking that 'every person who walked down the street was subjected to an electronic police line-up without their consent'. (Given that the technology just didn't work, this was maybe a tad zealous.) Face recognition was also tried at Boston's Logan airport; passengers passing through security screening were observed and matched. The system was found to be impractical, with no useful balance between false matches and false alarms [222].

Yet facial recognition is already the second largest-selling biometric with a nominal 19% of the market. However, much of this relates to the automated storage of facial images that are compared by humans — for example, the photos stored in the chips on the new biometric passports. The market for automated recognition is much smaller. Maybe as time passes and technology improves, both its potential (and the privacy worries) will increase.

15.4 Bertillonage

Inventors in the nineteenth century spent quite a lot of effort trying to identify people by their bodily measurements. The most famous of these, Alphonse Bertillon, started out as a clerk in the police records department in Paris, where an important task was to identify serial offenders. In 1882 he published a system based on bodily measurements, such as height standing and sitting, the length and width of the face, and the size and angle of the ear. These were principally used to index a collection of record cards that also held mugshots and thumbprints, which could be used to confirm an identification. This system was known as 'anthropometry', and also as 'Bertillonage' in honour of its creator. Eventually it fell out of favour, once police forces understood how to index and search for fingerprints.

This technique has made a comeback in the form of hand-geometry readers. 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 simple to implement and fairly robust, and the NPL trials found a single-attempt equal error rate of about one percent [834]. (Passport inspection is a less critical application than one might initially think, as airline staff also check passports against passenger lists and provide these lists to the homeland security folks.) Hand geometry is now reported to have 8.8% of the biometric market.

15.5 Fingerprints

Automatic fingerprint identification systems (AFIS) are by far the biggest single technology. In 1998, AFIS products accounted for a whopping 78%

of the $50 m sales of biometric technology; the huge growth of the industry since then has cut this in percentage terms to 43.5% of $1,539m by 2005, but it leads all other automated recognition options. AFIS 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 recent technical reference book on automatic fingerprint identification systems is [832].

The use of fingerprints to identify people was discovered independently a number of times. Mark Twain mentions 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; his 1894 novel *Pudd'nhead Wilson* made the idea popular in the States. Long before that, fingerprints were accepted in a seventh century Chinese legal code as an alternative to a seal or a signature, and required by an eighth century Japanese code when an illiterate man wished to divorce his wife. They were also used in India centuries ago. Following the invention of the microscope, they were mentioned by the English botanist Nathaniel Grew in 1684, by Marcello Malpighi in Italy in 1686; in 1691, 225 citizens of Londonderry in Ireland used their fingerprints to sign a petition asking for reparations following the siege of the city by King William.

The first modern systematic use was in India from 1858, by William Herschel, grandson of the astronomer and a colonial magistrate. He introduced handprints and then fingerprints to sign contracts, stop impersonation of pensioners who had died, and 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 came up with the idea of using latent prints from crime scenes to identify criminals. Faulds brought fingerprints to the attention of Charles Darwin, who in turn motivated Francis Galton to study them. Galton wrote an article in *Nature* [510]; this got him in touch with the retired Herschel, whose data convinced Galton that fingerprints persisted throughout a person's life. Galton went on to collect many more prints and devise a scheme for classifying their patterns [511]. The Indian history is told by Chandak Sengoopta, whose book also makes the point that fingerprinting saved two somewhat questionable Imperial institutions, namely the indentured labor system and the opium trade [1145].

The practical introduction of the technique owes a lot to Sir Edward Henry, who had been a policeman in Bengal. He wrote a book in 1900 describing a simpler and more robust classification, of *loops*, *whorls*, *arches* and *tents*, that he had developed with his assistants Azizul Haque and Hem Chandra Bose, and that is still in use today. In the same year he became Commissioner of the Metropilitan Police in London from where the

technique spread round the world[2]. Henry's real scientific contribution was to develop Galton's classification into an indexing system. By assigning one bit to whether or not each of a suspect's ten fingers had a whorl — a type of circular pattern — he divided the fingerprint files into 1024 bins. In this way, it was possible to reduce the number of records that have to be searched by orders of magnitude. Meanwhile, as Britain had stopped sending convicted felons to Australia, there was a perceived need to identify previous offenders, so that they could be given longer jail sentences.

Fingerprints are now used by the world's police forces for essentially two different purposes: identifying people (the main use in the USA), and crime scene forensics (their main use in Europe).

I'll now look at these two technologies in turn.

15.5.1 Verifying Positive or Negative Identity Claims

In America nowadays — as in nineteenth-century England — quite a few criminals change their names and move somewhere new on release from prison. This is fine when offenders go straight, but what about fugitives and recidivists? American police forces have historically used fingerprints to identify arrested suspects to determine whether they're currently wanted by other agencies, whether they have criminal records and whether they've previously come to attention under other names. The FBI maintains a large online system for this purpose; it identifies about eight thousand fugitives a month [1208]. It is 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, and checks are also run on some people applying to work with children or the elderly. Up to 100,000 fingerprint checks are made a day, and 900,000 federal, local and state law enforcement officers have access. There's now a project to expand this to contain other biometrics, to hold data on foreign nationals, and to provide a 'rap-back' service that will alert the employer of anyone with a clearance who gets into trouble with the law — all of which disturbs civil-rights groups [927]. Since 9/11, fingerprints are also used in immigration. The US-VISIT program fingerprints all aliens arriving at U.S. ports and matches them against a watch list of bad guys, compiled with the help of other police forces and intelligence services worldwide.

These are examples of one type of identity verification — checking an (implicit) claim not to be on a blacklist. The other type is where the system

[2]In the Spanish version of history, 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 USA 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. Success truly has many fathers!

checks a claim to have a certain known identity. Fingerprints are used for this purpose in the USA for building entry control and welfare payment [405]; and banks use them to identify customers in countries such as India and Saudi Arabia, where the use of ink fingerprints was already common thanks to high levels of illiteracy.

Fingerprints 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 [497]. These applications are not routine identity verification, though, so much as an attempt to identify customers who later turn out to be bad — another example being the large British van-hire company that demands a thumbprint when you rent a van. If the vehicle isn't returned, or if it's used in a crime and then turns out to have been rented with a stolen credit card, the thumbprint is given to the police. They are thus really a 'crime scene forensics' application, which I'll discuss in the following section.

So how good are automatic fingerprint identification systems? A good rule of thumb (if one might call it that) is that to verify a claim to identity, it may be enough to scan a single finger, while to check someone against a blacklist of millions of felons, you had better scan all ten. In fact, the US-VISIT program set out to scan just the two index fingers of each arriving visitor, and has been overwhelmed by false matches. With 6,000,000 bad guys on the database, the false match rate in 2004 was 0.31% and the missed match rate 4% [1347]. Although these numbers could be improved somewhat by using the best algorithms we have now in 2007, the program is now moving to '10-prints', as they're called, where each visitor will present the four fingers of each hand, and then both thumbs, in three successive scans.

This is all about the trade-off between false negatives and false positives — the receiver operating characteristic, described in the previous section. In 2001, the NPL study found a 1% false match and 8% false accept rate for common products; by now, the better ones have an equal error rate of slightly below 1% per finger. False accepts happen because of features incorporated to reduce the false reject rate — such as allowance for distortion and flexibility in feature selection [1080]. Spotting returning fugitives with high enough probability to deter them and high enough certainty to detain them (which means keeping false alarms at manageable levels) will require several fingers to be matched — perhaps eight out of ten. But requiring every finger of every passenger to be scanned properly at immigration may cause delays; a UK Passport Office study found that about 20% of participants failed to register properly when taking a 10-print, and that 10-print verification took over a minute [1274]. This will come down with time, but with even an extra

30 seconds per passenger, an airport getting a planeload of 300 international arrivals every 15 minutes would need an extra 10 working immigration lanes. The extra building and staffing costs could swamp anything spent on hardware and software. (For more on algorithms and systems, see [832, 656, 831].)

Errors are not uniformly distributed. 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 [275]. 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 [764]. Fingerprint damage can also impair recognition. When I was a kid, I slashed my left middle finger while cutting an apple, and this left a scar about half an inch long. When I presented this finger to the system used in 1989 by the FBI for building entry control, my scar crashed the scanner. (It worked OK with the successor system from the same company when I tried again ten years later.) Even where scars don't cause gross system malfunctions, they still increase the error rate.

Fingerprint identification systems can be attacked in a number of ways. An old trick was for a crook to distract (or bribe) the officer fingerprinting him, 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 [764]. The most recent batch of headlines was in 2002, when Tsutomu Matsumoto caused much alarm in the industry; he and his colleagues showed that fingerprints could be molded and cloned quickly and cheaply using cooking gelatin [845]. He tested eleven commercially available fingerprint readers and easily fooled all of them. This prompted the German computer magazine C'T to test a number of biometric devices that were offered for sale at the CeBIT electronic fair in Hamburg — nine fingerprint readers, one face-recognition system and one iris scanner. They were all easy to fool — the low-cost capacitative sensors often by such simple tricks as breathing on a finger scanner to reactivate a latent print left there by a previous, authorized, user [1246]. Latest fingerprints can also be reactivated — or transferred — using adhesive tape. The more expensive thermal scanners could still be defeated by rubber molded fingers.

However, fingerprint systems still dominate the biometric market, and are rapidly expanding into relatively low-assurance applications, from entry into golf club car parks to automatic book borrowing in school libraries. (Most European countries' privacy authorities have banned the use of fingerprint scanners in schools; Britain allows it, subject to government guidelines, with the rationale that fingerprints can't be reverse engineered from templates and thus privacy is protected [132]. As I'll discuss later, this reasoning is bogus.)

An important aspect of the success of fingerprint identification systems is not so much their error rate, as measured under laboratory conditions, but their deterrent effect. This is particularly pronounced in welfare payment

systems. Even though the cheap fingerprint readers used to authenticate welfare claimants have an error rate as much as 5% [267], they have turned out to be such an effective way of reducing the welfare rolls that they have been adopted in one place after another [890].

15.5.2 Crime Scene Forensics

The second use of fingerprint recognition is in crime scene forensics. In Europe, forensics are the main application. Prints found at a crime scene are matched against database records, and any that match to more than a certain level are taken as hard evidence that a suspect visited the crime scene. They are often enough to secure a conviction on their own. In some countries, fingerprints are required from all citizens and all resident foreigners.

The error rate in forensic applications has become extremely controversial in recent years, the critical limitation being the size and quality of the image taken from the crime scene. The quality and procedure rules vary from one country to another. The UK used to require that fingerprints match in sixteen *points* (corresponding minutiae), and a UK police expert estimated that this will only happen by chance somewhere between one in four billion and one in ten billion matches [764]. Greece accepts 10, Turkey 8, while the USA has no set limit (it certifies examiners instead). This means that in the USA, matches can be found with poorer quality prints but they can be open to doubt.

In the UK, fingerprint evidence went for almost a century without a successful challenge; a 16-point fingerprint match was considered to be incontrovertible evidence. The courts' confidence in this was shattered by the notorious McKie case [867]. Shirley McKie, a Scottish policewoman, was prosecuted on the basis of a fingerprint match on the required sixteen points, verified by four examiners of the Scottish Criminal Records Office. She denied that it was her fingerprint, and found that she could not get an independent expert in Britain to support her; the profession closed ranks. She called two American examiners who presented testimony that it is not an identification. The crime scene print is in Figure 15.1, and her file print is at Figure 15.2.

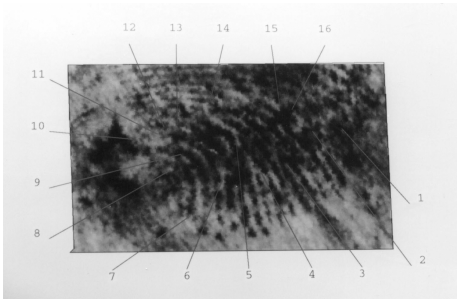

Figure 15.1: Crime scene print

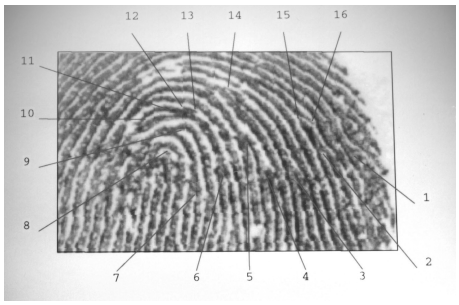

Figure 15.2: Inked print

She was acquitted [866], which led to a political drama that ran on for years. The first problem was the nature of the case against her [867]. A number of senior police officers had tried to persuade her to make a false statement in order to explain the presence, at the scene of a gruseome murder, of the misidentified print. Her refusal to do so led to her being prosecuted for perjury, as a means of discrediting her. Her acquittal said in effect that Glasgow police officers were not reliable witnesses. An immediate effect was that the man convicted of the murder, David Asbury, was acquitted on appeal and sued the police for compensation. A longer term effect was to undermine confidence in fingerprints as forensic evidence. The government then prosecuted its four fingerprint experts for perjury, but this didn't get anywhere either. The issue went back to the Scottish parliament again and again. The police refused to reinstate Shirley, the officers involved got promoted, and the row got ever more acrimonious. Eventually she won £750,000 compensation from the government [130].

The McKie case led to wide discussion among experts of the value of fingerprint identification [522]. It also led to fingerprint evidence being successfully challenged in a number of other countries. Two high-profile cases in the USA were Stephan Cowans and Brandon Mayfield. Cowans had been convicted of shooting a police officer in 1997 following a robbery, but was acquitted on appeal six years later after he argued that his print was a misidentification and saved up enough money to have the evidence tested for DNA. The DNA didn't match, which got the Boston and State police to reanalyze the fingerprint, whereupon they realised it was not a match after all. Brandon Mayfield was an Oregon lawyer who was mistakenly identified by the FBI as one of the perpetrators of the Madrid bombing, and held for two weeks until the Madrid police arrested another man whose fingerprint was a better match. The FBI, which had called their match 'absolutely incontrovertible', agreed to pay Mayfield $2 m in 2006.

In a subsequent study, psychologist Itiel Dror showed five fingerprint examiners a pair of prints, told them they were from the Mayfield case, and asked them where the FBI had gone wrong. Three of the examiners decided that the prints did not match and pointed out why; one was unsure; and one maintained that they did match. He alone was right. The prints weren't the Mayfield set, but were in each case a pair that the examiner himself had matched in a recent criminal case [402]. Dror repeated this with six experts who each looked at eight prints, all of which they had examined for real in the previous few years. Only two of the experts remained consistent; the other four made six inconsistent decisions between them. The prints had a range of difficulty, and in only half of the cases was misleading contextual information supplied [403].

How did we get to a point where law enforcement agencies insist to juries that forensic results are error-free when FBI proficiency exams have long had

an error rate of about one percent [141], and misleading contextual information can push this up to ten percent or more?

Four comments are in order.

- As Figure 15.1 should make clear, fingerprint impressions are often very noisy, being obscured by dirt. So mistakes are quite possible, and the skill (and prejudices) of the examiner enter into the equation in a much bigger way than was accepted until the McKie case, the Mayfield case, and the general uproar that they have caused. Dror's work confirmed that the cases in which misidentifications occur tend to be the difficult ones [403]. Yet the forensic culture was such that only certainty was acceptable; the International Association for Identification, the largest forensic group, held that testifying about ''possible, probable or likely identification shall be deemed . . . conduct unbecoming.'' [141]
- Even if the probability of a false match on sixteen points were one in ten billion (10^{-10}) as claimed by police optimists, once many prints are compared against each other, probability theory starts to bite. A system that worked fine in the old days as a crime scene print would be compared manually with the records of a hundred and fifty-seven 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 criminals who will not be able to muster the resolute defence that Shirley McKie did, I would be surprised if there hadn't already been other wrongful convictions. Indeed, things may get worse, because of a 2007 agreement between European police forces that they will link up their biometric databases (both fingerprints and DNA) so that police forces can search for matches across all EU member states [1261]. I expect they will find they need to develop a better understanding of probability, and much more robust ways of handling false positives.
- The belief that any security mechanism is infallible creates 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 sixteen to (say) twenty 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' experts. In the UK, all the experts were policemen or former policemen, so there were no independents available for hire. Even so, it would have been possible to use randomised matching with multiple experts; but if the fingerprint bureau had had to tell the defence in the perhaps 5–10% of cases when (say) one of four experts disagreed, then

many more defendants would have been acquitted and the fingerprint service would have been seen as less valuable.

- A belief of infallibility ensures that the consequences of the eventual failure will be severe. As with the Munden case described in section 10.4.3, 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 will thus have serious effects when it can no longer be postponed. In the Scottish case, there appears to have arisen a hierarchical risk-averse culture in which no-one wanted to rock the boat, so examiners were predisposed to confirm identifications made by colleagues (especially senior colleagues). This risk aversion backfired when four of them were tried for perjury.

However, even when we do have a correct match its implications are not always 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 [179]). Of course, even if the villain wasn't framed, he can always claim that he was and the jury might believe him.

In the USA, the Supreme Court's Daubert judgment [350] ruled that trial judges should screen the principles and methodology behind forensic evidence to ensure it is relevant and reliable. The judge ought to consider the refereed scientific literature — and in the case of fingerprints this has been somewhat lacking, as law enforcement agencies have been generally unwilling to submit their examination procedures to rigorous double-blind testing. A number of Daubert hearings relating to forensic fingerprint evidence have recently been held in U.S. trials, and the FBI has generally prevailed [523]. However, the bureau's former line that fingerprint examination has a zero error rate is now widely ridiculed [1208].

15.6 Iris Codes

We turn now from the 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. Research on the subject was funded by the Department of Energy, which wanted the most secure possible way of controlling entry to premises such as plutonium stores, and the

technology is now being used in applications such as immigration. The latest international standards for machine-readable travel documents mandate the use of photographs, and permit both fingerprints and irises.

So far as is known, every human iris is measurably unique. It is fairly easy to detect in a video picture, it does not wear out, and it 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.

John Daugman found signal processing techniques that extract 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 15.3). The resulting iris codes have the neat property that two codes computed from the same iris will typically match in 90% of their bits [351]. 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 [1327]. Iris codes provide the lowest false accept rates of any known verification system — zero, in tests conducted by both the U.S. Department of

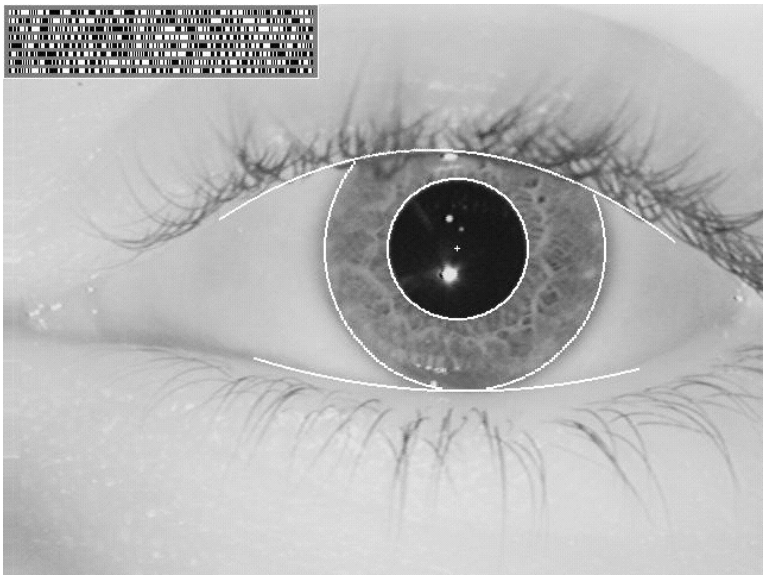

Figure 15.3: An iris with iris code (courtesy John Daugman)

Energy and the NPL [834]. 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. In practice, the false reject rate is significantly higher than this; many things, from eyelashes to hangovers, can cause the camera to not see enough of the iris. The U.S. Department of Defense found a 6% false reject rate in its 2002 field trials [852]; the Passport Office trial found 4% for normal users and 9% for disabled users [1274]. A further problem is failure to enrol; the Passport Office trial failed to enrol 10% of participants, and the rate was higher among black users, the over-60s and the disabled.

One practical problem with iris scanning used to be getting the picture cheaply 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. All current iris scanning systems use infrared light, and some people feel uncomfortable when this is shone in their eyes. (The Chinese government gave this as an excuse for rejecting iris scanning for the latest Hong Kong identity cards, going for a thumbprint instead [771].) Given more sophisticated cameras, with automatic facial feature recognition, pan and zoom, it is now possible to capture iris codes from airline passengers covertly as they walk along a corridor [841], and no doubt the cost will come down in time (especially once the key patent runs out in 2011). This is likely to make overt uses less objectionable; but covert identification of passers-by has Orwellian overtones, and in Europe, data protection law could be a show-stopper.

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. There are terminals available that will detect such simple fakes, for example by measuring *hippus* — a natural fluctuation in the diameter of the pupil that happens at about 0.5 Hz. But the widely-sold cheap terminals don't do this, and if liveness detection became widespread then no doubt attackers would try more sophisticated tricks, such as printing the target's iris patterns on a contact lens.

As iris recognition is fairly new, we don't have as much experience with it as we have with fingerprints. The biggest deployment so far is in the United Arab Emirates where it's used to screen incoming travelers against a blacklist of people previously deported for illegal working. The blacklist has 595,000 people as of July 2007 — 1.19 million irises — and so far 150,000 deportees have

been caught trying to re-enter the country. The typical arrestee is a lady with a previous conviction for prostitution, who returns with a genuine (but corruptly issued) passport, in a new name, from a low or middle income Asian country. A typical attack was for the returning deportee to take atropine eyedrops on the plane, dilating her pupils; nowadays such travelers are held in custody until their eyes return to normal. Nonetheless, the atropine trick might be a problem for blacklist applications in developed countries. There might also be evidentiary problems, as iris recognition depends on computer processing; there are no 'experts' at recognising eyes, and it's doubtful whether humans could do so reliably, as the information that John Daugman's algorithms depend on is mostly phase information, to which the human eye is insensitive. (In developed countries, however, the typical application is a frequent-traveler program that allows enrolees to bypass passport control at an airport; there the users want to be recognised, rather than wanting not to be. The UK, for example, has such a scheme with 200,000 enrolees. Here, evidence isn't really an issue.)

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 front of you is the same human as the one who was initially registered on the system. They alone can meet the goal of automatic recognition with zero false acceptances.

15.7 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.

As with fingerprints, the technology is used for both identification and forensics. In *forensic phonology*, the task 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 [721]. 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% [264]; and the NSA

maintains a standard corpus of test data for evaluating speaker recognition systems [655]. A recent application is the use of voice recognition to track asylum seekers in the UK; they will be required to ring in several times every week [1260]. Such systems tend to use caller-ID to establish where people are, and are also used for people like football hooligans who're under court orders not to go to certain places at certain times.

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 [506] there is a brief description of a system fielded in U.S. EP-3 aircraft that 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 that 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. There has been research aimed at improving them to the point that call centers can have the same 'person' always greet you when you phone; and audio remixing products improve all the time. Remote voice biometrics look less and less able to withstand a capable motivated opponent.

15.8 Other Systems

Many other biometric technologies have been proposed [890]. 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). Vein patterns have been used in one or two systems but don't seem to have been widely sold (in the NPL trials, the vein recognition ROC curve was almost all outside the other curves; it was the worst of the lot) [834].

There has been growing interest recently in identifying anonymous authors from their writing styles. Literary analysis of course goes back many years; as a young man, the famous cryptologist William Friedman was hired by an eccentric millionaire to study whether Bacon wrote Shakespeare. (He eventually debunked this idea but got interested in cryptography in the process.) Computers make it possible to run ever more subtle statistical tests; applications range from trying to identify people who post to extremist web fora to such mundane matters as plagiarism detection [3]. It's possible that such software will move from forensic applications to real-time monitoring, in which case it would become a biometric identification technology.

Other proposals include *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. Bertillon used the shape of the ear in nineteenth century Paris, but most of the rest of these exotica don't seem to have been marketed as products. Other technologies may provide opportunities 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.

One final biometric deserves passing mention — DNA typing. This has become a valuable tool for crime scene forensics and for determining parenthood in child support cases, but it is still too slow for applications like 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 his DNA sample. There have been major procedural problems, with false matches resulting from sloppy lab procedure. And there are also major data quality problems; the UK police have the biggest DNA database in the world, with records on about four million people, but have got the names misspelled or even wrong for about half a million of them [588]. The processes that work for local policing don't always scale nationally — small errors from mistyped records, to suspects giving false names that were never discovered because they weren't prosecuted, accumulate along with lab errors until the false-positive rate becomes a serious operational and political issue. For a survey of forensic DNA analysis, and suggestions of how to make national DNA databases consistent with privacy law, see [1124].

15.9 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 scared off some police forces which sent in samples from different volunteers and got back false matches, but also led to disputed court cases and miscarriages of justice. This is reminiscent of the fingerprint story, and brings to mind the quote from Lars Knudsen at the head of Chapter 5: *'if it's provably secure, it probably isn't'*. Any protection measure that's believed to be infallible will make its operators careless enough to break it.

Biometrics are also like many other physical 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 stand-alone, from cooperative to recalcitrant subjects, and from verification to identification, can all undermine a system's viability.

There are a number of interesting attacks that are more specific to biometric systems and that apply to more than one type of biometric.

- 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 he claims that he had gone there three days previously, he 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 moulds of fingerprints. Even simpler still, in countries like South Africa 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. The lesson to be learned here is that unattended operation of biometric authentication devices is tricky. Attacks aren't always straightforward; although it's easy to make a mold from a good fingerprint [281], the forensic-grade prints that people leave lying around on doorknobs, beer glasses and so on are often too smudged and fragmentary to pass an identification system. However, attacks are definitely possible, and definitely happen.
- 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 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 foolish and offensive. A biometric system that is (or is seen to be) socially regressive — that 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 [1046]. 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?

- A point that follows from this is that systems may be vulnerable to collusion. Alice opens a bank account and her accomplice Betty withdraws money from it; Alice then complains of theft and produces a watertight alibi. Quite apart from simply letting Betty take a rubber impression of her fingertip, Alice might voluntarily decrease handwriting ability; by giving several slightly different childish sample signatures, she can force the machine to accept a lower threshold than usual. She can spend a couple of weeks as a bricklayer, building a wall round her garden, and wear her fingerprints flat, so as to degrade registration in a fingerprint system. She might register for a voice recognition system when drunk.
- 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, 1609 people) enrolled. So identification is a tougher task than verification [352]. The practical consequence is that a system designed for authentication may fail when you try to rely on it for evidence.
- 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 [352].
- Many vendors have claimed that their products protect privacy, as what's stored is not the image of your face or fingerprint or iris, but rather a template that's derived from it, somewhat like a one-way hash, and from which you can't be identified. It's been argued from this that biometric data are not personal data, in terms of privacy law, and can thus be passed around without restriction. These claims were exploded

by Andy Adler who came up with an interesting *hill-climbing attack* on face recognition systems. Given a recogniser that outputs how close an input image is to a target template, the input face is successively altered to increase the match. With the tested systems, this led rapidly to a recognizable image of the target — a printout of which would be accepted as the target's face [14]. He then showed how this hill-climbing technique could be used to attack other biometrics, including some based on fingerprints [15].

- Automating biometrics can subtly change the way in which security protocols work, so that stuff that used to work now doesn't. An example is the biometric passport or identity card that contains your digital photo, and perhaps your fingerprint and iris data, on an RFID chip. The chip can be cloned by copying the contents to another RFID chip (or replaying them through a phone with an NFC interface.) The world's passport offices took the view that this wasn't a big deal as the data are signed and so the chip can't be altered. However, the police have another use for passports — if you're on bail they insist that you leave your passport with them. That protocol now breaks if you can leave the country via the fast track channel by replaying your iris data through your mobile phone. There was also some embarrassment when researchers discovered that despite the digital signature, they could modify the RFID contents after all — by replacing the JPEG facial image with a bitstring that crashed the reader [1374]. This in turn raises the question of whether a more cunningly designed bitstring could modify the reader's behaviour so that it accepted forged passports. I suppose the moral is that when passport offices digitized their systems they should have read all of this book, not just the chapters on biometrics and crypto.
- It's worth thinking what happens when humans and computers disagree. Iris data can't be matched by unaided humans at all; that technology is automatic-only. But what happens when a guard and a program disagree on whether a subject's face matches a file photo, or handwriting-recognition software says a bank manager's writing looks like a scrawled ransom note when they look quite different to the human eye? Psychologists advise that biometric systems should be used in ways that support and empower human cognition and that work within our social norms [404]. Yet we engineers often find it easier to treat the users as a nuisance that must adapt to our technology. This may degrade the performance of the humans. For example when an automated fingerprint database pulls out what it thinks is the most likely print and presents it to the examiner: is he not likely to be biased in its favour? Yet if the computer constantly tested the examiner's alertness by giving

him the three best matches plus two poor matches, would that work any better?

- Finally, Christian fundamentalists are uneasy about biometric technology. They find written of the Antichrist in Revelation 13:16-18: '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 may arouse political opposition on the right as well as the left.

So there are some non-trivial problems to be overcome as biometrics tiptoe towards mass-market use. 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.

15.10 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 iris patterns and voiceprints. 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 much 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 synthesisers, photographs of irises, fingerprint moulds 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 large error rates. That particular case teaches us that context matters; even a weak biometric can be effective if its use is well embedded in the social and legal matrix.

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. Forensic uses are problematic, and courts are much less blindly trusting of even fingerprint evidence than they were ten years ago. Finally, many biometric systems achieve most or all of their result by deterring criminals rather than actually identifying them.

Research Problems

Many practical research problems relate to the design, or improvement, of biometric systems. Is it possible to build a system — other than iris scanning — which 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 which has a good enough error rate (say 1%) for it to be used for screening images of all checks, rather than just as a pre-screening stage to human inspection of high-value checks? Are there any completely new biometrics that might be useful in some circumstances?

One I thought up while writing this chapter for the first edition in 2000, 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 operated the gears and the clutch. If your car thinks it's been stolen, it phones a GPS fix to a control center which then calls you to check. Recently this has come to pass; there is now research showing that users of haptic systems can be recognised by the way in which they use tools [990].

Further Reading

The history of fingerprints is good reading. The standard reference is Lambourne [764], while Block has a good collection of U.S. case histories [195] and the history of fingerprints in India is told by Sengoopta [1145]. The McKie case is described in a book by Ian McKie and Michael Russella [867]. A good technical reference on automated fingerprint identification systems is the book by Maltoni, Maio, Jain and Prabhakar [832]; there's also an earlier book by Jain, Bolle and Pankanti [655]. As for facial and handwriting recognition in the text, there's also an IBM experimental system described at [684] and a survey of the literature at [288]. The standard work on iris codes is Daugman [351]. For voice recognition, there is a tutorial in [264] which focuses on speaker identification while for the forensic aspects, see Klevans and Rodman [721]. Snapshots of the state of the technical art can be found in two journal special issues of the *Proceedings of the IEEE* on biometric systems — volume 85 no 9 (September 1997) and volume 94 no 11 (November 2006).