Chapter 12
Vulnerabilities

The assessment of vulnerability is vital for ensuring biometric security, and is a concept distinct from system accuracy. A perfectly accurate biometric system may still be highly vulnerable to attack, as unauthorized users may find alternates ways by which they can be falsely accepted by a system.

Compared with the effort expended on determining performance accuracy, significantly less effort has been given to the problem of determining if a presented biometric is real or fake. With the increasing use of biometric systems, the understanding of vulnerability related risks and their appropriate treatment will be a vital part of future biometric deployments.

All the attack methods described in this chapter are vulnerabilities that are publicly known. As a general principle, the public dissemination of points of vulnerably is an important step towards ensuring system designers can put in place appropriate risk mitigations. Secrecy about avenues of attack can help potential fraudsters more than the disclosure of risks, since where the risks are not understood by the system owners, attack methods may be easily exploited. The principle of security through transparency is accepted practice in the cryptographic community.

There are four high-level factors that contribute to a biometric system's vulnerability to a determined attack: the security of the computing infrastructure, trust in the human operators, the accuracy of the matcher, and the ability to detect fake biometrics. There are long established standards and practices for assessing the first two factors, which include ensuring the security of communication channels and storage, tamper-proofing devices, and establishing usage policies. However, the non-deterministic nature of biometric matching and its interface to the real world means the latter factors create a variety of new security threats to be mitigated against in order to reduce the chance of a malicious attack being successful.

The goals of this chapter are to:

- Introduce the analysis of biometric vulnerabilities (Sect. 12.1).
- Look at the research that has been undertaken in detecting fake biometrics (Sect. 12.2.2).
- Outline the different points of attack in a biometric system (Sect.12.3).

- Describe different fraud types, including enrollment, covert and cooperative (Sect.12.4.1).
- Discuss methods for assessing vulnerabilities (Sect.12.5).
- List mitigations that can be used to address biometric vulnerabilities (Sect.12.6).

Definitions for terms related to vulnerability used in this chapter can be found in Chap. 6.

12.1 Introduction

In February 2002, former U.S. Secretary of Defense Donald Rumsfeld said in response to questions about military threats:

> "... there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know." [5]

The detection and mitigation of biometric threats involves many of the same questions about what is known and unknown. The list of "known threats" is those that can be, or already have been, identified. Each "known threat" that has not been able to be investigated or evaluated fully can then be categorized as a one of the "known unknowns". In other words, the threat is known, but its impact or likelihood is not well understood. Threats may also be "yet to be discovered", and these are the "unknown unknowns". The discovery of these new threats requires active intelligence on the activities of attackers and researchers. It also requires creative and knowledgeable vulnerability evaluators. One thing that seems certain is further attack methods will be discovered, so the quick identification and mitigation of vulnerabilities is increasingly important to both the security of systems and the credibility of the industry.

The cost of any particular vulnerability is proportional to the value of the assets protected by the biometric, which might range from secure access to an office to the launch control for a missile, multiplied by the total risk of compromise, which is the chance that an attack will be successful. It is the total factor risk or the "spoofing risk" that is often of interest when examining a system's vulnerability.

Biometrics is a probabilistic science. Every time an individual has their biometric acquired it will be slightly different. This variation is caused by a combination of user behavior, environmental conditions and physical aging of the biometric, and means we can never be absolutely certain of identity through biometric means alone. However, the vulnerability of a biometric system should not be confused with its accuracy. It is possible to have a system that is extremely accurate at distinguishing between any two individuals, but which may be highly vulnerable to simple methods used to circumvent the security, either by mimicking physical biological characteristics, or by bypassing or altering the information used as part of the matching process.

In order to develop commercially useful biometric systems, past effort has focused on improving the ability of the biometric algorithms to distinguish between

different people. Most large-scale evaluations of biometric technology conduct tests to determine the probability that a random person will match successfully against them self, or be mistaken for someone else. As a consequence, matching engines have become increasingly specialized in undertaking this distinguishing task, and deliberately ignoring transitory factors that do not aid in the identification process. However, this focus can potentially increase the system's vulnerability to attack by reducing the number of aspects of a biometric presentation that an attacker needs to fake. Therefore, vulnerability mitigations should focus on techniques that are supplementary, or orthogonal, to improving algorithm performance. These techniques are often termed 'liveness' or 'spoof' detection. The goal of the liveness detection is to prevent the acceptance, regardless of the match score, of fake biometrics (Fig. 12.1).

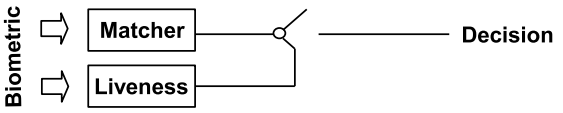

Fig. 12.1 The relationship between liveness detection, biometric matching and the match decision.

12.2 History

A variety of biometric vulnerabilities have been exposed both formally through academic research [13, 20, 27] and informally through magazine articles [26], hacker groups [15] and even television shows [6]. However, many of the demonstrated exploits are undertaken on less advanced systems, and examinations of vulnerability have been largely ad-hoc rather than systematic searches for all threat vectors.

12.2.1 Common Criteria

The international standard used for computer security, particularly by governments, is called the Common Criteria (CC) [2, 4]. The goal is to define rigorous standard processes that will determine a level of assurance, known as an Evaluation Assurance Level (EAL), in the security of computer products. For each security technology the requirements to be assessed are listed in a document called a Protection Profile (PP). Several such PP's are available from national standards bodies [1, 3, 14], however the most influential is the United Kingdom PP [25].

Unfortunately, the Common Criteria has not been very successful for the evaluation of biometric devices, as these are rapidly evolving technologies and have

not been compatible with the sometimes ponderous and costly Common Criteria. The non-deterministic nature of the technology makes it harder to undertake formal testing, as this relies on strict repeatability and restrictions in the set of valid input parameters. However, this approach is not suitable for exploring the space of potential vulnerabilities for a biometric system.

12.2.2 Liveness Research

The detection of liveness is an active area of research, and for each biometric modality different techniques have been suggested for assessing liveness. Fingerprints have had the largest amount of research undertaken; suggested techniques include optical properties [18], pulse [23], perspiration [22], electric resistance [23], subepidermis structure [23], skin deformation [18], papillary lines[11], pores [18] or a combination thereof [18, 24].

For face recognition, the incorporation of head motion has long been known as a method to prevent the use of a static pictures [26]. Furthermore, the natural blinking rate of eyes can be used [21], as well as multi-spectral imaging. Face recognition systems can be particularly vulnerable to poor quality image enrollment [17], so ensuring quality control assists in preventing certain attack mechanisms.

The detection of pupil movement and saccade (eye movement) [9, 10] is used in some iris systems. Other techniques include the use of controlled light to check pupil response, the detection of infrared reflections off the cornea [9], and multi-spectral sensing [8] .

12.3 Points of Attack

A biometric system is composed of a number of different subsystems (see Sect. 1.6). Each subsystem may have a number of different points of attack, and for each point of attack there may be one or more potential exploits. Although such attack points exist in all matching systems, not all are equally vulnerable. In general, the less distributed the system, the easier it is to secure.

For instance, consider a secure biometric smart-card with all of the subsystems, including the sensor, integrated on the card. In the case that only the authentication decision need to be securely transmitted, this would significantly reduce the potential points of attack. Such cards may soon be practical with advances in manufacturing and sensing technologies. The other major factor is the degree of trust in the network and environment, as this affects the likelihood of a vulnerability being exploited.

The points of attack are shown in Fig 12.2. This diagram shows a biometric system both at the subsystem level as well as each individual component, and highlights the potential points of vulnerability:

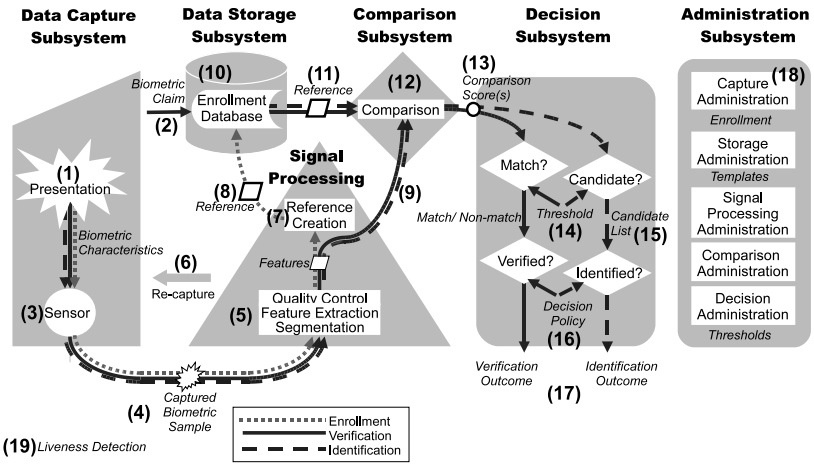

Fig. 12.2 Vulnerability points in a general biometric system. Derived from a diagram used with permission from Tony Mansfield, National Physical Laboratory, UK.

1. **Presentation**: The use of a fake biometric to enroll is the principal threat at the presentation point. This may be through the use of an *artifact* instead of the live biometric, or through the modification of an existing characteristic made to look more like an enrolled one, for instance the use of makeup. This process may occur at either the enrollment or verification stage. Another vulnerability is that a user might be coerced into presenting a biometric.
   Example: A latent fingerprint obtained from a glass surface and a fake fingerprint made from gelatin that replicates its ridge pattern is created and used to fool a fingerprint sensor.
2. **Identity Claim**: At the point of enrollment or verification, the use of a fake or stolen identifier to create a false claim of identity may result in either an invalid enrollment or a potential false accept.
   Example: An attacker creates a fake passport and uses this proof of identity to create a new identity in a government application that uses iris recognition.
3. **Sensor**: The integrity of the sensor allows trust in the integrity of the acquired biometric sample. If the sensor can be faked or compromised, the system may see what it believes to be a valid biometric sample, but which is actually an artifact or replay attack. Ideally the signal processing subsystem will be able to check using cryptographic techniques[1] that the sensor has not been tampered with and is operating properly.
   Example: An attacker removes the camera from a laptop and replaces it with a fake that always sends the same image.
4. **Transmission - Sample**: If the transmission of the biometric sample from the sensor to the signal processor travels over an insecure connection it may be intercepted for later use. Alternatively, the substitution of a fake biometric sample

[1] In smart-cards this is known as a SAM (Security Authentication Module)

for the real one may be undertaken.
Example: An attacker intercepts a fingerprint image coming from a sensor and stores it for later use in a replay attack.

5. **Quality control and feature extraction**: Low quality enrollments or verification samples can be a source of creating lamb or chameleon templates (those that are easy to spoof). Hence, it is important that tight control over the quality of the enrolled biometric data is maintained. The extraction of the features is at the heart of processing for a biometric algorithm: if this process can be compromised, it may lead to a significant threat.
Example: A person enrolls in a fingerprint system with a dirty finger. The poor quality enrollment allows others to more easily spoof this identity.
6. **Re-capture**: Through the continual re-acquisition and re-capture of a biometric an attacker can refine attack mechanisms by altering the biometric to discover which techniques work best. This may be especially the case for algorithms that are widely available for general purchase.
Example: A hand geometry system allows an attacker to attempt unlimited retries. This allows the attacker to figure out how to spoof the sensor.
7. **Reference creation**: If the creation of the reference feature and generation of a template can be compromised, this then can create a significant threat.
Example: A hacker inserts or changes code in the reference creation to ensure that whenever a particular palm vein pattern is seen it always generates a high score. The hacker can then distribute copies of this reference template to allow system access for fellow hackers.
8. **Transmission - Reference to enrollment**: If the transmission of the template from the reference creation process is over an insecure channel, the enrollment might be substituted for another before it is stored in the database.
Example: A hacker has infiltrated the database connectivity layer and substitutes templates as they are inserted into the database.
9. **Transmission - Features to database**: As for the enrollment reference transmission, when the features from the reference creation process are sent over an untrusted channel, the sample might be substituted for another before it is stored in the database.
Example: A hacker has infiltrated the database connectivity layer and substitutes templates before they are matched.
10. **Enrollment database**: The enrollment database is the source of the authentication data; if the enrollment database is compromised, this would allow any number of potential alterations and substitutions.
Example: A malicious database administrator inserts new templates for attackers.
11. **Transmission - Reference from database**: When the reference is transmitted from the database if it is over an insecure channel a hacker would be able to substitute templates before they are compared.
Example: A hacker substitutes a template retrieved from the database before it can be compared.
12. **Comparison process**: The comparison process creates a similarity score between the reference template and the verification sample features. If hackers can

compromise this process, they can output high scores for selected identities.
Example: A hacker changes the comparison process so that high scores are always given during a specific time period.

13. **Transmission - Score**: If the score is not transmitted securely it may be altered before it reaches the decision subsystem.
Example: A hacker substitutes high scores for people with particular identities.
14. **Threshold process**: If the threshold process is compromised, the match threshold may be lowered, making it easier for attackers to be accepted by the system.
Example: A hacker sets the system threshold to zero, allowing all individuals to pass.
15. **Candidate list**: During identification the candidate list results could be modified or re-ranked to exclude specific individuals**.**
Example: A hacker ensures that particular identities are never ranked highly enough to be presented to the operator.
16. **Decision policy**: The decision policy uses business rules to convert the match results into a final acceptance or rejection. An attack who had the ability to change this policy would be able allow acceptance decisions at will.
Example: Modification by a rogue administrator of the business rules around exception cases, falsely labeling an individual as someone who could not enroll, in order to bypass the biometric security mechanisms.
17. **Transmission - Outcome**: The final decision needs to be transmitted for action; if this transmission protocol is compromised then the matching outcome could be altered to generate a successful match.
Example: A fingerprint sensor used for access control on a secure door transmits the unlock code to the door lock using a simple power relay. The attacker removes the sensor from the wall and shorts the open wires together, causing the door to unlock.
18. **Administration**: The administration subsystem potentially controls all aspects of decisions from acquisition and quality setting through to business rules and threshold settings. The security and audit of administrative access is hence a critical component.
Example: A malicious administrator substitutes a fake enrollment and reduces the thresholds to allow an attacker to pass under a false identity.
19. **Liveness detection**: Liveness detection is the mitigation strategy used to protect against the use of prosthetic artifacts. However, the liveness detection process itself may also be open to attack.
Example: A fingerprint system that uses heat for liveness detection may be spoofed by warming the artifact fingerprint or by creating a thin film to put over the top of a real finger.

12.4 Fraud

In addition to the attack points described above, fraud in biometric systems can be broken into a number of different classes. These dependend on when the fraud occurs (enrollment or verification) and the type of attack that has been mounted (covert or cooperative). Covert verification fraud is the most commonly considered fraud type, however it is the processes around enrollment that are the most crucial to ensure system integrity.

12.4.1 Enrollment Fraud

One of the most vulnerable points in any biometric system is the enrollment process. If poor control is maintained over the enrollment process then the overall integrity of the system can be seriously compromised.

Ensuring the enrollment processes has integrity usually means providing a reliable link to other identity credentials. This credentialing is commonly achieved through the use of proof of identity documents such as a birth certificate, passport or driver's license. *The strength of subsequent authentications using a biometric is dependent on the integrity and strength of the enrollment process.*

Where it is important to have high credential strength, biometric enrollment should always be supervised, as this helps mitigate against the use of artifact attacks since it is harder to use a fake biometric when you are being watched, and a human can also look for other suspicious activity. Also important is maintaining a strong audit trail of the enrollment process, including who undertook the enrollment and when it occurred.

12.4.2 Covert Fraud

Covert fraud is when an attack is undertaken without the knowledge of the person to be spoofed. This is the most common scenario to educate people about since it can lead to identity theft.

An attacker can covertly obtain an individual's biometric through several mechanisms. The most commonly considered is the creation of an artifact by use of either a biometric impression, for instance dusting for a fingerprint left on a surface and then creating an artifact, or through some form of surveillance activity. Often the covertly acquired biometric will be degraded through noise or missing features, and the creation of the artifact for attack will be of lower quality or have significant quality variations. The creation of artifacts will seldom have a perfect success rate, so any attacker needs to consider the risk that the artifact produced may in fact fail when tried on the target systems.

Raising Latent Fingerprint Prints

The fingerprints of the last person to use a sensor are sometimes visible on the surface of the sensor for some time. In early implementations of fingerprint sensors it was discovered that these 'latent prints' could be raised through simply breathing on the sensor or using a bag filled with water [26]. Modern systems should not be susceptible to such attacks since detecting the presence of a fingerprint that is almost identical to the one used previously is relatively easy. However, not all sensors explicitly check for this vulnerabillity.

Other methods of covert fraud involve reverse engineering the template. Since the template contains enough information for the algorithm to recognize a person, it follows that it should be possible to reconstruct a biometric that would successfully pass using only this information. Two methods have been used for this purpose. The first is called a *hill-climbing attack*. In this attack, the attacker must have access to the output of matching algorithm. They then use this to compare the template to some sample input. Random successive changes are made to the sample, and those that improve the match score are maintained and iteratively modified. Eventually, this may result in an fake biometric sample that doesn't necessarily look like the original, but is able to successfully authenticate. Examples using this technique have been shown to be effective in defeating face recognition systems [7]. The other form of attack is called a *masquerade attack*. This attack utilizes knowledge about the structure of the template (for instance, the position and location of fingerprint minutiae) to attempt to explicitly recreate the source biometric. It relies on the attacker understanding how to decode and interpret the template structure. Fingerprint systems have been successfully spoofed using this technique [16].

12.4.3 Cooperative Fraud

When the party to be imitated is colluding in the attack, for instance by allowing other people to use their identity, it is considerably easier for the attacker. Cooperative attacks might include the deliberate creation of poor quality templates (lambs), or the use of an artifact during enrollment that can be given to others.

'Insider fraud' can also be cooperative. This could be due to the collusion of the operator undertaking the enrollment, or of other administrative staff. Catching this sort of fraud can be particularly difficult, and relies on strong audit trails and the correct corporate culture where such activities are regarded extremely seriously.

When people are assessing attack likelihoods it needs to be understood that cooperative attacks are likely to have a much higher success rate than covert attacks. For instance, the success rates will obviously be much higher from a fake fingerprint created from a cooperative party, than compared to a covert acquisition (e.g. from a

partial latent print). *Although the threat being assessed might be the same, the risk will vary depending on whether the attack is cooperative or covert.*

12.5 Assessing Vulnerabilities and Attack Methods

The evaluation of biometric threats should provide a reliable estimate of the vulnerability for a particular threat using specific technology. It is necessary to ensure that the assessment meets resource constraints, and can flexibly adjust to different biometric modalities in an evolving threat landscape. Ideally, it will also fit in with other security assessment processes.

Biometrics Institute Vulnerability Assessment Methodology

The Biometrics Institute vulnerability assessment methodology provides a principled methodology for assessing the vulnerability of biometric systems to deliberate attacks. A key part of this methodology is to separate the total risk factor into two separate components: *Exploitation Potential*, which relates to properties of the biometric system itself; and *Attack Potential*, which is primarily a function of the capabilities of an attacker [12].

The set of potential threats and threat variants is complex, and a wide variety of factors need to be considered during the evaluation. Each biometric product will have different levels of vulnerability to a range of threats, and each threat is dependent on the attributes of an attacker. Potential threats against a biometric system range from the presentation of artifacts, such as simple printed picture of a biometric, through to the reconstruction of a biometric from stolen biometric templates. The protection profiles established provide one such baseline list of threats [1]. However, the threat list is not static and may continually be expanding as new techniques and materials become available upon which to base attacks.

A general threat list for a biometric system will include many threats that would apply to any complex computer security environment. Given the wide scope of such investigations it is seldom possible to investigate all known threats. One method of dealing with this is to rank the threats in order of how likely the exploit will be for a given system, and investigate those with the highest priority first. The experience of the evaluator to make informed judgments about how the different threats compare, dependending on the nature of the application and what is being protected, is relied upon to ensure this ranking does not yield misleading results.

During testing, vulnerability to a threat is indicated when both the liveness detection is defeated (where it is available) and also high match scores are observed during an attack. A score above the minimum threshold that could be practically

set represents a potentially successful attack. The testing process is generally concerned with looking for artifacts that can defeat the liveness tests and obtain the highest similarity scores.

The components of a vulnerability assessment process are to select threats (see Fig. 12.3) and then apply an assessment methodology. This methodology will undertake a testing process for each threat and provide standardized reporting [12].

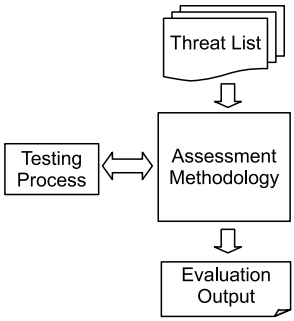

Fig. 12.3 The process of the assessment and reporting of threats.

12.5.1 Attacker Strength

The strength of an attacker is an important consideration when assessing the risk of an attack being successful. Attackers span from a casual impostor or an uneducated criminal with access to information available on the Internet, through to a skilled programmer with access to state sponsorship and full administrator rights. For each increase in attacker strength the likelihood of a system penetration may increase dramatically. However, it can be known a priori that some attacks are impossible without higher level access to the system in order to be able to change settings or inject new information. Similarly, some attacks can be accomplished with very few resources, whilst others require significant skill, time and manufacturing expertize. This difference can be characterized by the level of information that an attacker has, and their level of access to the system. The different levels of information relate to the knowledge of an attacker about mechanisms to attack the system, and span from no particular knowledge through to a detailed knowledge of matching algorithm internals. This could also be put in terms of attacker resources, from bedroom hacker to state sponsored terrorists. The different levels of access an attacker may have span from simple user access (verification, enrollment or both) through to administrator-level access, including source code and re-compilation access [12].

The degree of effort needed to discover, characterize and implement the vulnerability is also a relevant parameter in characterizing the threat. For many simple biometric systems with no functional liveness detection, this discovery effort may

be minimal, requiring only hours or days. More complex attacks involving surgery or complex prosthetic construction may take months or years.

12.5.2 The Test Object Approach

One of the most well known and influential biometric vulnerability studies was undertaken by Matsumoto in 2001 [20]. This research involved demonstrating the simple creation of artificial fingerprints from latent fingerprints.[2] Subsequent work by Matsumoto [19] has concentrated on what he calls the *test object* approach to testing.

The test object approach categorizes the classes of artifact attacks using set theory. Consider that from the universe of all physical objects, there is a set that are able to be enrolled and verified in a given system. Within the set of enrolled objects some (in the perfect system - all) will be humans. However, in practice it is highly likely that the universe of all physical objects will also contain some artificial objects that can also be enrolled and verified. There are hence four methods to create artifact attacks (Fig. 12.1)

Enrollment	Verification	Example
human	human	A person whose biometrics are naturally similar enough to be able to pass as someone else.
artificial	human	The use of an artifact to enroll as an impostor, then allowing that impostor later access as themselves.
human	artificial	Co-operative attack where a user allows someone to make an replica of the enrolled biometric for attack purposes.
artificial	artificial	An artifact is used during the enrollment, which can be transferred to another individual later.

Table 12.1 Methods to create artifact attacks [19].

12.6 Vulnerability Mitigations

Mitigations are the steps that can be taken to prevent a specific threat from being exploited. These might involve new sensing mechanisms, changes to the matching algorithm, cryptography, alterations to the environment or usage policy. Developing

[2] Latent prints are fingerprint impressions left on a surface after it has been touched.

a mitigation to treat identified vulnerability risks is a vital task for critical systems. Mitigation strategies can be broken into the following categories:

- **Multi-factor Mitigations**: The use of multiple different authentication factors as part of the authentication significantly mitigates against a number of vulnerabilities, as an attacker needs to compromise more than one security mechanism. The use of either a smart-card and/or password/pin combination with a biometric is recommended practice for most secure authentication scenarios. Multimodal biometric solutions may also be used to mitigate risk.[3]
- **Sensor Mitigations**: Different types of sensors can be used to detect artifact attacks and ensure liveness. Examples include the detection of a pulse in fingerprint capture and the detection of reflex action of an iris to light.
- **Signal Processing Mitigations**: Without using any different sensor technologies additional signal processing can be applied to the detection of liveness. Examples include detecting the elasticity of skin as a fingerprint is pressed onto a sensor or noting the deformations expected by a real face compared to a photograph. Signal processing may also be used to check for a replay attack, by examining if a biometric sample is too similar to a biometric seen previously.
- **Behavioral Mitigations**: For biometrics that incorporate a behavioral element such as speech or typing dynamics, mitigation can be applied by asking the user to undertake some behavioral task that can be monitored. Examples include asking a user to speak a random digit string or getting a user to type a random word.
- **Coercion Mitigations**: Where a biometric characteristic has multiple instances such as a fingerprint or iris, one particular instance can be nominated to be used in a "panic" situation. For instance, when the "panic finger" is used it may still allow access but silently raise an alarm. Some biometrics can detect stress through changes to biological signals, such as pitch in voice or increasing pulse rate, however the false alarm rate is often unacceptably high due to the natural variation in such signals.
- **Environmental Mitigations**: The environment in which biometrics are captured can affect its biometric vulnerability. Where biometrics are captured in heavily monitored and policed areas, such as airports, they are more secure than when captured in a private and unmonitored area. Providing surveillance in areas where biometrics are used can greatly enhance the chance of detection of fraud and act as a deterrent to would be attackers.
- **Cryptography Mitigations**: Ensuring the secure transmission of data from each biometric component is vital where the transmission is sent over untrusted networks or insecure communication links. Mitigations may include using a public key infrastructure (PKI) to ensure match decisions are not altered after they are made.

[3] It is important to ensure that the multi-factor mitigations are, as much as possible, independent. Two different feature sets from the same physical region (e.g. fingerprint and skin pore, or face and iris) will make covert acquisition of both easier, and hence provide little additional additional security.

- **Tamper Mitigations**: The integrity of the sensor can be both electronically tested and physically secured to ensure that no modifications or substitution have been undertaken. Tamper-proofing might include physically sealing all the internal hardware in resin and using electronic sensors to detect if seals have been broken.
- **Policy Mitigations**: Policy mitigations are those instructions in the use of the system for both operators and users that ensure integrity. One of the most important is ensuring trust in the enrollment process by requiring human supervision, and having policies in place around system administration.
- **Monitoring Mitigations**: By installing an active monitoring system that looks for deviations in normal operational usage, potential fraud relating to system lambs can be determined, or attempted attacks can be determined through an analysis of time series data. Furthermore, the examination of audit logs can potentially reveal patterns of internal fraud.

Appropriate mitigations depend on the system requirements and the value of assets being protected. The trade-off may need to be carefully weighed against practicality since, in some cases, the imposition of a mitigation may have negative side-effects on usability, or may lead to falsely rejecting live input.

12.7 Conclusion

Quantifying the vulnerability of biometric systems and determining appropriate countermeasures is a vital area of research. This chapter has provided an overview of the potential attack points and fraud mechanisms in biometric systems, and an introduction to how they might be assessed for these vulnerabilities.

It is strongly argued that it is important to distinguish between the goal of the biometric matching algorithm, which is to robustly distinguish one person from all others, and the goal of anti-spoofing or liveness techniques, which is to ensure non-human objects are not matched. Both are necessary, but distinct, components of any secure system.

There is a bright future for biometrics systems to enhance and simplify all our interaction with technology. However, in this future the technology will need continued focus on both accuracy and vulnerability.

References

[1] Biometric device protection profile BDPP. `http://www.cesg.gov.uk/site/iacs/itsec/media/protection-profiles/bdpp082.pdf` (2001)

[2] Communications security establishment certification body canadian common criteria evaluation and certification scheme. `http://www.cse-cst.gc.ca/documents/services/ccs/ccs_biometrics121.pdf` (2001)

[3] U.S. government biometric verification mode protection profile for basic robustness environments. http://www.niap.bahialab.com/cc-scheme/pp/pp_bvm_mr_v1.0.pdf (2001)
[4] Common criteria common methodology for information technology security evaluation: Biometric evaluation methodology supplement BEM. http://www.cesg.gov.uk/site/ast/biometrics/media/BEM_10.pdf (2002)
[5] Transcript: Defense department briefing. http://www.america.gov/st/washfile-english/2002/October/20021017192919ross@pd.state.gov0.9141504.html (2002)
[6] Episode 59 - crimes and myth-demeanors 2. http://en.wikipedia.org/wiki/MythBusters_(season_4)#Episode_59_.E2.80.94_.22Crimes_and_Myth-Demeanors_2.22 (2006)
[7] Adler, A.: Sample images can be independently restored from face recognition templates. Electrical and Computer Engineering, 2003. IEEE CCECE 2003. Canadian Conference on **2** (2003)
[8] Boyce, C., Ross, A., Monaco, M., Hornak, L., Li, X.: Multispectral iris analysis: A preliminary study. Proc. Conf. Computer Vision and Pattern Recognition Workshop pp. 51–59 (2006)
[9] Czajka, A., Strzelczyk, P., Pacut, A.: Making iris recognition more reliable and spoof resistant. SPIE The International Society for Optical Engineering (2007)
[10] Daugman, J.: Iris Recognition and Anti-Spoofing Countermeasures. 7th International Biometrics Conference (2004)
[11] Drahansky, M., Lodrova, D.: Liveness detection for biometric systems based on papillary lines. International Conference on Information Security and Assurance, 2008. ISA 2008. pp. 439–444 (2008)
[12] Dunstone, T., Poulton, G., Roux, C.: Update, Biometrics Institute vulnerability assessment project. In: The Biometrics Institute, Sydney Conference (2008)
[13] Faundez-Zanuy, M.: On the vulnerability of biometric security systems. Aerospace and Electronic Systems Magazine, IEEE **19**(6), 3–8 (2004)
[14] Godesberger, A.: Common criteria protection profile biometric verification mechanisms, german federal office for information security (bsi). http://www.bsi.bund.de/zertifiz/zert/reporte/PP0016b.pdf (2005)
[15] Harrison, A.: Hackers claim new fingerprint biometric attack. http://www.securityfocus.com/news/6717 (2003)
[16] Hill, C.: Risk of masquerade arising from the storage of biometrics. Bachelor of science thesis, Dept. of CS, Australian National University (2002)
[17] Kryszczuk, K., Drygajlo, A.: Addressing the vulnerabilities of likelihood-ratio-based face verification. Proceedings of 6th International Conference on Audio-and Video-Based Biometric Person Authentication (AVBPA), T. Kanade and NR (AK) Jain, Eds., vol. LNCS **3546**, 426–435 (2005)
[18] Maltoni, D., Maio, D., Jain, A., Prabhakar, S.: Handbook of Fingerprint Recognition. Springer (2003)

[19] Matsumoto, T.: The test object approach in measuring security of fingerprint and vein pattern authentication systems. In: The Biometrics Institute, Sydney Conference (2008)
[20] Matsumoto, T., Matsumoto, H., Yamada, K., Hoshino, S.: Impact of artificial gummy fingers on fingerprint systems. In: Proc. of the SPIE, Optical Security and Counterfeit Deterrence Techniques IV, vol. 4677 (2002)
[21] Pan, G., Sun, L., Wu, Z., Lao, S.: Eyeblink-based anti-spoofing in face recognition from a generic webcamera. Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on pp. 1–8 (2007)
[22] Parthasaradhi, S., Derakhshani, R., Hornak, L.A., Schuckers, S.: Time-series detection of perspiration as a liveness test in fingerprint devices. Systems, Man and Cybernetics, Part C, IEEE Transactions on **35**(3), 335–343 (2005)
[23] van der Putte, T., Keuning, J., Origin, A.: Biometrical fingerprint recognition: Don't get your fingers burned. Smart Card Research and Advanced Applications: Ifip Tc8/Wg8. 8 Fourth Working Conference on Smart Card Research and Advanced Applications, September 20-22, 2000, Bristol, United Kingdom (2000)
[24] Schuckers, S.: Spoofing and anti-spoofing measures. Information Security Technical Report **7**(4), 56–62 (2002)
[25] Statham, P.: UK government biometrics security assessment programme, cesg biometrics. `http://www.biometrics.org/bc2004/CD/PDF_PROCEEDINGS/bc247a_Statham.ppt` (2003)
[26] Thallheim, L., Krissler, J., Ziegler, P.: Body check: biometrics defeated. `http://www.extremetech.com/print_article/0,3998,a=27687,00.asp` (2002)
[27] Uludag, U., Jain, A.: Attacks on biometric systems: a case study in fingerprints. Proceedings of SPIE **5306**, 622–633 (2004)