This article was downloaded by: *10.2.97.136*
On: 10 Jun 2023
Access details: *subscription number*
Publisher: *CRC Press*
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK

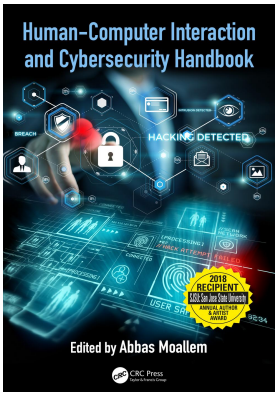

Human-Computer Interaction and Cybersecurity Handbook

Abbas Moallem

Biometrics

Publication details
https://test.routledgehandbooks.com/doi/10.1201/b22142-2
Francisco Corella, Karen Lewison
Published online on: 24 Oct 2018

How to cite :- Francisco Corella, Karen Lewison. 24 Oct 2018, *Biometrics from:* Human-Computer Interaction and Cybersecurity Handbook CRC Press
Accessed on: 10 Jun 2023
https://test.routledgehandbooks.com/doi/10.1201/b22142-2

PLEASE SCROLL DOWN FOR DOCUMENT

Full terms and conditions of use: https://test.routledgehandbooks.com/legal-notices/terms

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

chapter two

Biometrics

Francisco Corella and Karen Lewison

Contents

2.1 Introduction

After more than a century of biometric usage for forensics and physical identification, and decades of evolution of biometric technology at a measured pace, we are now in the middle of a biometric revolution, fueled by the availability of biometric sensors in smartphones, tablets, and laptops and, more recently, by breakthroughs in biometric technology. The usage of biometrics is becoming routine in everyday life, as fingerprints are used for unlocking smartphones, selfies are used for online identity proofing, and banks use speaker verification for customer authentication; and machines are now claimed to be better than humans at identifying faces (Lu and Tang 2015).

At the same time, biometric technology is facing severe security and privacy challenges. Security challenges come from spoofing techniques, which are improving just as

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

quickly as biometric accuracy. It is possible to spoof a fingerprint reader with an artifact constructed after photographing a finger with a high resolution camera from several meters away (Khandelwal 2014); a highly accurate deep neural network for face recognition may be deceived by colored eyeglass frames (Sharif et al. 2016); real-time voice morphing may allow an impersonator to fool not only speaker recognition software, but also even human verifiers (Mukhopadhyay et al. 2015). A privacy challenge arises from the fact that biometric characteristics used for biometric verification may be used to link the activities of the subject across both cyberspace and the physical world. If, for example, a selfie is used for authentication at a web site, the site operator or an adversary who breaches the database of users of the site can use the selfie to link each user's account to the user's activities on social networks and the user's visits to physical stores equipped with customer identification cameras.

Biometric usage for human–computer interaction is mostly concerned with *biometric recognition*, or *matching*. Biometric recognition can be divided into two categories, one-to-many matching, also called *identification*, and one-to-one matching, also called *verification*. In identification, the biometric task is to assign a biometric sample as belonging to one out of a large collection of individuals, while in verification, the task is to decide whether or not a sample belongs to a given individual. Use cases of biometric recognition include forensics, surveillance, photo tagging, and identification of customers who walk into a store. Use cases of biometric verification include physical access control, authentication of automated teller machine customers, phone unlocking, remote identity proofing, authentication, and privilege escalation. Of the two categories, biometric verification is the one most relevant to human–computer interaction, and is therefore the focus of this chapter.

The rest of the chapter is organized as follows: Section 2.2 introduces basic concepts related to biometric verification. Section 2.3 describes four paradigms commonly used today for biometric verification, which make use of biometric templates, statistical models, deep neural networks, and biometric cryptosystems (a.k.a. revocable biometrics or biometric key generation). Section 2.4 discusses biometric security, including zero-effort attacks, presentation attacks and their mitigation, and security architectures. Section 2.5 describes in some detail biometric modalities most commonly used today and other modalities in less detail; it also discusses the fusion of biometric modalities. Section 2.6 concludes by suggesting possible avenues for future research in biometric verification.

2.2 Biometric verification concepts

A *biometric characteristic*, or *trait*, is a measurable aspect of the human body that can be used to distinguish individuals from each other, such as a fingerprint, an iris image, a facial image, or acoustic features of the human voice. A *biometric sample* is a sample of a biometric characteristic. A *biometric modality* is a class of biometric systems that deal with a particular biometric characteristic.

In biometric verification, a verifier compares two biometric samples and decides whether they come from the same individual. Biometric verification is a two-phase protocol. In an *enrollment phase*, an enrollment sample is acquired from a subject. In a subsequent *verification phase*, the subject provides biometric input comprising a verification sample to a verifier, which compares it to the enrollment sample or to data derived from the enrollment sample. In addition to the verification sample, the biometric input may provide clues that the verifier can use for *presentation attack detection*, as described in Section 2.4.2.3.

There are two kinds of biometric verification, which are often called *authentication* and *identity proofing*. In authentication, the subject presents the enrollment sample to the

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

verifier. If the verification sample later matches the enrollment sample (or a template or other enrollment data derived from the enrollment sample), the verifier learns that the individual presenting the verification sample is the same subject who provided the enrollment sample, but nothing else. In identity proofing, the verifier, which may have no prior relationship with the subject, obtains the enrollment sample or other enrollment data from an *identification authority* with which the subject has enrolled earlier, together with a binding of the enrollment data to attributes of the subject. If the verification sample matches the enrollment data, the verifier learns that the attributes belong to the subject presenting the verification sample and uses those attributes to identify the subject.

The verification sample is said to be *genuine* if it comes from the same subject as the enrollment sample. For a given configuration of the verifier, the accuracy of the verification process may be defined by two probabilities: the probability that the presented sample is accepted as genuine when it is not genuine, called the *false accept rate* (FAR) or the *false match rate* (FMR), and the probability that it is rejected when it is in fact genuine, called the *false reject rate* (FRR) or *false nonmatch rate* (FNMR).

The verifier may be configured to be more or less forgiving of differences between the presented sample and the enrollment sample. At one extreme, FAR = 0, while FRR = 1. At the other extreme, FAR = 1, while FRR = 0. As FAR increases, FRR decreases; therefore, if the FAR and the FRR are modeled as continuous functions of a real-valued configuration parameter, there is a value of the parameter for which the FAR and the FRR have the same value, which is called the *equal error rate* (ERR). If the FAR is further modeled as a strictly monotonic function of the configuration parameter, each FAR determines a value of the parameter, which in turns determines a value of the FRR. The function that thus maps each FAR to the corresponding FRR is called the *receiver operating characteristic* (ROC). (The term *accuracy rate* is typically used in the context of identification rather than verification, to refer to the complement of the classification error rate.) Accuracy metrics are highly dependent on the sample space and can only be assigned precise numeric values when the sample space has been defined, for example, with reference to a database of biometric samples used for benchmarking, such as the Labeled Faces in the Wild (LFW) database (University of Massachusetts 2017).

2.3 Biometric matching paradigms

There is a great variety of techniques for comparing biometric samples, across biometric modalities as well as within each modality. They can be roughly classified into four different paradigms, based on whether they use biometric templates, statistical models, deep neural networks, or biometric cryptosystems. For some modalities, techniques pertaining to multiple paradigms are available.

2.3.1 Biometric matching using biometric templates

In this paradigm, a *biometric template* is derived from the enrollment sample and matched against the verification sample or against a template derived from the verification sample.

A biometric template is an encoding of characteristic features of a biometric sample. The order of the features encoded in the template may or may not be significant. If it is significant, the template is a *feature vector*, or an encoding of a feature vector. If not, it is a *feature set*, or an encoding of a feature set. An example of a feature set is a fingerprint template consisting of a set of minutiae. A fingerprint minutia is either the end of a friction ridge or a bifurcation of a friction ridge. Each minutia is described in the template by its

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

type (end or bifurcation), its position, and its orientation. An example of a feature vector is an iris code (Daugman 2003) described in Section 2.5.2.

2.3.2 Biometric matching using statistical models

In this paradigm, multiple enrollment samples are used to construct a statistical model of the subject's biometric characteristic. A general model of the biometric characteristic is also constructed using samples from a large number of individuals, and a statistical test is used to estimate the likelihood that the verification sample comes from the subject rather than a random individual. An example of this paradigm is the Gaussian mixture model–universal background model (GMM-UBM) verification method (Reynolds et al. 2000) often used in voice biometrics.

2.3.3 Biometric matching using a deep neural network

Deep neural networks, further described in Section 2.5.3, are multilayer artificial neural networks that are being used very successfully in applications such as facial and speech recognitions. When a deep neural network is used for face verification, the network is trained with millions of labeled faces belonging to thousands of people, but there is no need to specifically train the network with enrollment samples of the subject. The enrollment and verification samples are separately input to the network, which produces a mathematical output for each sample. The outputs are then compared according to some similarity metric and deemed to belong to the same person if their similarity metric is above a certain threshold. In the case of Google's FaceNet (Schroff et al. 2015), the output is a vector with 128 coordinates, each of which is a single byte, and the similarity metric used to compare the vectors derived from enrollments and verification samples is the Euclidean distance between the two vectors.

2.3.4 Biometric matching using a biometric cryptosystem

As illustrated in Figure 2.1, in a biometric cryptosystem (ISO/IEC 2011, Rathgeb and Uhl 2011), error correction techniques are used to consistently generate a *biometric key* from varying but genuine biometric samples. At enrollment time, an enrollment biometric template is derived from an enrollment sample, and a random biometric key and *helper data* are generated from the enrollment template and random bits produced by a random or pseudorandom bit generator (NIST 2016). At verification time, a verification biometric template is derived from a verification sample, and an error correction algorithm attempts to recover the biometric key from the verification template and the helper data. If the verification sample is genuine, the error correction algorithm is able to recover the key with a probability equal to the complement of the FRR, 1 − FRR.

Even though the helper data are derived from the enrollment template, randomization makes it computationally unfeasible to derive any useful biometric information from it. Thus, the confidentiality of the subject's biometric information is preserved even if an adversary captures the helper data. By contrast, traditional biometric templates reveal biometric information (Cappelli et al. 2007, Ross et al. 2007).

Different kinds of biometric key generation techniques are used with different kinds of biometric templates. Techniques based on the concept of a *fuzzy commitment* (Juels and Wattenberg 1999) may be used with feature vectors, where the order of the features matters, while techniques based on the concept of a *fuzzy vault* (Juels and Sudan 2006) may be used with feature sets, where the order does not matter.

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

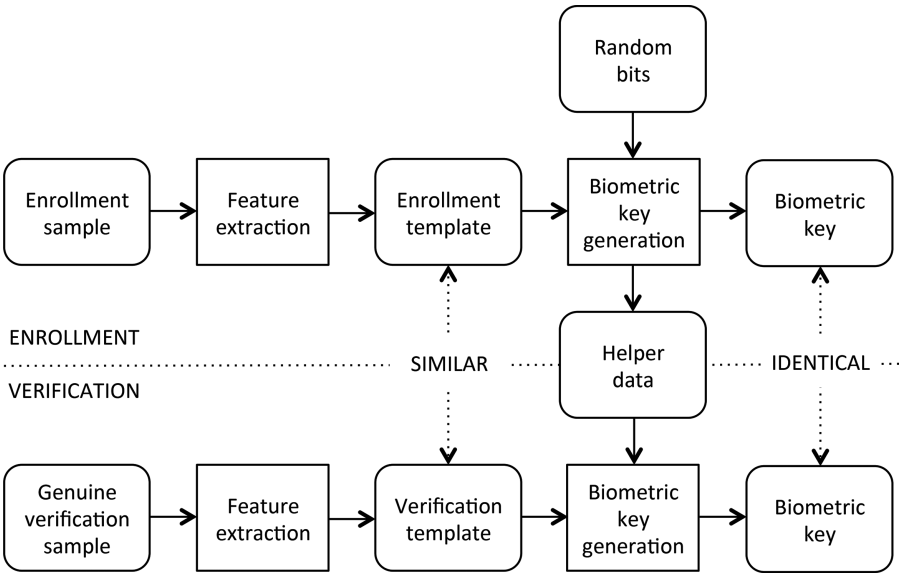

Figure 2.1 Biometric cryptosystem.

A biometric cryptosystem can be used for a variety of purposes. For example, the biometric key can be used to encrypt data. In such a use case, if the biometric key is compromised, it can be replaced with a different random key generated from the same biometric characteristic of the subject, and the data can be encrypted anew with the replacement key. The biometric key is said to be revocable, and this motivates referring to biometric cryptosystem technology as *revocable biometrics*, or *cancelable biometrics*.

On the other hand, when a biometric cryptosystem is used for biometric matching as discussed here, the biometric key is not used for encryption or any other cryptographic use. It is used to check whether the verification sample is genuine, by verifying that the error correction algorithm is able to produce the same key that was generated at enrollment time. The biometric key cannot be stored along with the helper data for that purpose, because the helper data and the biometric key together do reveal biometric information. But a cryptographic hash of the biometric key can be stored together with the helper data, and the biometric key produced at verification time can be verified by hashing it and comparing the resulting hash to the stored hash.

The use of a biometric cryptosystem for biometric matching raises a practical difficulty. Neither the enrollment sample nor the enrollment template is available at verification time. Therefore, it is not possible to perform any geometric alignment of the enrollment and verification samples or templates, in modalities that require such alignment. Methods for solving this difficulty are discussed in Section 2.5 in connection with fingerprint and iris modalities.

2.4 Biometric security

The goal of biometric verification is to prevent the impersonation of the subject by an adversary, and that requires protecting against a variety of attacks that may be carried out by the adversary. The adversary may carry out a *zero-effort attack* by presenting a biometric sample from his or her own body and hoping that it will be accepted as genuine.

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

Biometric accuracy mitigates zero-effort attacks. However, the sample presented by the adversary may not come straight from the adversary's body. It may come from an artifact, or the adversary may wear a disguise, or the sample may be a digital transformation of a sample originating from the adversary, or it may be a digital copy of a genuine sample coming from the subject's body. Attacks with such samples are *presentation attacks*, informally known as *spoofing attacks*.

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

2.4.1 Presentation attacks

To understand presentation attacks and how to provide protection against them, it helps to classify them along the following dimensions:

- A presentation attack may be *physical* or *digital*, according to whether it is performed before or after a sensor has digitized the biometric sample.
- The presented sample may be *artificial*, if it is produced by a physical artifact or is digitally generated; *disguised*, if it comes from a physically or digitally disguised adversary; or *genuine*, if it originates from the impersonation victim.
- The target, i.e., the biometric characteristic of the subject that the adversary wants to impersonate, may be *known* or *unknown*.

These classification facets are illustrated by the following examples.

When a fake finger is used to hack the fingerprint sensor of a smartphone (Chaos Computer Club 2013), or a photo of the impersonation victim is presented to a smartphone camera in an attack against face verification, the attack is *physical*, the sample is *artificial*, and the target is *known*.

When a MasterPrint (Roy et al. 2017) is used to make a fake finger that is presented to a partial fingerprint sensor, the attack is *physical*, the sample is *artificial*, and the target is *unknown*.

When a wig, a fake nose, and makeup are used to disguise an adversary against face verification (Pavlidis and Symosek 2000), the attack is *physical*, the sample is *disguised*, and the target is *known*.

When colored eyeglass frames are used in a perturbation attack against a deep neural network (Sharif et al. 2016), the attack may be *physical* or *digital*, the sample is *disguised*, and the target is *known*.

When voice morphing is used to disguise the voice of an adversary reading prompted text in an attack against speaker verification (Mukhopadhyay et al. 2015), the attack is *digital*, the sample is *disguised*, and the target is *known*.

When a video of the impersonation victim is obtained by a malicious verifier and replayed to another verifier, the attack is *digital*, the sample is *genuine*, and the target is *known*.

2.4.2 Protection against presentation attacks

2.4.2.1 Biometric confidentiality

Biometric characteristics of an individual are not secrets, and hoping that the adversary does not know the target characteristic should not be the only defense against impersonation. However, the biometric characteristics used in some biometrics modalities, such as iris or retina verification, may be difficult for an adversary to acquire without the subject's consent. Furthermore, the adversary may not know the target characteristic because the

adversary does not know the identity of the target subject. Therefore, efforts to protect the confidentiality of a biometric characteristic are a useful mitigation against attacks that require the target characteristic to be known, besides being motivated by privacy considerations.

Biometric confidentiality can be protected by presenting samples only to trusted verifiers over secure connections and by protecting databases containing biometric enrollment samples or templates against security breaches. It can also be protected by using a biometric cryptosystem as described in Section 2.3.4.

2.4.2.2 Combination with a password

A biometric sample can acquire secrecy by combining it with a password or passphrase. In text-dependent speaker verification, a short sample text becomes a passphrase simply by treating it as a shared secret between the subject and the verifier (Novoselov et al. 2014). A password has also been used in combination with lip reading (Cheung 2017) and was used, two decades ago, in combination with behavioral biometrics based on keystroke dynamics (Monrose et al. 1999).

2.4.2.3 Presentation attack detection

While the confidentiality protection and combination with a password are useful mitigations, the best defense against presentation attacks is presentation attack detection.

Different techniques have been proposed for detecting different kinds of attacks against different modalities. In some modalities such as fingerprint verification, presentation attack detection is performed by the sensor. In other modalities, such as iris, face, or speaker verification, presentation attack detection is performed on the digital output of the sensor.

Challenge–response is a technique available against attacks where the adversary presents a genuine sample obtained from the subject. In the face verification modality, for example, the verifier may ask the subject to stream a video of him/herself reading a challenge sequence of digits, rather than a static selfie. To detect a possible replay attack, the verifier may then use a lip reading technique to check that the digits being read are those in the challenge sequence (Kollreider et al. 2007). In any challenge–response interaction, the challenge should be chosen by the verifier at random with high entropy.

To detect attacks that use an artificial sample coming from a physical artifact, the verifier may check for the presence or absence of signs indicating that the sample comes from a live human body. The absence of such signs indicates a presentation attack. For example, a fingerprint sensor may look for indications of perspiration coming from the pores on the friction ridges (Schuckers and Johnson 2014), or the contraction of the pupil in response to brighter light may indicate liveness in iris scanning.

To detect a disguise worn by the adversary, the verifier may look for specific kinds of disguise, which may require using additional sensors. For example, disguises, such as makeup, a fake nose, or even a wig made out of human hair, are revealed by imaging in the upper near-infrared (IR) spectrum (0.8–1.4 μm) (Pavlidis and Symosek 2000). Some smartphones have near-IR cameras, which are used for iris scanning (Mayhew 2016) but could be used for other purposes in the future.

Protection against presentation attacks with digitally generated or modified samples is an open area of research.

Multiple presentation attack detection techniques may be used together for protection against different kinds of attacks on the same modality. For example, face verification may be exposed both to replay attacks and disguise attacks. The preceding lip reading

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

challenge–response technique may be used together with imaging from a near-IR camera to protect against both attacks.

2.4.2.4 *Remark: Liveness detection*

The term *liveness detection* is sometimes used as a synonym of presentation attack detection. Strictly speaking, however, liveness detection should refer to presentation attacks with samples that are not live. The word *live* may refer to the real-time presentation of a sample or to a sample that comes from a live body. Hence, liveness detection may refer to the detection of replay attacks with genuine samples or the detection of samples that come from artifacts.

2.4.3 *Security and privacy implications of biometric verification architectures*

Biometric verification for human–computer interaction involves components, such as a sensor, enrollment data, and biometric matching software, and devices such as a smartphone, a personal computer, a smartcard, or a server. A biometric verification architecture determines what components reside on what device. There is a wide variety of possible architectures, ranging from an old-fashioned one where a fingerprint is obtained by a sensor attached to a desktop and compared to a template stored in a smartcard plugged into a card reader also attached to the desktop to more recent ones such as where a credit card is equipped with a fingerprint sensor (Mastercard 2017). This section examines the security and privacy implications of four architectures commonly used today, as illustrated in Figures 2.2 through 2.5.

In Figures 2.2 and 2.3, the subject locally authenticates to a personal device, such as a smartphone or a laptop, by presenting a biometric sample to a sensor located on the device. The purpose of the authentication may be to unlock the device or to authorize a secondary nonbiometric authentication to a remote server. The latter purpose is the goal of the Fast IDentity Online (FIDO) Universal Authentication Framework (FIDO Alliance 2016), where the secondary authentication to the remote server is by means of an uncertified key pair.

Figures 2.2 and 2.3 differ by the kind of presentation attack detection that is performed, if any, which depends on the kind of modality and sensor that are used. In Figure 2.2,

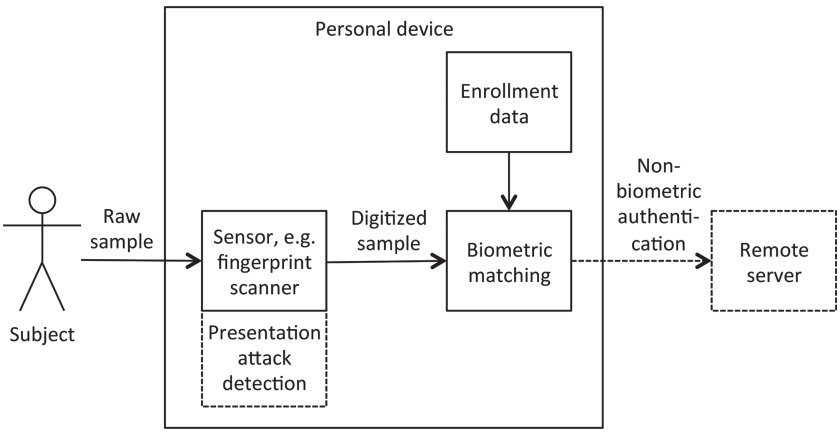

Figure 2.2 Local authentication with presentation attack detection by sensor.

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

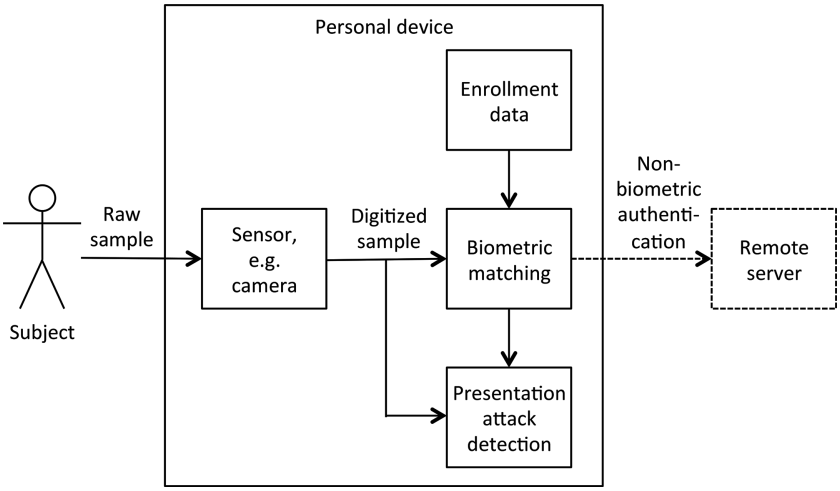

Figure 2.3 Local authentication with presentation attack detection performed on a digitized sample.

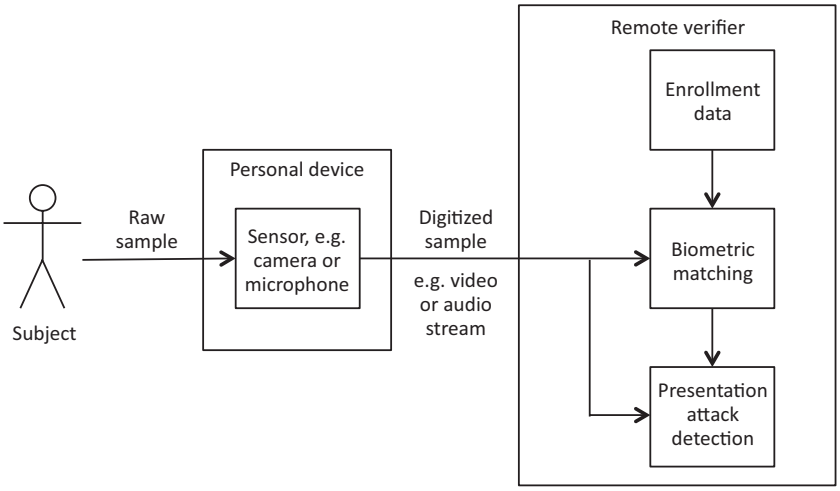

Figure 2.4 Remote authentication or identity proofing against a database.

presentation attack detection is carried out by a sensor such as a fingerprint scanner, or is omitted. Today, fingerprint scanners found on smartphones do not perform any presentation attack detection, but they may do so in the future, as discussed in Section 2.5.1. In Figure 2.3, presentation attack detection is carried out on the digitized output of the sensor, which may be, for example, a video captured by a camera.

The architectures in Figures 2.2 and 2.3 provide strong biometric privacy protection because biometric samples are never sent outside the personal device. The protection is even stronger if the enrollment data are stored in tamper-resistant hardware such as a secure element or a Trusted Platform Module (TPM), as is sometimes the case for the fingerprint template in some smartphones equipped with a fingerprint sensor. In that case, the subject's biometric information is protected against an adversary who physically captures the subject's device. Equivalent protection can be achieved by using a biometric

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

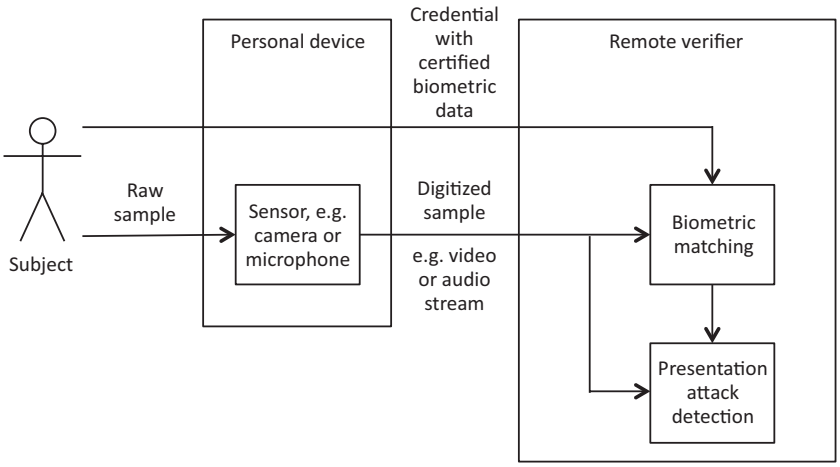

Figure 2.5 Remote authentication or identity proofing against certified enrollment data.

cryptosystem and using helper data together with a hash of a biometric key to verify the authentication sample, as described in Section 2.3.4.

The security provided by the architectures in Figures 2.2 and 2.3 depends on the modality, the sensor, and the efficacy of the presentation attack detection. It is poor today when a fingerprint sensor is used without presentation attack detection, as discussed in Section 5.1.

In Figures 2.4 and 2.5, the subject presents a biometric sample to a sensor located in a personal device, which digitizes the sample and forwards it to a remote verifier. The sensor may be, for example, a camera that streams a video of the subject's face to the verifier for face verification or a microphone that streams speech uttered by the subject for speaker verification. Presentation attack detection is performed on the digital output of the sensor transmitted to the remote verifier. These architectures may be used both for recurring authentication against an enrolled sample and for identity proofing to a verifier that has no prior relationship with the subject.

Figures 2.4 and 2.5 differ in the location of the biometric enrollment data.

In Figure 2.4, the enrollment data are stored in a database that is accessed by the remote verifier. In the authentication case, the database is built by the verifier as subjects enroll biometric samples. For example, the subjects may be users of an online service, and the database may be the user database of the service; each user record may then contain enrollment data used for authenticating the user. In the identity proofing case, the database is not specific to the verifier. It is provided instead by an identification authority such as, in the United States, a department of motor vehicles of a state (Gamlin 2016).

In Figure 2.5, the enrollment data are included in a credential and certified by a digital signature on credential data. The credential is submitted by the subject together with the biometric verification sample. It may be a physical token, such as a passport with an embedded near-field communication (NFC) chip (International Civil Aviation Organization 2015), a Personal Identity Verification (PIV) card carried by employees of the US Federal Government (NIST 2013), or a national identity card of a country that issues such cards. It may also be a purely digital credential, which may provide multifactor authentication, such as a biometric certificate (Dulude and Musgrave 2001) or a rich credential (Lewison and Corella 2016).

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

By contrast with the architectures in Figures 2.2 and 2.3, those of Figures 2.4 and 2.5 are not vulnerable to the physical capture of the subject's personal device.

The biometric architecture in Figure 2.4 has a serious privacy drawback, because, given the current state of cybersecurity, the database containing the enrollment data may be breached, and if so, the adversary may capture at once biometric information pertaining to a large number of subjects. This drawback may be eliminated by using a biometric cryptosystem for biometric matching and using helper data for verification together with a hash of a biometric key, as described in Section 2.3.4. In Figure 2.5, by contrast, an adversary can only attempt to capture biometric information of one subject at a time.

2.5 Biometric modalities

2.5.1 Fingerprint verification

The biometric characteristic that is measured in fingerprint verification is the pattern of the friction ridges of a finger, which can be observed by different kinds of sensors at different resolutions. Most sensors used today, for example, those on smartphones, are capacitance-based sensors that measure the variations in electrical capacitance between the finger and an array of microelectrodes embedded in the sensor, the capacitance measured under ridges being higher than the capacitance measured under valleys.

Different features can be extracted from the friction ridge pattern at different resolutions. At 500 pixels per inch (PPI), which is the resolution of today's smartphone sensors, the extracted features are the minutiae already described in Section 2.3.1. At lower resolutions, the extracted features are ridge shapes known as arches, loops, and whorls. At higher resolutions, it is possible to measure the thickness of the ridges and to count the perspiration pores located on the ridges. It is even possible to detect pores that are open and emitting perspiration onto adjacent valleys, which can be used for liveness detection (Schuckers and Johnson 2014).

Higher resolutions may be available in future smartphones by using optical rather than capacitance sensors. US patent 9,570,002 (Sakariya and Nauta 2017) refers to inserting IR light-emitting diodes and sensing IR diodes between the subpixels of the phone screen in order to image the ridges of a fingerprint placed on the screen, potentially achieving higher resolution.

At 500 PPI, a fingerprint template may be a list of minutiae, each described by its type, its x and y coordinates, and the angular orientation θ of the ridge that ends or bifurcates at the minutia. Matching two templates requires finding a correspondence between some of the minutiae in one template and some of the minutiae in the other. A matching score can then be computed from parameters including, among others, the number of minutiae that have been matched.

Minutiae can be viewed as an unordered set of features, and a biometric cryptosystem based on the concept of a fuzzy vault can be used for biometric matching in the fingerprint modality. However, this requires corresponding minutiae to be mapped to elements of a finite field. If the mapping is based on the x and y coordinates and orientation θ of the minutiae, measured at some granularities, enrollment and verification samples must be digitally aligned precisely enough for x, y, and θ to have the same values at those granularities. However, as pointed out in Section 2.3.4, the enrollment sample and template are not available at verification time. The first biometric cryptosystem proposed for fingerprint matching (Clancy et al. 2003) avoided this problem by requiring the physical alignment of the samples, but this is not practical. A subsequent proposal (Uludag and Jain 2006)

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

addressed the problem by adding ridge shape information to the helper data. But this goes counter to the essential tenet that helper data should not reveal any useful biometric information. More recently, there have been proposals to make the mapping of minutiae to elements of the finite field independent of geometric alignment (Wencheng et al. 2014, Li and Hu 2016).

The fingerprint sensors used today to unlock smartphones provide little security, for two reasons. The first reason is that due to the absence of any presentation attack detection, the sensor can be hacked with an artifact constructed from a latent print lifted from the phone itself. The iPhone TouchID was thus hacked shortly after it was introduced (Chaos Computer Club 2013, Rogers 2013). A fingerprint can also be obtained for that purpose from a photograph of a finger (Khandelwal 2014). The sensor on the Samsung Galaxy S5 was also hacked, four days after being introduced (Storm 2014). An improved version of TouchID was introduced for the iPhone 6, but it was also hacked in the same way as that for the iPhone 5S (Bort 2014).

The second reason is that fingerprint sensors on today's smartphones capture only partial fingerprints. A partial fingerprint template has less entropy and is more likely to be matched by a zero-effort attack than a full template. Furthermore, in order to avoid false rejection when the partial verification fingerprint does not match the partial enrollment fingerprint, users are allowed to record multiple partial fingerprints, further reducing entropy and increasing exposure to zero-effort attacks. Also, some partial templates have been shown to occur with higher probability than others. An adversary can further increase his/her chances by constructing an artifact that produces such a template and applying it to the sensor (Roy et al. 2017). In the future, it should be possible to scan full rolled fingerprints with a sensor embedded in the smartphone screen as described in the aforementioned US patent 9,570,002.

Fingerprint verification is also used on smartphones for purposes other than unlocking the phone. It is used in particular to secure payments both online and in stores equipped with NFC payment terminals. Mastercard has recently announced an alternative way of using a fingerprint to secure payments in stores, by means of a sensor located on a Europay, Mastercard, and Visa (EMV) chip card (Mastercard 2017). The fingerprint is matched against an "encrypted" template stored in the chip. However, if the template is encrypted, it is not clear what key can be used to decrypt it that would not also be available to an adversary who tampers with the card to obtain the template.

2.5.2 *Iris verification*

In the iris verification modality, the biometric characteristic that is used is the texture of the iris, as imaged in the near-IR spectrum where it exhibits a richer structure than in the visible wavelength spectrum, particularly for brown eyes. Iris verification is a very accurate modality. Based on 200 billion cross comparisons performed on a database of 632,500 iris images acquired at border crossings (Daugman 2006), the FMR was estimated to be 10^{-15} when setting a threshold of 0.225 for a matching score (best of seven Hamming distances of iris codes at different relative rotations, as described in the following) for which the FNMR was estimated to be less than 1%. Until recently, the most common use of iris verification was for traveler identification at border crossings, but some phones now feature a near-IR camera that is used for unlocking the phone by means of iris verification (Mayhew 2016).

Iris verification uses algorithms for computing a 2048-bit biometric template, called an *iris code*, invented and patented by John Daugman in the 1990s (Daugman 1994) and still

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

in use today. An iris code is computed in three steps. First, the pixels of the iris are located in the near-IR image of the eye by identifying the boundaries between the iris and the pupil on one hand and the iris and the limbus on the other hand. Second, areas where the view of the iris is obstructed by eyelids, eyelashes, or specular reflections are identified. Third, the 2048-bit iris code is constructed by demodulating the iris pattern with pairs of quadrature two-dimensional Gabor wavelets, extracting spatial phase information, but discarding amplitude information that depends on imaging contrast, illumination, and camera gain (Daugman 2003). Each bit of the iris code corresponds to a relative position in the iris defined by pseudopolar coordinates. The pseudopolar coordinates eliminate the dependency on the dilation and contraction of the pupil, but not on its rotation, which depends on head tilt, torsional eye rotation within its socket, camera angle, etc. An iris image is verified against an enrollment image by comparing an iris code computed from the enrollment image to seven iris codes computed from the verification image using seven angular origins for the pseudopolar coordinates. A modified Hamming distance is computed between the enrollment iris code and each verification iris code by counting the number of bits that differ between the codes, ignoring bits that belong to areas where the view of the iris is obstructed in either of the images. The smallest of the seven modified Hamming distances is used as a matching score.

A biometric cryptosystem for iris verification has been described (Hao et al. 2006). Since an iris code is a feature vector where order matters, it is based on the concept of a fuzzy commitment rather than a fuzzy vault. At enrollment time, a random biometric key K is generated; a commitment C to K is computed using a cryptographic hash function, $C = \text{hash}(K)$; a 2048-bit error correction codeword W is obtained by adding redundancy to K; an enrollment iris code E is computed from the enrollment iris image; and the helper data H are computed by x-oring E and W, $H = E \text{ xor } W$. At verification time, seven iris codes are generated from the verification image for seven angular origins of the polar coordinates, and each verification code V is xored with the helper data to compute $W' = V \text{ xor } H = V \text{ xor } (E \text{ xor } W) = (V \text{ xor } E) \text{ xor } W$. If the verification image is genuine and the angular origin of V results in good alignment, then the iris codes V and E differ only in a few bits, W and W' differ only in those same bits, and the error correction system may be able to recover W from W', compute K by removing the redundancy from W, and compute $C = \text{hash}(K)$. Verification succeeds if this process successfully produces C for one of the seven iris codes.

This verification method is not able to ignore the bits that correspond to areas where the view of the iris is obstructed in one of the iris images, because the enrollment image is not available at verification time. Hao et al. compensate for this by combining two error correction techniques, one for random errors and one for burst errors caused by the presence of such areas of obstruction.

While iris verification has a low FMR, it is only secure against an adversary who attempts to impersonate the subject if the adversary does not have an iris image of the subject or effective presentation attack detection is implemented. The detection of presentation attacks against iris verification is a nascent area of research (Raghavendra and Busch 2015, Czajka 2016).

2.5.3 *Face verification*

Following three decades of the steady development of traditional techniques in academia, the field of face verification has been disrupted in the past five years by two independent phenomena: the advent of deep learning and the appearance of commercial face verification systems. Deep learning systems and commercial systems have both furthered the

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

state of the art of face verification, each in its own way, but both have been lacking in protection against presentation attacks.

Traditional techniques for face verification rely on methods for face detection such as Viola–Jones (Viola and Jones 2004) or Histograms Oriented Gradients (HOG) (Dalal and Triggs 2005) and a broad variety of methods for face matching. Face matching methods may be based on linear subspace analysis, where the facial images to be compared are projected onto a linear subspace determined by training images, and a similarity score is computed from the resulting coefficients (Yambor et al. 2002); or on comparing the texture of the facial images as described by Local Binary Patterns (LBP) (Ojala et al. 1996); or on identifying landmark points on each face, linking them into graphs that model the faces and comparing the models (Wiskott et al. 1997). Traditional techniques are challenged by pose, illumination, and expression variations. Three-dimensional (3D) imaging techniques have been proposed to help with those challenges, but some of them require complex imaging setups, such as using multiple cameras, which are not practical for most use cases.

Deep learning refers to machine learning using a deep neural network. Machine learning using artificial neural networks goes back to the 1950s, but its performance has improved spectacularly over the last few years for applications including face and speech recognition. These improvements are based on the use of deep neural networks (Zeiler and Fergus 2013, Goodfellow et al. 2016), composed of many layers of neurons (e.g., eight layers in Facebook's DeepFace or 22 layers in Google's FaceNet) and millions of parameters (120 million in DeepFace, 140 million in FaceNet) that are adjusted by training on millions of inputs. The training of such networks has been made possible by the use of arrays of graphics processing units (GPUs) with thousands of cores each. Deep neural networks are used by social networks and search engines for photo tagging, where they match or surpass human-level performance for facial image classification (Taigman et al. 2014, Lu and Tang 2015, Schroff et al. 2015). Photo tagging is an identification task rather than a verification task, but FaceNet is explicitly intended for use in both identification and verification.

While deep neural networks provide surprisingly accurate face recognition when used in a nonadversarial setting, they seem to be surprisingly weak when under attack. It is possible to compute quasi-imperceptible perturbations that cause natural images to be misclassified with high probability by state-of-the-art deep neural networks (Moosavi-Dezfooli et al. 2017). Independently, it has been shown that such a perturbation can be inconspicuously achieved with colored eyeglass frames and can be targeted for the impersonation of specific individuals (Sharif et al. 2016).

Face verification is now used commercially for many purposes, including identity proofing (Brinded 2016, Gamlin 2016), authentication (http://zoomlogin.com, http://keylemon.com), unlocking of phones (Mayhew 2016), and unlocking of password managers (http://biomids.com, http://1uapps.com). Most commercial face verification systems now use some form of presentation attack detection, but even advanced detection techniques were defeated by displaying a 3D model of a face reconstructed from photographs on a virtual reality system (Xu et al. 2016).

2.5.4 *Speaker verification*

In speaker verification, or voice verification, the biometric characteristic being measured consists of the acoustic features of the human voice. Speaker verification may be text dependent or text independent. Text-dependent verification has the advantage that, as noted earlier in Section 2.4.2.2, the text that is spoken may be a password, injecting secrecy into the biometric characteristic. Text-independent verification, on the other hand,

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

has the advantage that the text to be spoken may be prompted by the verifier for challenge–response protection against replay attack.

Since the year 2000 and to this day, most techniques used for text-independent speaker verification have been based on the GMM-UBM verification method (Reynolds et al. 2000). In this method, a statistical GMM is constructed for a hypothesized speaker based on cepstral analysis of training speech, and a UBM is constructed from speech samples obtained from a large number of speakers representative of the expected range of possible speakers. At verification time, a likelihood ratio test is applied to the verification speech sample to compare how much better it fits the hypothesized speaker model than the UBM. The resulting statistic is then compared to a configured threshold to decide whether to accept or reject the sample as genuine.

Recently, deep neural networks have been very successful at speech recognition, and this has motivated attempts at using them for speaker recognition as well. Today, deep neural networks are being used in combination with GMM-UBM techniques to improve performance (Richardson et al. 2015).

Text-independent speaker verification is being used by financial institutions and call centers to authenticate speakers during phone calls, using previously recorded calls as training speech (Barclays 2016, Citigroup 2017, Pindrop 2017a). Text-dependent speaker verification is being used as well (OCBC 2016).

The speaker verification modality has two drawbacks. One is practical: the acoustics of human voice change with age much faster than a fingerprint or an iris code. Coping with this may require adjusting the statistical speaker model after each successful verification.

The other is a serious security vulnerability. Voice morphing may be used to change the sound of human voice in real time. This is done for fun or game playing, using widely available and inexpensive commercial software to make the voice sound as coming from a younger or older person or from a person of a different gender. However, it can also be done to mount an impersonation attack against a specific individual. This can be done by building a model of the victim's voice from a very limited number of speech samples using acoustic-to-articulatory inversion mapping, according to a study (Mukhopadhyay et al. 2015). In the study, voice morphing was used to attack several speaker verification systems, which were only able to reject the fake voices at a rate of 10–20%. Even humans had difficulty identifying fake voices, which they rejected at a rate of 50%. The authors assert that the voice morphing technique that they used can be used in real-time communications, but this was not part of the study. Whether victim-specific voice morphing can be done in real time is important, because if so, voice morphing could be used against speaker verification systems currently used by financial institutions and call centers. Anecdotal evidence that voice morphing attacks are already being carried out in the wild can be found in a call center fraud report that attributes a 113% fraud rate increase from 2015 to 2016 at least in part to "voice distortion software" (Pindrop 2017b).

2.5.5 Other biometric modalities

A variety of biometric characteristics are used by other biometric modalities:

- *Retinal scanning* is based on the pattern of blood vessels in the retina, observed using a beam of IR light that scans the retina as the subject looks into the scanner.
- *Eye vasculature biometrics* is based on the pattern of veins in the sclera (the white part of the eye).

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

- *Finger vein recognition* is based on the pattern of surface veins in the finger, imaged with IR light while the finger is inside a scanner.
- *Electrocardiogram biometrics* is based on the cardiac rhythm, which may be measured by an NFC-enabled wristband.
- *Behavioral biometrics* is based on a pattern of human activity detected by a collection of signals such as keystroke dynamics, mouse movements, movements of the hand that holds a smartphone, details of touchscreen gestures, etc.
- *Gait* as observed by a camera is not a practical biometric for human–computer interaction, but it becomes practical if observed by the accelerometer and gyroscope in a smartphone carried by the subject while walking. Gait may then be a component of a broader collection of behavioral biometric signals.

2.5.6 Biometric fusion

Biometric fusion refers to the observation of multiple biometric characteristics, resulting in multiple biometric samples, for biometric verification. There are many ways of combining the multiple samples to reach a decision as to whether they are genuine or not. The samples may be different instances of the same biometric modality, such as fingerprints from multiple fingers or images of both irises, or pertain to different modalities. They may be acquired by one sensor or multiple sensors. If biometric templates are used for matching, the samples may be processed together to produce a joint template or separately to produce multiple templates. The decision may be based on a joint matching score or on separate matching scores. If multiple matching scores are computed, they may be used together to reach the decision, or they may be separately compared to thresholds to reach separate decisions and then combine the decisions, for example, by requiring all the matching scores to exceed their thresholds or only some of them. The following are notable examples of biometric fusion systems:

- Veridium's (2017) product 4 Fingers Touchless ID combines four fingerprints captured together by a camera.
- Derakhshani (2012) describes a biometric cryptosystem that combines eye vasculature with micro features found in the tear duct and below the lower eyelid.

2.6 Conclusion

The preceding sections have hopefully shed light on the security and privacy challenges mentioned in the introduction and possible ways of addressing them.

The security challenge is the need to provide protection against presentation attacks. There are two aspects to this challenge. One is the need to spread awareness of the threat of presentation attacks among implementers and users of biometric verification systems. Progress has been made in that respect over the last few years, but work remains to be done. The other is the need to address presentation attacks that are particularly difficult to cope with.

One class of such attacks is illustrated by virtual reality and voice morphing attacks in Sections 2.5.3 and 2.5.4. In those attacks, the adversary digitally constructs an alternate reality where the verifier sees what it needs to see in order to accept the biometric evidence as genuine. If the subject's biometric characteristic is known to the adversary, as it is prudent to assume, the adversary may have all the information needed to construct a digital sample indistinguishable from a genuine one. If the alternate reality is constructed in real

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

time, the adversary may also have all the information needed to respond to a challenge presented by the verifier. Protection against this class of attacks is an open area of research.

Another class of challenging attacks is illustrated by the universal perturbation and the eyeglass frame attacks in Section 2.5.3. Here the difficulty seems to come from a weakness in the current deep neural network technology. This calls for research in understanding the weakness and finding ways of eliminating or mitigating it. Such research is already under way.

In biometric verification, there will always be an arms race between verifiers and impersonators. Therefore, biometric verification should be used in combination with other methods of identity proofing and authentication, so that the emergence of an unforeseen method of attack is not catastrophic for verifiers.

As noted in the introduction, a privacy challenge arises from the linkability of a biometric characteristic across cyberspace and the physical world. When biometric matching is used for authentication, biometric information can be kept from the verifier by using the architectures in Figure 2.2 or 2.3. When biometric matching is used for identity proofing, a credential containing certified biometric data can be used to obviate the need for storing biometric information in a database vulnerable to hacking, as illustrated in Figure 2.5. If a database is used for identity proofing as illustrated in Figure 2.4, a biometric cryptosystem would make it possible to store in the database helper data that would reveal no useful biometric information to a hacker, instead of a traditional biometric template that leaks such data.

This suggests several areas of research that may lead to mitigations of the privacy challenge.

Digital credentials with certified biometric data protect biometric information against hackers by enabling the identity proofing architecture in Figure 2.5 and may further enhance privacy if they accommodate biometric cryptosystems and provide selective disclosure of attributes and selective presentation of verification factors, as rich credentials do (Lewison and Corella 2016). However, the widespread deployment of such credentials would require an ecosystem of issuers and verifiers who agree on standard protocols for issuance, presentation, and validation of the credentials. Research is needed into such protocols.

Biometric cryptosystems can be used to protect biometric information against hackers in the architecture in Figure 2.4, but little work has been done on biometric cryptosystems for face or voice verification. More research is needed in those areas.

Traditional biometric templates were once thought to hide biometric information, but, as noted in Section 2.3.4, this was shown not to be the case. That led to the development of biometric cryptosystems, where helper data are deemed not to leak biometric information. Research needs to be done in determining the biometric information that may be leaked to an adversary by the output of a deep neural network on a facial image, assuming that the adversary has access to the trained network.

References

Barclays. 2016. Barclays launches voice security technology to all customers. Press release. Barclays, London. http://www.newsroom.barclays.com/r/3383/barclays_launches_voice_security_technology_to_all_customers (accessed April 29, 2017).

J. Bort. 2014. The guy who just hacked Touch ID in the iPhone 6 says it's safe . . . for now. *Business Insider*. http://www.businessinsider.com/iphone-6-touchid-is-safe-for-now-2014-9 (accessed April 26, 2017).

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

L. Brinded. 2016. HSBC is letting customers verify their bank accounts like Airbnb does with selfies. *Business Insider*. http://www.businessinsider.com/hsbc-implements-selfie-security-verification-technology-for-banking-2016-9 (accessed April 28, 2017).

R. Cappelli, A. Lumini, D. Maio, and D. Maltoni. 2007. Fingerprint image reconstruction from standard templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(9):1489–1503.

Chaos Computer Club. 2013. Chaos Computer Club breaks Apple Touch ID. Chaos Computer Club, Germany. https://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid (accessed April 10, 2017).

Y. Cheung. 2017. HKBU scholar invents world's first "lip password." Hong Kong Baptist University News, Hongkong. http://hkbuenews.hkbu.edu.hk/?t=enews_details/1758&acm=50_726 (accessed April 20, 2017).

Citigroup. 2017. Citi tops 1 million mark for voice biometrics authentication for Asia Pacific consumer banking clients. Citigroup, New York. http://www.citigroup.com/citi/news/2017/170321b.htm (accessed April 29, 2017).

T. C. Clancy, N. Kiyavash, and D. J. Lin. 2003. Secure smartcard-based fingerprint authentication. In *Proceedings of the ACM SIGMM Multimedia, Biometrics Methods and Applications Workshop*, 45–52.

A. Czajka. 2016. Iris liveness detection by modeling dynamic pupil features. In *Handbook of Iris Recognition*, K. W. Bowyer and M. J. Burge (eds), 439–467. Springer, London.

N. Dalal and B. Triggs. 2005. Histograms of oriented gradients for human detection. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 886–893. IEEE Computer Society, Washington, DC.

J. Daugman. 1994. Biometric personal identification system based on iris analysis. US Patent 5,291,560. US Patent and Trademark Office, Alexandria, VA.

J. Daugman. 2003. The importance of being random: Statistical principles of iris recognition. *Pattern Recognition*, 36:279–291.

J. Daugman. 2006. Probing the Uniqueness and Randomness of IrisCodes: Results From 200 Billion Iris Pair Comparisons. *Proceedings of the IEEE*, 94(11):1927–1935.

R. Derakhshani. 2016. Biometric template security and key generation. US Patent 9,495,588. US Patent and Trademark Office, Alexandria, VA.

R. S. Dulude and C. Musgrave. 2001. Biometric certificates. US Patent 6,310,966. US Patent and Trademark Office, Alexandria, VA.

FIDO (Fast IDentity Online) Alliance. 2016. Fast IDentity Online (FIDO) Universal Authentication Framework (UAF). FIDO Alliance, Wakefield, MA. https://fidoalliance.org/download/ (accessed April 25, 2017).

R. Gamlin. 2016. Alabama's new weapon in the war against tax refund fraud? The selfie. *Birmingham Business Journal*, Birmingham, AL. http://www.bizjournals.com/birmingham/news/2016/05/27/alabamas-new-weapon-in-the-war-against-tax-refund.html (accessed April 25, 2017).

I. Goodfellow, Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press, Cambridge, MA.

F. Hao, R. Anderson, and J. Daugman. 2006. Combining crypto with biometrics effectively. *IEEE Transactions on Computers*, 55(9):1081–1088.

International Civil Aviation Organization. 2015. Doc 9303—Machine readable travel documents, seventh edition. International Civil Aviation Organization, Montreal, QB http://www.icao.int/publications/pages/publication.aspx?docnum=9303 (accessed April 25, 2017).

ISO (International Organization for Standardization)/IEC (International Electrotechnical Commission). 2011. ISO/IEC 24745: Information technology—Security techniques—Biometric information protection. ISO, Geneva; IEC, Geneva.

A. Juels and M. Sudan. 2006. A fuzzy vault scheme. *Designs, Codes and Cryptography*, 38:237–257.

A. Juels and M. Wattenberg. 1999. A fuzzy commitment scheme. In *Proceedings of the 6th ACM Conference on Computer and Communications Security*, 28–36. ACM, New York.

S. Khandelwal. 2014. Hacker clones German defense minister's fingerprint using just her photos. *The Hacker News*. http://thehackernews.com/2014/12/hacker-clone-fingerprint-scanner.html (accessed April 18, 2017).

K. Kollreider, H. Fronthaler, M. I. Faraj, and J. Bigun. 2007. Real-time face detection and motion analysis with application in "liveness" assessment, *IEEE Transactions on Information Forensics and Security*, 2(3–2):548–558.

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

K. Lewison and F. Corella. 2016. Rich credentials for remote identity proofing. Pomcor, Carmichael, CA. https://pomcor.com/techreports/RichCredentials.pdf (accessed April 25, 2017).

C. Li and J. Hu. 2016. A security-enhanced alignment-free fuzzy vault-based fingerprint cryptosystem using pair-polar minutiae structures. *IEEE Transactions on Information Forensics and Security*, 11(3):543–555.

C. Lu and K. Tang. 2015. Surpassing human-level face verification performance on LFW with GaussianFace. *AAAI Conference on Artificial Intelligence, North America*. https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9845 (accessed April 23, 2017).

Mastercard. 2017. Thumbs up: Mastercard unveils next generation biometric card. Johannesburg and Purchase, New York. https://newsroom.mastercard.com/press-releases/thumbs-up-mastercard-unveils-next-generation-biometric-card/ (accessed April 24, 2017).

S. Mayhew. 2016. Princeton Identity to license its patented iris recognition technology to Samsung. Biometrics Research Group, Inc, Toronto. http://www.biometricupdate.com/201609/princeton-identity-to-license-its-patented-iris-recognition-technology-to-samsung (accessed April 2017).

F. Monrose, M. K. Reiter, and S. Wetzel. 1999. Password hardening based on keystroke dynamics. In *Proceedings of the 6th ACM Conference on Computer and Communications Security*, 73–82. ACM, New York.

S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard. 2017. Universal adversarial perturbations. In *IEEE Conference on Computer Vision and Pattern Recognition*, 86–94. Available at https://arxiv.org/pdf/1610.08401.pdf.

D. Mukhopadhyay, M. Shirvanian, and N. Saxena. 2015. All your voices are belong to us: Stealing voices to fool humans and machines. In *Computer Security—ESORICS 2015*, G. Pernul, Y. A. P. Ryan, and E. Weippl (eds). Springer International Publishing, Cham.

NIST (National Institute of Standards and Technology). 2013. Federal Information Processing Standard (FIPS) 201-2. NIST, Gaithersburg, MD. http://csrc.nist.gov/publications/PubsFIPS.html (accessed April 25, 2017).

NIST. 2016. Special Publication (SP) 800-90A-C. 2016. NIST, Gaithersburg, MD. http://csrc.nist.gov/publications/PubsSPs.html#SP%20800 (accessed April 24, 2017).

S. Novoselov, T. Pekhovsky, A. Shulipa, and A. Sholokhov. 2014. Text-dependent GMM-JFA system for password based speaker verification. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 729–737.

OCBC (Oversea-Chinese Banking Corporation). 2016. OCBC Bank is first in Singapore to use voice recognition technology to enhance customer experience. Press release. OCBC, Singapore https://www.ocbc.com/assets/pdf/media/2016/may/ocbc%20media%20release%20-%20ocbc%20rolls%20out%20voice%20recognition.pdf (accessed April 29 2017).

T. Ojala, M. Pietikäinen, and D. Harwood. 1996. A comparative study of texture measures with classification based on feature distributions. *Pattern Recognition*, 29:51–59.

I. Pavlidis and P. Symosek. 2000. The imaging issue in an automatic face/disguise detection system. In *Proceedings of the IEEE Workshop on Computer Vision Beyond the Visible Spectrum: Methods and Applications*, 15–24. IEEE Computer Society, Washington, DC.

Pindrop. 2017a. Phoneprinting technology explained—How the largest global contact centers stop fraud and protect customers. Pindrop. https://www.pindrop.com/phoneprinting-webinar/ (accessed April 29, 2017).

Pindrop. 2017b. Call center fraud report. Pindrop. https://www.pindrop.com/wp-content/uploads/2017/04/Fraud-Report-Global-4-24-17-FINAL.pdf (accessed April 29, 2017).

R. Raghavendra and C. Busch. 2015. Robust scheme for iris presentation attack detection using multiscale binarized statistical image features. *IEEE Transactions on Information Forensics and Security* 10(4):703–715.

C. Rathgeb and A. Uhl. 2011. A survey on biometric cryptosystems and cancelable biometrics. *EURASIP Journal on Information Security* 2011:3.

D. A. Reynolds, T. F. Quatieri, and R. B. Dunn. 2000. Speaker verification using adapted Gaussian mixture models. *Digital Signal Processing*, 10(1):19–41.

M. Rogers. 2013. Why I hacked Apple's TouchID, and still think it is awesome. https://blog.lookout.com/blog/2013/09/23/why-i-hacked-apples-touchid-and-still-think-it-is-awesome/ (accessed April 26, 2017).

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2

F. Richardson, D. Reynolds, and N. Dehak. 2015. Deep neural network approaches to speaker and language recognition. *IEEE Signal Processing Letters*, 22(10):1671–1675.

A. Ross, J. Shah, and A. K. Jain. 2007. From template to image: Reconstructing fingerprints from minutiae points. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(4):544–560.

A. Roy, R. Memon, and A. Ross. 2017. MasterPrint: Exploring the vulnerability of partial fingerprint-based authentication systems. *IEEE Transactions on Information Forensics and Security*, 12(9): 2013–2025.

K. V. Sakariya and T. Nauta. 2017. Interactive display panel with IR diodes. US Patent 9,570,002. US Patent and Trademark Office, Alexandria, VA.

F. Schroff, D. Kalenichenko, and J. Philbin. 2015. FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.

S. Schuckers and P. Johnson. 2014. Fingerprint pore analysis for liveness detection. US patent application 14/243,420. US Patent and Trademark Office, Alexandria, VA.

M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter. 2016. Accessorize to a Crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 1528–1540. ACM New York.

D. Storm. 2014. Researchers spoof fingerprint, bypass Samsung Galaxy S5 security, access PayPal. Computerworld, IDG Communications Inc, Boston. http://www.computerworld.com/article/2476130/cybercrime-hacking/researchers-spoof-fingerprint—bypass-samsung-galaxy-s5-security—access-paypal.html (accessed April 26, 2017).

Y. Taigman, M. Yang, M. A. Ranzato, and L. Wolf. 2014. DeepFace: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1701–1708. IEEE Computer Society, Washington, DC.

U. Uludag and A. K. Jain. 2006. Securing fingerprint template: Fuzzy vault with helper data. In *Proceedings of the IEEE Workshop on Privacy Research in Vision*, 163–163. Full paper available at http://biometrics.cse.msu.edu/Publications/SecureBiometrics/UludagJain_FFVHelper_PRIV06.pdf.

Veridium. 2017. https://info.veridiumid.com/hubfs/Content/Datasheets/datasheet-4-fingers-touchless-id.pdf (accessed April 29, 2017).

University of Massachusetts. 2017. Labeled Faces in the Wild. http://vis-www.cs.umass.edu/lfw/ (accessed April 28, 2017).

P. Viola and M. Jones. 2004. Robust real-time object detection. *International Journal of Computer Vision*, 57(2):137–154.

Y. Wencheng, J. Hu, S. Wang, and M. Stojmenovic. 2014. An alignment-free fingerprint bio-cryptosystem based on modified Voronoi neighbor structures. *Pattern Recognition*, 47(3):1309–1320.

L. Wiskott, J.-M. Fellous, N. Krüger, and C. von der Malsburg. 1997. Face Recognition by Elastic Bunch Graph Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):775–779.

Y. Xu, T. Price, J.-M. Frahm, and F. Monrose. 2016. Virtual U: Defeating face liveness Detection by building virtual models from your public photos. In *Proceedings of the 25th USENIX Security Symposium*, 497–512.

W. S. Yambor, B. A. Draper and J. R. Beveridge. 2002. Analyzing PCA-based face recognition algorithm: Eigenvector' selection and distance measures. In *Empirical Evaluation Methods in Computer Vision*, ed. H. I. Christensen, P. J. Phillips, 39–60. World Scientific.

M. D. Zeiler and R. Fergus. 2013. Visualizing and understanding convolutional networks. In *Proceedings of the 13th European Conference on Computer Vision*, 818–833. Springer International Publishing, Cham.

Downloaded By: 10.2.97.136 At: 17:53 10 Jun 2023; For: 9781315184319, chapter2, 10.1201/b22142-2