Pattern Recognition Letters
journal homepage: www.elsevier.com

50 Years of Biometric Research: Accomplishments, Challenges, and Opportunities

Anil K. Jain[a,**], Karthik Nandakumar[b], Arun Ross[a]

[a]Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, USA
[b]IBM Research Collaboratory - Singapore, 9 Changi Business Park Central 1, Singapore 486048

ABSTRACT

Biometric recognition refers to the automated recognition of individuals based on their biological and behavioral characteristics such as fingerprint, face, iris, and voice. The first scientific paper on *automated* fingerprint matching was published by Mitchell Trauring in the journal Nature in 1963. The first objective of this paper is to document the significant progress that has been achieved in the field of biometric recognition in the past 50 years since Trauring's landmark paper. This progress has enabled current state-of-the-art biometric systems to accurately recognize individuals based on biometric trait(s) acquired under controlled environmental conditions from cooperative users. Despite this progress, a number of challenging issues continue to inhibit the full potential of biometrics to automatically recognize humans. The second objective of this paper is to enlist such challenges, analyze the solutions proposed to overcome them, and highlight the research opportunities in this field. One of the foremost challenges is the design of robust algorithms for representing and matching biometric samples obtained from uncooperative subjects under unconstrained environmental conditions (e.g., recognizing faces in a crowd). In addition, fundamental questions such as the distinctiveness and persistence of biometric traits need greater attention. Problems related to the security of biometric data and robustness of the biometric system against spoofing and obfuscation attacks, also remain unsolved. Finally, larger system-level issues like usability, user privacy concerns, integration with the end application, and return on investment have not been adequately addressed. Unlocking the full potential of biometrics through inter-disciplinary research in the above areas will not only lead to widespread adoption of this promising technology, but will also result in wider user acceptance and societal impact.

1. Introduction

> "It is the purpose of this article to present, together with some evidence of its feasibility, a method by which decentralized automatic identity verification, such as might be desired for credit, banking or security purposes, can be accomplished through automatic comparison of the minutiae in finger-ridge patterns."
> *– Mitchell Trauring, Nature, March 1963*

In modern society, the ability to reliably identify individuals in real-time is a fundamental requirement in many applications including forensics, international border crossing, financial transactions, and computer security. Traditionally, an exclusive possession of a token, such as a passport or an ID card, has been extensively used for identifying individuals. In the context of computer systems and applications, knowledge-based schemes based on passwords and PINs are commonly used for person authentication[1]. Since both token-based and knowledge-based mechanisms have their own strengths and limitations, the use of two-factor authentication schemes that combine both these authentication mechanisms are also popular.

Biometric recognition, or simply biometrics, refers to the automated recognition of individuals based on their biological and behavioral characteristics (Jain et al., 2011). Examples of biometric traits that have been successfully used in practical applications include face, fingerprint, palmprint, iris, palm/finger

[**] Corresponding author: Tel.: +1-517-355-9282; fax: +1-517-432-1061; *e-mail:* jain@cse.msu.edu (Anil K. Jain)

[1] Authentication involves verifying the claimed identity of a person.

vein, and voice. The use of DNA, in the context of biometrics (as opposed to just forensics), is also beginning to gain traction. Since biometric traits are generally inherent to an individual, there is a strong and reasonably permanent link between a person and his/her biometric traits. Thus, biometric recognition can be used to identify individuals in surveillance operations where covert recognition[2] is required or in scenarios where a person may attempt to conceal their true identity (e.g., by using forged documents to claim social welfare benefits). Consequently, the application domain of biometrics far exceeds that of passwords and tokens. In applications such as border control, forensics, surveillance, de-duplication[3] and chain-of-custody[4], the use of biometric solutions has clear-cut advantages over passwords or tokens.

However, the emergence of biometrics does not necessarily supplant the use of passwords or tokens in authentication applications. While biometrics can mitigate some of the limitations associated with the use of passwords, biometric systems themselves are vulnerable to spoof attacks, linkability attacks (linking users across applications based on their biometric data), and can incur additional hardware and software costs. Further, the acquisition process introduces variations in the biometric data of an individual (referred to as intra-subject variations) that may lead to false non-matches and false matches. False matches can lead to identity creep, where an adversary, after repeated attempts, manages to take on the identity of a legitimate user of the system. The lack of secrecy (e.g., face images on social media sites) and distinctiveness (e.g., face images of identical twins) of biometric traits pose additional problems to biometric-based authentication schemes. Given the above limitations, a multi-factor authentication mechanism that judiciously combines biometrics with passwords and/or tokens may be a better approach to security in many applications (O'Gorman, 2003).

1.1. Motivation and Objectives

The first known research publication on automated biometric recognition was the one published by Mitchell Trauring in the journal Nature in 1963 on fingerprint matching (Trauring, 1963). The development of automated biometric systems based on other traits such as voice (Pruzansky, 1963), face (Bledsoe, 1966), and signature (Mauceri, 1965) also started in the 1960s. Subsequently, biometrics systems based on traits like hand geometry (Ernst, 1971) and iris (Daugman, 1993) were developed. In this sense, 50 years have passed since the first paper on automated biometric recognition was published. Coincidentally, modern computer vision[5] also had its beginnings approximately 50 years ago, with Roberts' PhD dissertation on machine vision of 3D solids (Roberts, 1963).

In a 2007 article, Wayman (Wayman, 2007) tracked the major developments in biometrics in the United States from the 1960's to the 1990's, and observed the following: *"A quick overview of biometric history shows that much of what we consider to be "new" in biometrics was really considered decades ago. There is much left to be done, but the most efficient route will be to consider that which is really yet undiscovered, not wasting time repeating the studies of years ago. Even in 2005, it is much too early to speculate on what the first decade of the new millennium will ultimately hold for biometrics. It seems clear, however, that the industry will continue to grow and that technical and human improvements to the systems will be made."*

In line with the above observation from Wayman, the objective of this paper is to summarize the progress in biometric recognition so as to understand *how this field emerged*, *where we are now*, and *where we should go from here*. We believe that this assessment of biometrics research would shed light on the cross-disciplinary nature of problems in biometric recognition, highlight the tremendous opportunities for both basic and applied research in biometrics, and motivate budding scientists and engineers to consider biometric recognition as their field of study.

2. Biometric Recognition Framework

A typical biometric recognition system has two stages of operation, namely, the enrollment stage and the recognition stage (see Figure 1). In the enrollment stage, the biometric system acquires the biometric trait of an individual, extracts a salient feature set from it and stores the extracted feature set in a database (often referred to as a template), along with an identifier associating the feature set with an individual. During the recognition stage, the system once again acquires the biometric trait of an individual, extracts a feature set from it, and compares this feature set against the templates in the database in order to determine a match or to verify a claimed identity.

In the enrollment stage, a biometric sensor scans the biometric trait ($\mathbf{B}$) of a user (Y) to obtain a digital representation ($\mathbf{M}$). Since the scanned biometric trait may be affected by various sources of noise ($\boldsymbol{\eta}$) during the sensing process, a quality check is generally performed to ensure that the acquired biometric data can be reliably processed by successive modules. In order to facilitate recognition, the raw biometric data is further processed by a feature extractor (denoted by the function f_e) to generate a compact but expressive representation, called a feature set, which is stored as a template ($\mathbf{X}_E$) in the system database ($\mathbf{D}$) for future comparison. During the recognition stage, when the user needs to be authenticated or identified, a new sample of the biometric trait is obtained. Features ($\mathbf{X}_R$) are extracted from this query biometric sample and compared (denoted by the function f_m) to the templates stored in the database in order to determine the identity ($\hat{Y}$) associated with the query (sometimes referred to as the probe) biometric sample.

If the objective is to verify the claimed identity of an individual, the query biometric sample needs to be compared only to the template corresponding to the claimed identity (one-to-one match). The identity claim is accepted if the resulting similarity

[2] In a covert scenario, the subject's biometric traits are acquired without the subject's explicit knowledge and surreptitiously used for recognition purposes

[3] De-duplication involves the removal of duplicate "identities", where, for example, a single individual may have multiple passports under different names

[4] This is to keep track of individuals who handle the physical evidence collected during the course of a legal proceeding

[5] J. Malik, "The Three Rs of Vision," http://www.di.ens.fr/willow/events/cvml2013/materials/slides/wednesday/Malik-paris-CVML-2013.pdf

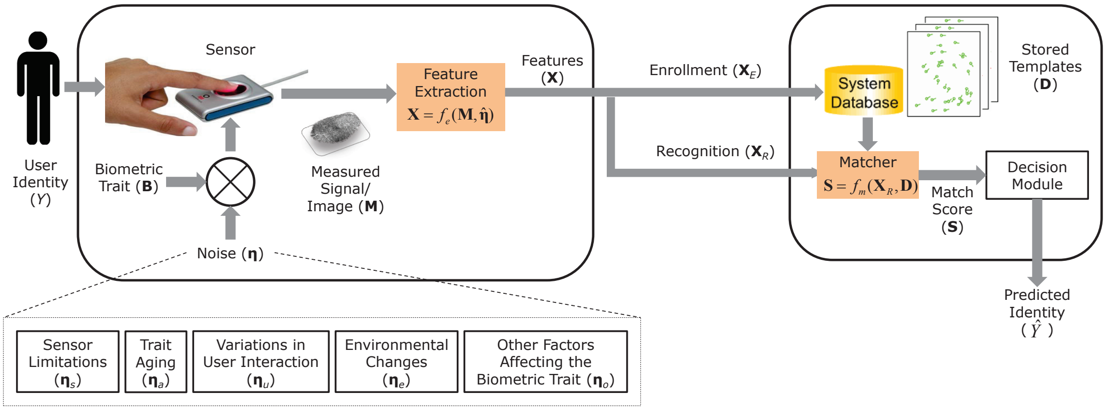
Sensor
User Identity (Y)
Biometric Trait ($\mathbf{B}$)
Noise ($\boldsymbol{\eta}$)
Measured Signal/ Image ($\mathbf{M}$)
Feature Extraction
$\mathbf{X} = f_e(\mathbf{M}, \hat{\boldsymbol{\eta}})$
Features ($\mathbf{X}$)
Enrollment ($\mathbf{X}_E$)
Recognition ($\mathbf{X}_R$)
System Database
Stored Templates ($\mathbf{D}$)
Matcher
$\mathbf{S} = f_m(\mathbf{X}_R, \mathbf{D})$
Match Score ($\mathbf{S}$)
Decision Module
Predicted Identity ($\hat{Y}$)
Sensor Limitations ($\boldsymbol{\eta}_s$)
Trait Aging ($\boldsymbol{\eta}_a$)
Variations in User Interaction ($\boldsymbol{\eta}_u$)
Environmental Changes ($\boldsymbol{\eta}_e$)
Other Factors Affecting the Biometric Trait ($\boldsymbol{\eta}_o$)

Fig. 1. Operation of a typical biometric system. The two fundamental problems in biometric recognition involve finding an invariant feature representation and designing a robust matcher for a given representation scheme.

value (also called the match score **S**) is above a preset threshold. In this case, the biometric system is said to operate in the *verification* or *authentication* mode. If the goal is to determine the user's identity without the user having to claim an identity, the query needs to be compared against all the templates in the database (one-to-many match). This latter functionality is commonly referred to as *identification* and the result of an identification operation will be one of the following two decisions: (i) output the identity of one or more users whose templates exhibit high similarity with the query biometric sample or (ii) a response indicating that the query biometric sample does not match with the templates of any of the users already enrolled in the database. If the biometric identification system is forced to output an identity, it is referred to as *closed-set* identification. On the other hand, the option of having a reject response results in an *open-set* identification system.

As shown in Figure 1, the measured biometric signal **M** may exhibit intra-subject variations, i.e., the signal does not remain stable across measurements. These sources of ira-subject variations can be broadly divided into five categories: (i) sensor limitations ($\boldsymbol{\eta}_s$), (ii) intrinsic aging of the biometric trait ($\boldsymbol{\eta}_a$), (iii) variations in user interaction ($\boldsymbol{\eta}_u$), (iv) changes in the acquisition environment ($\boldsymbol{\eta}_e$), and (iv) all other factors affecting the biometric trait ($\boldsymbol{\eta}_o$). As an illustrative example, let us consider a face recognition system where the face is captured using a 2-dimensional (2D) camera operating in the visible spectrum. In this context, sensor limitations ($\boldsymbol{\eta}_s$) may include low spatial resolution and frame rate of the camera, inability to capture the full 3D structure of the human face, and inability to capture the details of the face under low illumination conditions. Changes in a person's facial structure and appearance ($\boldsymbol{\eta}_a$) can occur over time due to the effects of biological aging. Variations in the person's facial pose and expression ($\boldsymbol{\eta}_u$) can be introduced when the user interacts with the sensor. This type of variation is more pronounced in biometric applications where cooperation from the users cannot be expected (e.g., covert surveillance). Illumination changes ($\boldsymbol{\eta}_e$) in the acquisition environment will also affect the quality of the captured face images. Finally, other factors such as make-up and accessories (e.g., sunglass, hat, etc.) worn by the person and occlusion of a person's face by other objects or individuals ($\boldsymbol{\eta}_o$) will also the potential adversely affect the face image quality. Table 1 presents a summary of different sources of intra-subject variations encountered in biometric systems based on commonly used biometric traits such as fingerprint, face, iris, and voice.

2.1. How to Choose a Biometric Trait?

A critical issue in biometric system design is the choice of biometric trait. In theory, any anatomical, behavioral, or physiological characteristic of an individual can be used as a biometric trait. However, the choice of a biometric trait for a particular application usually depends on the degree to which the following properties are satisfied: (i) uniqueness or distinctiveness, (ii) permanence, (iii) universality, (iv) collectability, (v) performance, (vi) user acceptance, (vii) invulnerability, and (viii) integration (Jain et al., 2011). A biometric trait is said to be unique to an individual only if every pair of individuals in the target population can be differentiated based on this trait. Since uniqueness is difficult to guarantee, the term *distinctiveness* is often used. Ideally, a biometric trait or its representation (extracted features) should be permanent and should retain its discriminatory power over the lifetime of an individual. Since the distinctiveness and permanence of a biometric trait constitute the fundamental premise of biometric recognition, they play a major role in determining the value of biometric trait.

While a number of biometric traits have been proposed for person recognition (see Figure 2), fingerprint, face, and iris are the three most popular biometric traits in deployed systems. One of the reasons for the popularity of fingerprint and face is the availability of large legacy databases (e.g., driver license and immigration databases), which have been collected by law enforcement and other government agencies all over the world. While iris is being increasingly adopted for large-scale identification (e.g., the iris recognition border crossing system in the United Arab Emirates) due to its high accuracy in applications requiring de-duplication, there are relatively fewer legacy iris databases. Another major reason for the adoption of fingerprint, face, and iris modalities is the periodic technology evaluations

Table 1. Summary of various sources of intra-subject variations in the measured biometric signals for different biometric traits.

Source of intra-subject variations	**Fingerprint**	**Face**	**Iris**	**Voice**
Sensor limitations ($\boldsymbol{\eta}_s$)	Resolution (dots per inch), signal to noise ratio, sensor cleanliness	Spatial resolution, frame rate, acquisition spectrum (visible vs. infrared), distance from camera, 2D vs 3D	Acquisition spectrum (visible vs. near infra-red), distance from sensor	Signal to noise ratio
Intrinsic aging ($\boldsymbol{\eta}_a$)	Variations in ridge thickness & height due to changes in skin elasticity & sebaceous gland activity	Geometric changes during childhood & adolescence, wrinkles & saggy skin in old age	Myotic pupil (pupil constricts)	Voice changes during childhood & adolescence, pitch changes, voice shakiness in old age
Variations in user interactions ($\boldsymbol{\eta}_u$)	Rotation, translation, finger pressure	Pose, expression	Pupil dilation, partially closed eyes (blinking), gaze angle	Speed, intensity, accent variations
Environment changes ($\boldsymbol{\eta}_e$)	Indoor vs. outdoor	Illumination, background scene	Illumination	Background noise
Other factors ($\boldsymbol{\eta}_o$)	Cuts, worn-out fingers, dry/wet fingers	Make-up, accessories, occlusion	Eye diseases, influence of alcohol	Common cold

for these traits (along with voice) conducted by the National Institute of Standards and Technology (NIST). These large scale evaluations on operational biometric data have been responsible for documenting the significant progress that has been made in the matching accuracy of these biometric traits.

Leaving aside face, fingerprint and iris, the other biometric traits that have been either deployed in biometric systems or proposed in the research literature can be grouped under three broad categories:

1. Traits such as palmprint and deoxyribonucleic acid (DNA) are beginning to play a major role in law enforcement and forensic applications, mainly because of their value in large-scale identification. While DNA has a relatively short history (first put into operational use in 1986 (New England Innocence Project, 2011)) and a relatively small database (current size of the National DNA Index in the United States is about 12 million (The Federal Bureau of Investigation, 2013a)), often this may be the only reliable forensic evidence available at crime scenes. The role of DNA in exonerating wrongly convicted and incarcerated individuals through the efforts of the Innocence Project (Innocence Project, 2013) is well known. While friction ridge pattern on the surface of human palm, similar to a finger ridge pattern, is claimed to be unique (Ashbaugh, 1999), fingerprints are easier to capture (due to their relatively small size compared to palmprints) and provide acceptable solutions for person recognition. This explains why fingerprints are more popular than palmprints in biometric systems. However, given the fact that many friction ridge impressions left at crime scenes are those of palms, law enforcement agencies have started to collect palmprints of suspects at the time of booking. This is the rationale behind the decision by the Federal Bureau of Investigation (FBI) to include palmprint modality in the Next Generation Identification (NGI) system.[6]
2. Biometric traits such as voice, signature, hand geometry, and vascular patterns (palm vein, hand vein, or finger vein) have been deployed in commercial applications, mostly as a tool for verification or authentication, but their use so far is rather limited.
3. Traits like gait, ear, retina/sclera, keystroke dynamics, electrocardiogram (ECG), and electroencephalogram (EEG) signals have been proposed by researchers for person recognition in niche applications, but are yet to attain sufficient level of technological maturity and acceptance.

Biometric traits discussed above have varying degrees of distinctiveness and permanence for person recognition in a target population. Since the basic objective of a biometric system is to correctly establish whether the two given samples of a biometric trait belong to the same user, the recognition performance or matching accuracy is often used as the primary criterion for selecting a biometric trait. However, it is important to realize that accuracy is not the only factor that determines the utility of a biometric trait or the biometric system itself in a particu-

[6] `http://www.fbi.gov/about-us/cjis/fingerprints_biometrics/ngi`

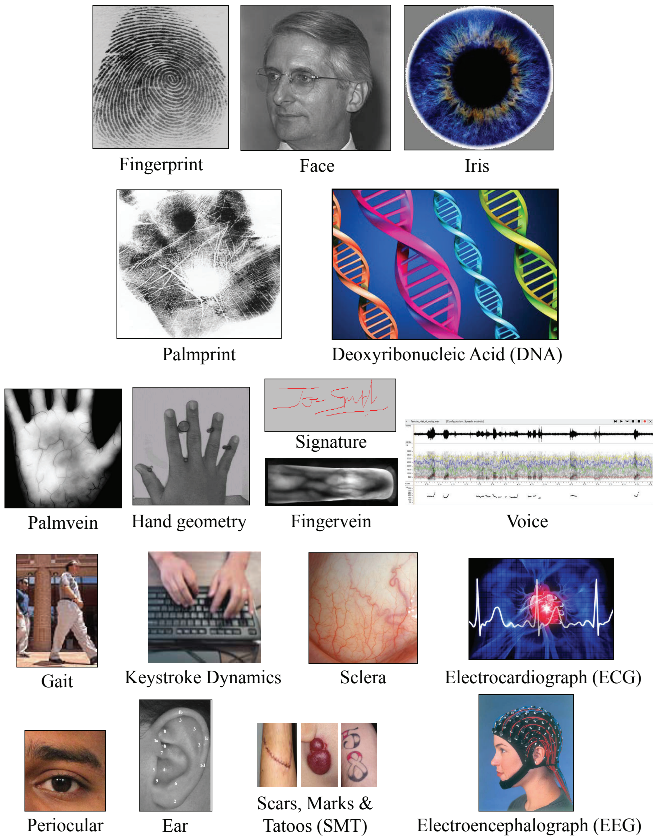

Fig. 2. A large number of body traits have been proposed and used for person recognition. Fingerprint, face, and iris modalities shown in the first row are the three most popular biometric traits in deployed systems. Traits such as palmprint and DNA (depicted in the second row) have legacy databases and are currently being used primarily in law enforcement and forensics. The third row shows traits that have been deployed in commercial applications, primarily for verification operation (one-to-one matching). Finally, the last two rows show traits like gait, ear, sclera, keystroke dynamics, ECG, and EEG signals, which have been proposed by researchers for person recognition in niche applications, but are yet to attain sufficient level of technological maturity for deployment.

lar application. Often, other practical issues such as throughput, cost, return on investment (ROI), user experience, template size, resistance to spoof and template attacks, and ease of system integration must also be given due consideration during the selection of a biometric trait.

Due to the diverse nature of biometric applications (e.g., mobile phone unlocking to international border crossing), no single biometric trait is likely to be *optimal* and satisfy the requirements of all applications. In many cases, a combination or fusion of multiple biometric traits may be required to attain the desired level of performance; such systems are generally referred to as *multibiometric systems* (Ross et al., 2006). One such example is the Aadhaar system in India, where there is a need to distinguish between individuals in a database involving hundreds of millions of identities[7]. Therefore, the Aadhaar system uses all ten fingerprints and both irides of subjects for de-duplication of identities.

It is important to emphasize that the design of a biometric system generally involves a complex interplay of factors related to user interaction (with the biometric sensor), the end-application, and the biometric recognition technology. For example, consider a financial application like Internet banking, where the overarching objective of using a biometric system is to minimize the losses incurred due to fraudulent transactions without causing too much inconvenience to the genuine customers. In this scenario, the bank needs to decide whether a particular transaction should be authorized or declined. Hence, the level of authentication required will depend on the risk associated with a transaction. A simple authentication scheme (e.g., account number and PIN) may be sufficient for an account balance inquiry, while a much higher level of identity assurance (e.g., a strong biometric match) may be required to perform a high-value funds transfer. It is also possible to combine the biometric match score with other contextual information such as customer's past transaction history and current location of the customer to generate an overall risk score, which can form the basis for the authorization decision. Thus, designing a biometric system not only requires knowledge of biometric technology, but also a good understanding of application requirements and issues related to human factors, ergonomics, and environmental variables.

2.2. Core Research Challenges in Biometrics

The main objective of a biometric system is to recognize individuals accurately. This in turn implies that a biometric system must have low recognition error rates. While false match rate (FMR) and false non-match rate (FNMR) quantify the errors in a verification system, false positive identification rate (FPIR) and false negative identification rate (FNIR) are used as the error metrics in an identification system. The conditional entropy[8] $H(Y|\hat{Y})$, where Y and $\hat{Y}$ are the true and predicted identities, respectively, is a function of the recognition

[7] As on 15^{th} December 2014, more than 720 million Aadhaar numbers have been issued.

[8] Intuitively, $H(Y|X)$ measures the uncertainty in Y given X.

error rates of the biometric system. In the case of biometric verification, $H(Y|\hat{Y}) = (H_b(\text{FMR}) + H_b(\text{FNMR}))/2$, where $H_b(p) = -(p\log_2 p + (1-p)\log_2(1-p))$ is the binary entropy function. For closed set identification, $H(Y|\hat{Y}) = -((1 - \text{FNIR})\log_2(1 - \text{FNIR}) + \text{FPIR}\log_2(\text{FPIR}/(N-1)))$. Since every stage of processing in a biometric system from the sensor to the matcher typically leads to loss of some discriminatory information, the following relationship is usually true: $H(Y|\hat{Y}) \geq H(Y|\mathbf{S}) \geq H(Y|\mathbf{X}) \geq H(Y|\mathbf{M}) \geq H(Y|\mathbf{B})$.

The primary challenge in a biometric recognition system is to design a suitable *sensor*, *feature representation scheme*, and *similarity measure* to minimize the recognition errors (or $H(Y|\hat{Y})$). This can be achieved by suppressing the effect of various noise sources without degrading the inherent identity information contained in a biometric trait. In particular, the following two conditions must be satisfied: (i) the similarity between different samples of the same biometric trait acquired from the same subject (intra-subject similarity) should be very high, and (ii) the similarity between different samples of a biometric trait acquired from different individuals (inter-subject similarity) should be very low. While advancements in sensor design can certainly benefit a biometric system by minimizing $H(Y|\mathbf{M})$, such improvements heavily rely on scientific and technological breakthroughs in related fields (e.g., optics). Consequently, most of the research in biometric recognition has rightly focused on the following two fundamental problems:

1. *The challenge of identifying the best representation scheme for a given biometric trait* - The desired set of features should retain all the discriminative information that is distinctive to a person and remain invariant to intra-subject variations. In other words, the feature extractor f_e must be designed such that it minimizes $H(Y|\mathbf{X})$, which is the conditional entropy of Y given the feature representation $\mathbf{X}$. Note that $H(Y|\mathbf{X}) = H(Y) - H(\mathbf{X}) + H(\mathbf{X}|Y)$, where $H(\mathbf{X})$ is the entropy of the biometric template and $H(\mathbf{X}|Y)$ quantifies the intra-subject variations. Thus, minimizing $H(Y|\mathbf{X})$ requires maximization of biometric template entropy and simultaneously minimizing intra-subject variations. While it may be relatively easy to enhance the entropy of the biometric template by extracting more features from the sensed images, there is no guarantee that these additional features will lead to better accuracy, unless these features also exhibit small intra-subject variations.
2. *The challenge of designing a robust matcher for a given representation scheme* - The desired matching algorithm must model the variations in the features belonging to the same individual, while accounting for variations between features of different individuals. Thus, the matcher f_m must minimize $H(Y|\mathbf{X}_R, \mathbf{D})$, the conditional entropy of Y given the query features $\mathbf{X}_R$ as well as the templates in database $\mathbf{D}$.

It is important to point out that there is no representation scheme or matcher that can be applied universally to all biometric traits. In fact, the feature extraction and matching algorithms must be carefully selected after taking into account the characteristics of the underlying biometric trait, the properties of the biometric samples captured by the sensor, and the requirements of the application (error rate, processor and memory constraints, throughput, etc.).

Since the inherent distinctiveness and permanence of a biometric trait determine the recognition accuracy of a biometric system to a large extent, analysis of these two properties for different biometric traits is also considered a core research problem in biometrics.

Genetic similarity between related individuals (e.g., twins, father and son) may contribute to the lack of distinctiveness for some biometric traits (e.g., facial appearance as shown in Figure 3). The iris texture and, to some extent, local fingerprint details, are known to be generated through random morphogenesis (phenotype characteristics). For this reason, fingerprints and irides of identical twins have been empirically shown to satisfy the distinctiveness property (Jain et al., 2002). The "distinctiveness" of a biometric trait is a quantifiable measure of the distinctiveness of the trait based on the selected feature representation. It can be mathematically defined as the mutual information between the user identity Y and the feature representation $\mathbf{X}$ derived from the biometric trait (denoted as $I(Y;\mathbf{X})$). Since $I(Y;\mathbf{X}) = H(\mathbf{Y}) - H(\mathbf{Y}|X)$, it is clear that the distinctiveness of a biometric trait is a useful theoretical measure, which can indicate the best recognition accuracy achievable based on the selected features $\mathbf{X}$. A rigorous evaluation of distinctiveness for different biometric traits and features derived from them is still an open research problem.

The effects of body growth on common identifiers like face, fingerprint, or iris (and the representations derived from them) have not been systematically studied in the literature. The notion of permanence (also referred as persistence) can be studied by modeling the variations caused by aging as a form of time-varying noise $\boldsymbol{\eta}_a(t)$. Ideally, one would like to precisely understand the effect of $\boldsymbol{\eta}_a(t)$ on $H(\mathbf{Y}|X)$ and design a feature extractor f_e such that the effect of aging on $H(\mathbf{Y}|X)$ is minimal. However, this is challenging in practice because it is very difficult to isolate the effect of aging phenomenon from other types of noise affecting a biometric measurement (Lui et al., 2009; Beveridge et al., 2009; Klare et al., 2012).

3. Evolution of Biometric Recognition

One trigger for the systematic use of biometric traits to recognize a person was the enactment of the Habitual Criminals Act in 1869 in the United Kingdom (Spearman, 1999). This Act made it mandatory to maintain a register of all persons convicted of a crime in the United Kingdom along with appropriate evidences of identity. This register was used to identify repeat offenders, who were generally incarcerated with a higher degree of punishment compared to first-time offenders. The need for such an identification scheme was expressed by a Home Office Committee as follows,

> "What is wanted is a means of classifying the records of habitual criminal, such that as soon as the particulars of the personality of any prisoner (whether *description, measurements, marks, or photographs*) are received, it may be possible to ascertain readily, and with certainty,

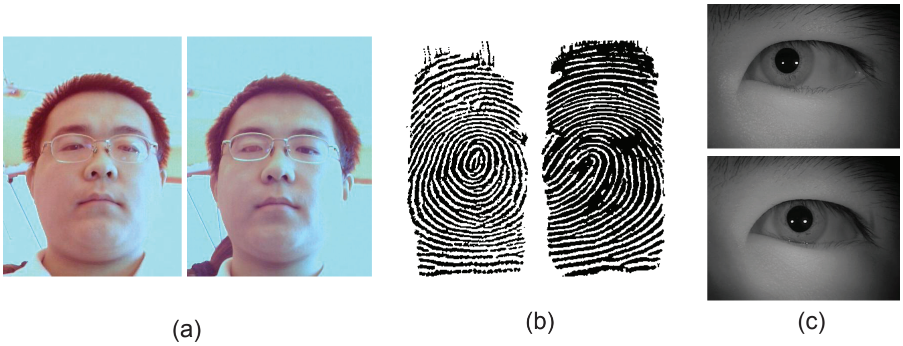

Fig. 3. Biometric samples obtained from identical twins. (a) Face, (b) fingerprint, and (c) iris images. While it is difficult to distinguish between these two users based on face images, it is indeed possible to distinguish between them using fingerprint or iris.

whether his case is in the register, and if so, who he is" (in page 257 of (Spearman, 1999), emphasis added).

In order to identify repeat offenders, a French police officer named Alphonse Bertillon introduced a system of person identification based on a set of anthropometric measurements (Bertillon, 1896). Additionally, he utilized multiple descriptive attributes such as eye color, scars and marks (referred to as soft biometrics in contemporary literature) in order to recognize an individual (see Figure 4). But the Bertillon system lacked automation, was cumbersome to administer uniformly (making it prone to error), and could not guarantee variations across individuals. Therefore, it was quickly abandoned in favor of a simpler and more accurate approach involving manual comparison of human fingerprints. This was made possible by the pioneering works of Henry Faulds, William Herschel, and Sir Francis Galton, who established the uniqueness of certain features in a fingerprint ridge pattern such as minutia points (Galton, 1892).

3.1. *Historical Developments in Fingerprint Recognition*

> "Perhaps the most beautiful and characteristic of all superficial marks (on human body) are the small furrows with the intervening ridges and their pores that are disposed in a singularly complex yet even order on the under surfaces of the hands and feet."
> *–Sir Francis Galton, Nature, June 28, 1888*

Traditionally, fingerprint images have been broadly classified into three categories, namely, (i) rolled/full, (ii) plain/flat and (iii) latent (see Figure 5). Typically, rolled and plain fingerprint images obtained using live-scan fingerprint sensors are of good quality (especially if the user is cooperative). In contrast, latent fingerprints are lifted from surfaces of objects that are inadvertently touched or handled by a person through a variety of means ranging from simply photographing the print to more complex dusting or chemical processing. While forensic applications typically require latent-to-rolled print comparison, most of the other applications involve comparisons between plain/rolled prints (Maltoni et al., 2009).

Fingerprint features can generally be categorized into three levels as shown in Figure 6. Level 1 features capture macroscopic details of the fingerprint such as ridge flow, ridge frequency, pattern type, and singular points (e.g., core and delta). Level 2 features refer to minutiae, such as ridge bifurcations and endings. Level 3 features capture the dimensional attributes of the ridge and include extended features such as ridge path deviation, width, shape, pores, edge contour, incipient ridges, breaks, creases, scars, and other permanent details. Level 1 and Level 2 friction ridge details are the most commonly used features by all deployed fingerprint recognition systems. Generally, Level 1 features are first extracted, followed by Level 2 features with the guidance of Level 1 features.

Numerous solutions have been proposed in the literature to tackle the problem of matching features extracted from two fingerprint images to determine if they were acquired from the same finger (Maltoni et al., 2009). Most of these solutions adopt one of the following three approaches: image correlation, matching of ridge features, and minutiae matching. Minutiae-based matching is the most commonly used approach, primarily due to the following reasons: (i) minutiae have been used successfully for fingerprint comparison by forensic examiners over the past 100 years and (ii) minutiae-based representation is storage efficient.

Some of the major milestones in the history of fingerprint recognition are summarized in Figure 7. In 1891, Argentine police officials initiated the fingerprinting of criminals and used fingerprint as an evidence in a homicide case in 1892 (Hawthorne, 2009). This is believed to be the first use of fingerprints in criminal proceedings. In 1901, the Scotland Yard in the United Kingdom began using fingerprint in law enforcement applications[9]. Fingerprints were accepted as an evidence of identity in a British criminal case for the first time in 1905. In 1924, the United States Congress authorized the Department of Justice to collect fingerprints along with the arrest information. This paved the way for the establishment of a fingerprint identification system by the Federal Bureau of Investigation (FBI), which started collecting fingerprints using tenprint cards (see Figure 8).

The FBI initiated the implementation of automated fingerprint identification system (AFIS) in the 1970s. Though this system is referred to as AFIS, it must be emphasized that the automation was not fully completed in the initial years of deployment. Human experts were still required to process the fin-

[9] http://onin.com/fp/fphistory.html

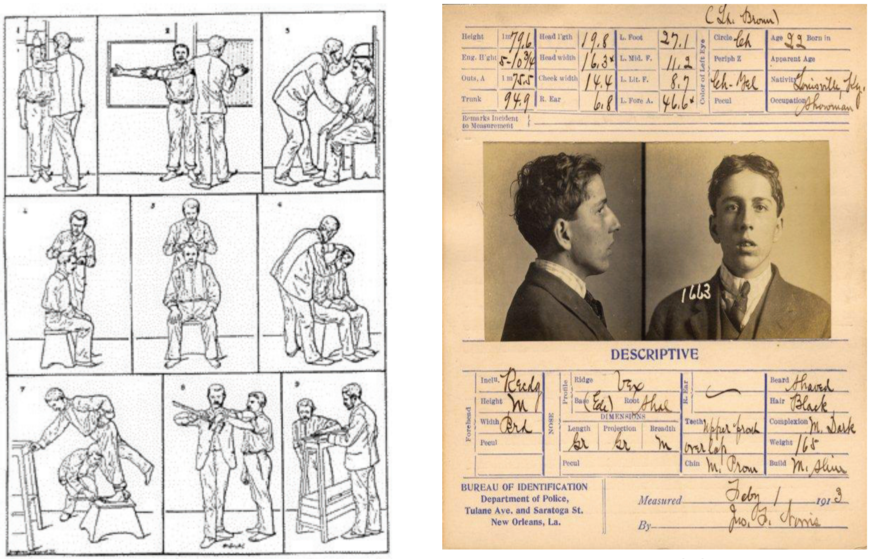

(a) (b)

Fig. 4. The Bertillon system, so named after its inventor Alphonse Bertillon (Bertillon, 1896), relied on the precise measurement of various attributes of the body for identifying recidivists. These measurements included, among others, the height of the individual, the length of the arm, geometry of the head, and the length of the foot. Some of the steps in the measurement process are depicted in (a) and the results were marked on a card as shown in (b).

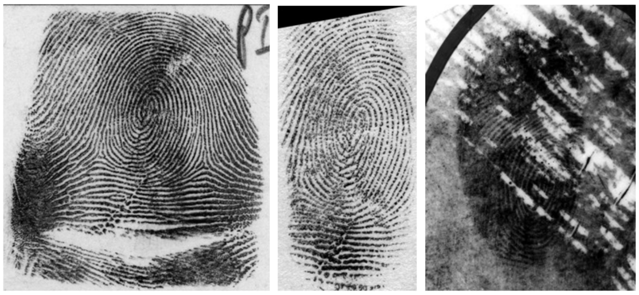

(a) (b) (c)

Fig. 5. Three different types of impressions of the same finger. (a) Rolled fingerprint, (b) plain fingerprint, and (c) latent fingerprint.

gerprint cards and identify the basic features such as minutia points, which were then matched automatically by the AFIS to retrieve a short-list of most similar matches from the database. The final match decision continued to remain in the hands of human experts. More recent large-scale deployments of fingerprint recognition systems such as the US-VISIT[10] program by the Department of Homeland Security (Department of Homeland Security, 2013), FBI's Next Generation Identification (NGI) program, and India's Aadhaar project (Planning Commission, Government of India, 2012) tend to be fully automated systems that use all ten fingers of the human hand as well as other modalities such as face, iris, and palmprint.

The growth in application areas for fingerprint recognition has coincided with the development of new sensors to capture the fingerprint (friction ridge) patterns (see Figure 9). In 1892, Juan Vucetich pioneered the use of inked fingerprint images, which are acquired by first applying ink to the subject's fingertip and then rolling or pressing the finger on paper, thereby creating an impression of the fingerprint ridges on paper. Later, the development of flatbed document scanners enabled the digitization of the inked fingerprints into images on a computer. Live-scan fingerprint sensors, which produce the digital image directly from a subject's fingertip via digital imaging technologies (e.g., optical, capacitive, and ultrasound) were developed in the 1990s (Xia and O'Gorman, 2003). Some of the recent advances in fingerprint sensing include the development of sensors that allow rapid ten-print capture (Department of Homeland Security, 2013), sensors that can record the 3-dimensional information of the ridge-valley patterns present on a fingertip (Parziale and Diaz-Santana, 2006), touchless fingerprint acquisition (e.g., Morpho's Finger-on-the-Fly) and imaging of fingerprints in multiple spectral bands (e.g., Lumidigm's Multi-Spectral Imaging (MSI) sensors).

With the advancements in the semiconductor industry, live-scan fingerprint scanners continue to become more compact and efficient, thereby enabling new applications in consumer electronic devices. For example, the Touch-ID fingerprint recognition system in iPhone-6 enables phone unlocking capability as well as mobile payments via the Apple Pay service. In the near future, it may be possible to capture face, fingerprint, iris, and voice biometric modalities using a commodity smartphone. The ability to securely authenticate a smartphone user using multibiometrics can be expected to open up a number of new applications involving mobile commerce and transactions.

[10] In March 2013, the United States Visitor and Immigration Status Indicator Technology (US-VISIT) was replaced by the Office of Biometric Identity Management (OBIM).

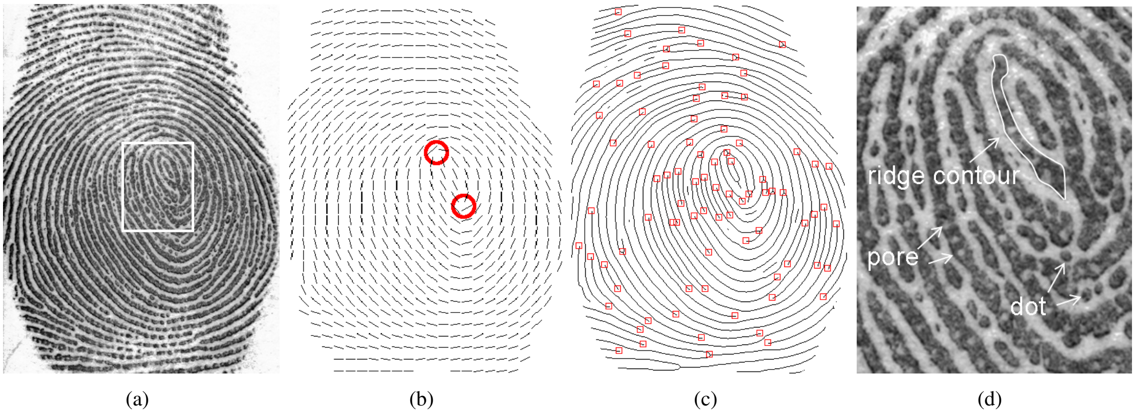

Fig. 6. Feature representation for fingerprint recognition. (a) A grayscale fingerprint image, (b) Level 1 features (orientation field or ridge flow and singular points), (c) Level 2 feature (ridge skeleton and minutiae), and (d) Level 3 features (ridge contour, pore, and dot).

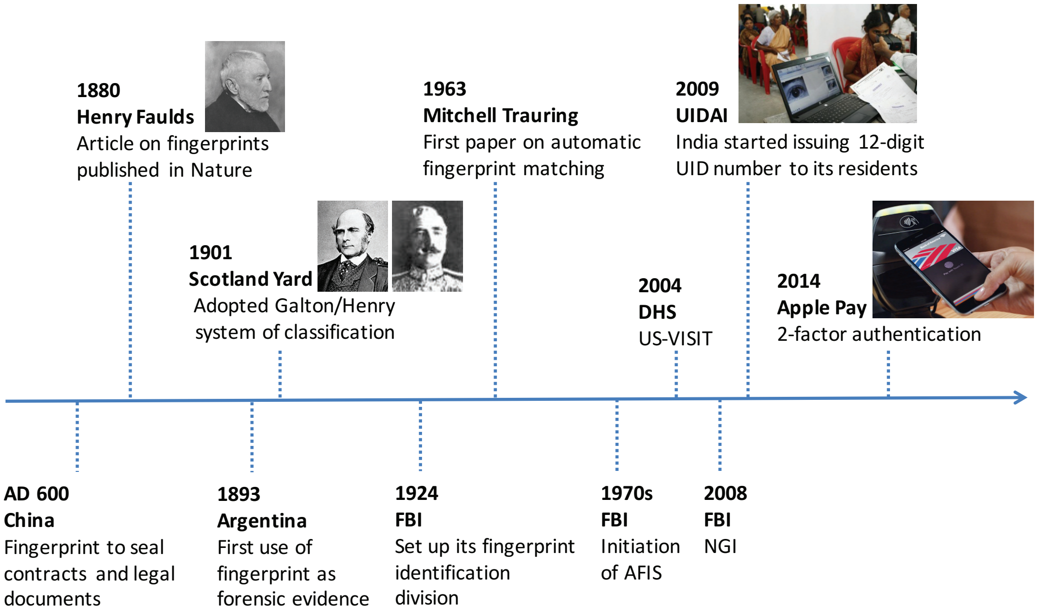
1880
Henry Faulds
Article on fingerprints published in Nature
1963
Mitchell Trauring
First paper on automatic fingerprint matching
2009
UIDAI
India started issuing 12-digit UID number to its residents
1901
Scotland Yard
Adopted Galton/Henry system of classification
2004
DHS
US-VISIT
2014
Apple Pay
2-factor authentication
AD 600
China
Fingerprint to seal contracts and legal documents
1893
Argentina
First use of fingerprint as forensic evidence
1924
FBI
Set up its fingerprint identification division
1970s
FBI
Initiation of AFIS
2008
FBI
NGI

Fig. 7. Some major milestones in the history of fingerprint recognition. Here, FBI stands for the Federal Bureau of Investigation, AFIS represents Automated Fingerprint Identification System, DHS indicates the Department of Homeland Security, US-VISIT stands for United States Visitor and Immigration Status Indicator Technology, NGI indicates FBI's Next-Generation Identification, and UIDAI represents Unique Identification Authority of India.

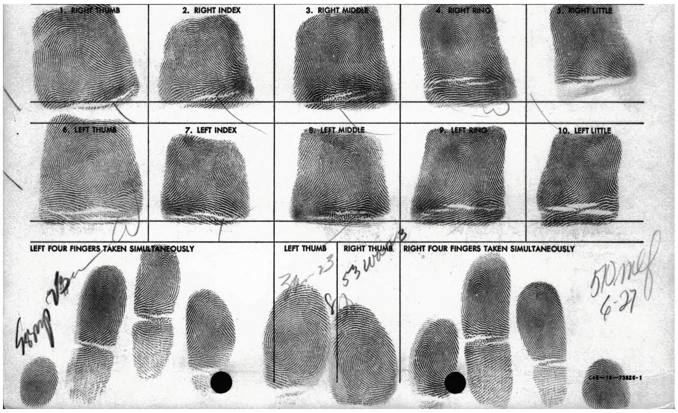

Fig. 8. A tenprint card used in law enforcement. The top two rows show fingerprints acquired by rolling each finger from one side to the other (so called *rolled* fingerprints). The bottom row shows *plain* or *slap* fingerprints: slap impressions of four fingers (little to index finger) of the left hand acquired simultaneously are shown on the left part of the bottom row, two thumb prints are shown in the middle, and the slap impressions of four fingers (index to little finger) of the right hand acquired simultaneously are shown on the right.

3.2. Historical Developments in Face Recognition

> "This (face) recognition problem is made difficult by the great variability in head rotation and tilt, lighting intensity and angle, facial expression, aging, etc."
> – *Woodrow Bledsoe, 1966*

Although human beings have been using faces to recognize one another since time immemorial, the work on enabling computers to recognize human faces was started in the mid-1960s by Woodrow W. Bledsoe and his colleagues at Panoramic Research (Bledsoe, 1966). Bledsoe qualified his face recognition system as a "man-machine" system, because it required human experts to first manually locate some facial landmarks on a photograph. The matching was then done automatically based on 20 normalized distances derived from these facial landmarks (e.g., width of the mouth, width of eyes, etc.). A system to automatically extract such facial landmarks was proposed in Takeo Kanade's Ph.D. thesis (Kanade, 1973) published in 1973, which can be considered to have presented the first fully automated face recognition system. Figure 10 presents a brief summary of the milestones in the development of face recognition algorithms.

While the earliest face recognition systems were based on geometric features (distances between pre-defined landmarks), the Eigenface approach popularized by Turk and Pentland in 1991 (Turk and Pentland, 1991) was based on holistic facial appearance.[11] Holistic appearance-based techniques generate a compact representation of the entire face region in the acquired image by mapping the high-dimensional (4,096 dimensions for a 64 × 64 image) face image into a lower dimensional sub-space. This sub-space is defined by a set of representative basis vectors, which are learned using a training set of images. The local feature analysis method of Penev and Atick (Penev and Atick, 1996) and the Fisherface method of Belhumeur et al. (Belhumeur et al., 1997) are other examples of holistic appearance-based face recognition.

The elastic bunch graph matching approach of Wiskott et al. (Wiskott et al., 1997) was a pioneering work in model-based face recognition. Model-based techniques try to derive a pose-independent representation of the face images by building 2D or 3D face models. These schemes typically require the detection of several fiducial or landmark points in the face (e.g., corners of eyes, tip of the nose, corners of the mouth, and the chin), which leads to increased complexity compared to appearance-based techniques. The morphable model proposed by Blanz and Vetter (Blanz and Vetter, 2003) advanced the use of 3D models in face recognition by exploiting both facial texture and shape features.

Since appearance-based schemes use the raw pixel intensity values, they are quite sensitive to changes in ambient lighting and facial expressions. Therefore, texture-based methods like Scale Invariant Feature Transform (SIFT) (Lowe, 2004) and Local Binary Patterns (LBP) (Ojala et al., 2002) were developed. These methods use more robust representations that characterize the texture of an image using the distribution of local pixel values. Sparse representation coding (SRC) (Wright et al., 2009) and face recognition based on deep learning (Sun et al., 2014; Taigman et al., 2014) are some of the more notable advances in the area of face recognition in the last decade.

Most of the face recognition techniques assume that faces can be aligned and properly normalized (both geometrically and photometrically). The alignment is typically based on the location of the two eyes in the face. The face detection scheme developed by Viola and Jones (Viola and Jones, 2004) is considered a milestone because it enabled faces to be detected in real-time even in the presence of background clutter, a situation commonly encountered in applications such as surveillance. Even though the Viola-Jones face detector has demonstrated excellent performance in real-time applications, it is still challenged when confronted with non-frontal facial poses, illumination changes, occlusion, etc.

While advancements in algorithms have contributed to improvements in face recognition accuracy, practical face recognition systems have also benefited due to improvements in face acquisition systems, be it 2-D (intensity image), 3-D (intensity and depth/range image), infrared, or video cameras.

One of the major turning points in the history of camera technology was the introduction of digital cameras[12] in the early 1990s. Due to improvements in semiconductor technology, the frame rate, spatial resolution (pixel density), and quality (pixel sensitivity) of image sensors has improved significantly (Suzuki, 2010) and it has been claimed that the performance of state-of-the-art digital cameras can match that of the human eye (Skorka and Joseph, 2011). At the same time, these image sensors have also become smaller and cheaper making it possible to embed them in many personal electronic devices such as computers, tablets, and mobile phones. Today, it is possible to cap-

[11] Earlier work by Sirovich and Kirby in 1987 had shown that faces could be represented by Principal Component Analysis (Sirovich and Kirby, 1987)

[12] Though Kodak invented the first digital camera in 1975, they did not become commercially available until 1990.

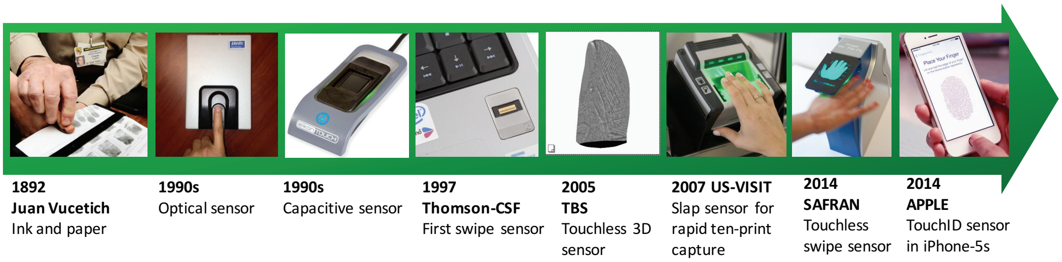

Fig. 9. Evolution of fingerprint sensing technology. Fingerprint sensors have evolved in two ways. On the one hand, they have become compact in size and cheaper in cost, which makes it possible to embed fingerprint sensors in devices such as laptops or mobile phones. While some applications still use a large surface area fingerprint sensor (for capturing a full fingerprint impression resulting in higher accuracy), they are equipped with several advanced functionalities. A typical example is the slap sensor used in the US-VISIT program (Department of Homeland Security, 2013), which can capture impressions of multiple fingers simultaneously, thereby facilitating rapid capture of ten-prints. Other examples include the 3D fingerprint sensor introduced by TBS in 2005 (Parziale and Diaz-Santana, 2006) and the touchless fingerprint sensor introduced by Safran in 2010 (Fourre et al., 2011), which can acquire the images of multiple fingers on-the-fly as the user moves his hand across the device.

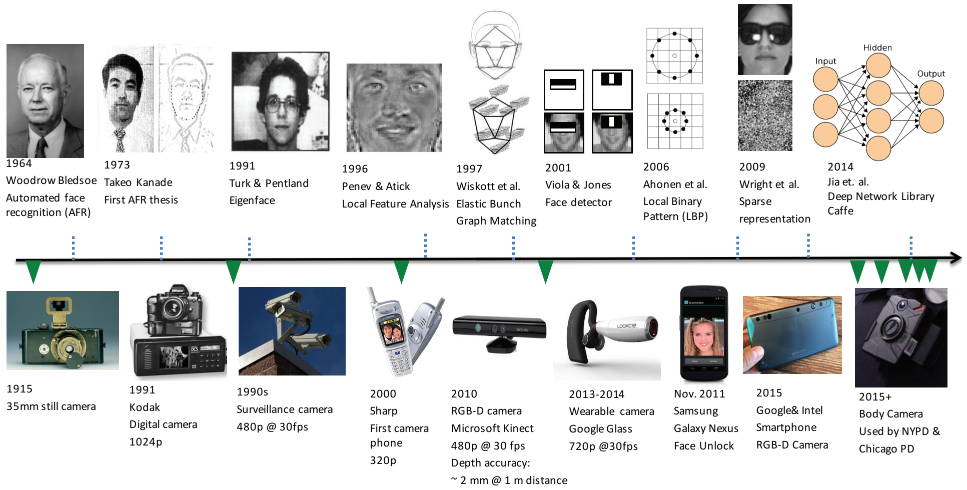

Fig. 10. Major milestones in the history of automated face recognition. While the events in the top row highlights the important landmarks in the evolution of face recognition algorithms, those in the bottom row correspond to the turning points in the development of face acquisition systems.

ture good quality face images using smartphones or wearable devices such as Google glass (Google, Inc., 2013). Furthermore, low cost cameras that can capture 3-dimensional images in real-time are also becoming available now (e.g., Microsoft Kinect (Khoshelham and Elberink, 2012)). Consequently, real-time face recognition has been made feasible in a wide range of applications where the user is cooperative and the face image is acquired in a controlled environment (e.g., access control, deduplication of driver licenses and passports). But, solutions to unconstrained face recognition such as in surveillance applications are still elusive.

3.3. Historical Developments in Iris Recognition

> "For purposes of rapid and reliable person identification,...it is hard to imagine one (unique identifier) better suited than a protected, immutable, internal organ of the eye (iris), that is readily visible externally and that reveals random morphogenesis of high statistical complexity."
> *– John Daugman, IEEE Transactions on PAMI, 1993*

The iris of the eye contains rich textural information that can be used for person recognition. The idea of using iris patterns for human identification was first proposed by Frank Burch in 1936 and the first patent for an iris recognition system was granted to Flom and Safir in 1985 (Flom and Safir, 1987). While Flom and Safir presented ideas for iris image capture, feature extraction, and matching, the first working iris recognition system was developed and implemented by John Daugman in the early 1990s (Daugman, 1993). In fact, Daugman was the first to develop (a) a camera to capture the iris images, (b) image processing algorithms to process the eye images and extract the iris region, and (c) the well-known IrisCode representation to characterize the iris images in the form of a compact binary code.

One of the first major deployments of iris recognition was the one implemented by United Arab Emirates (UAE) for border control in 2001. This was soon followed by the use of iris recognition to facilitate immigration control for frequent travelers at the Amsterdam Schipol airport in 2003. An iris recognition based immigration system was also operational at major airports in the United Kingdom for nearly a decade, before it was decommissioned in 2013. Iris-based border control systems are also being used to enable quicker immigration clearance for pre-approved travelers between the United States and Canada. Iris recognition was also extensively used by the United States military for field operations in Afghanistan and Iraq. Recently, several large-scale national identification systems for civilians such as India's Aadhaar project (Planning Commission, Government of India, 2012), Mexico's national ID program, and Indonesia's e-ktp program include iris as one of the primary biometric modalities.

Iris recognition systems have also benefitted greatly from the huge improvements in image sensors. The early iris cameras such as the one developed by Daugman and other commercial cameras developed in the 1990s were not only bulky and expensive, but also required high levels of cooperation from the user to provide good quality iris images (see Figure 11). Typically, the textural details of the iris are imaged using a camera that is sensitive to near infra-red (NIR) illumination. NIR illumination is required to capture the texture details of dark-colored irides (which are not clearly resolved under visible light) and to make the sensing less intrusive (NIR illumination cannot be perceived by the human eye). Moreover, the user needs to be cooperative and hold his head in a relatively stable position while looking directly into the camera so that the iris images are not degraded due to factors such as partially closed eyelids, intruding eyelashes, extremely dilated or constricted pupil, or off-axis acquisition.

However, iris cameras developed in the last decade are more portable, compact, affordable, and easy to use. For example, SecuriMetrics introduced a portable iris scanner in 2004. Sarnoff developed the "iris-on-the-move" system in 2006 (SRI International, 2013), which could capture iris images at a three meter standoff distance and from subjects walking at 1 meter per second. In 2013, companies such as A-Optix and Delta-ID have shown that it is possible to capture good quality iris images using smartphones.

3.4. Developments in Other Biometric Traits

1. Ear: The appearance, structure, and morphology of the human ear has been studied as a biometric cue for a number of years (Abaza et al., 2013). While most face recognition systems extract the attributes of the human face from frontal images, the visibility of the ear in non-frontal poses of the face (e.g., side view) makes it a viable biometric in many scenarios. The human ear is observed to exhibit variations across individuals as assessed by the curves, surfaces, and geometric measurements pertaining to the visible portion of the ear, commonly referred to as the pinna. As a biometric trait, the ear offers several advantages: (a) the structure of the ear has been observed to be stable despite aging, and ear growth is almost linear after the age of four; (b) the ear, unlike other facial features, is minimally impacted by changes in facial expression; and (c) image acquisition does not involve explicit contact with the sensor.

   Although several algorithms for ear detection and matching have been proposed in the literature, large-scale public evaluation of ear recognition algorithms has not been conducted. Further, there are not many commercial biometric systems at this time that explicitly utilize features of the ear for human recognition.[13] But the performance of ear recognition algorithms has been tested on some standard ear datasets. Experiments suggest that ear images obtained under controlled conditions can result in good recognition accuracy. However, the performance of ear recognition methods on non-ideal images obtained under varying illumination and occlusion conditions is yet to be established. Several challenges have to be overcome to make this possible.
2. Gait: The demand for human identification at a distance has gained considerable traction, particularly due to the need for covertly recognizing individuals in unconstrained

[13] An example application can be found at `http://www.descartesbiometrics.com/ergo-app/`.

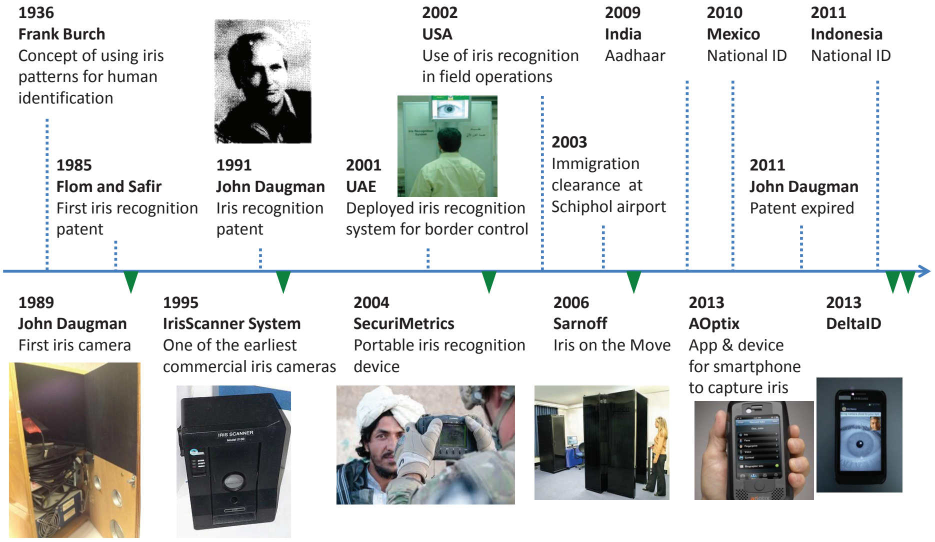
1936
Frank Burch
Concept of using iris patterns for human identification
1985
Flom and Safir
First iris recognition patent
1991
John Daugman
Iris recognition patent
2001
UAE
Deployed iris recognition system for border control
2002
USA
Use of iris recognition in field operations
2003
Immigration clearance at Schiphol airport
2009
India
Aadhaar
2010
Mexico
National ID
2011
Indonesia
National ID
2011
John Daugman
Patent expired
1989
John Daugman
First iris camera
1995
IrisScanner System
One of the earliest commercial iris cameras
2004
SecuriMetrics
Portable iris recognition device
2006
Sarnoff
Iris on the Move
2013
AOptix
App & device for smartphone to capture iris
2013
DeltaID

Fig. 11. Major milestones in the history of automated iris recognition. While the events in the top row highlight the important landmarks in the evolution of iris recognition algorithms, those in the bottom row correspond to the turning points in the development of iris sensors.

environments with uncooperative subjects. In such environments, the person of interest may not be interacting with the biometric system in a concerted manner. Further, the individual might be moving in this environment characterized by variable illumination and a non-uniform background. Biometric modalities such as fingerprint and iris cannot be easily acquired at large stand-off distances. On the contrary, the face and gait modalities can easily be acquired at a distance, although the smaller spatial resolution of the face at long distances can degrade accuracy of face recognition systems. As a result, gait-based human recognition has received some interest for biometric recognition at a distance (Nixon et al., 2006). *Gait* is defined as the pattern of locomotion in animals. Human gait, therefore, is the manner in which people walk. While the formal definition of gait refers to human *motion*, practical algorithms for gait recognition include both dynamic and static features (such as body shape) of the moving human body. It can be viewed as a behavioral trait that is impacted by the musculo-skeletal structure of the human body.

Gait recognition is perceived as an attractive solution for distance-based identification for a number of reasons. First and most importantly, human gait has been observed to have some person-specific characteristics. Psychological studies by Cutting and Kozlowski showed that humans are capable of deducing gender and recognizing known individuals based on gait. Second, the gait biometric can be acquired passively and, therefore, explicit subject interaction is not required for data acquisition. Passive collection is beneficial in an environment where subjects are being observed covertly. Finally, discriminatory features of human gait can be extracted in low resolution images. This suggests that expensive camera systems may not be required for gait recognition.

The matching performance of gait recognition algorithms is impacted by factors such as clothing, footwear, walking surface, walking speed, walking direction (with respect to the camera), etc. Further, the gait pattern of an individual can change over time, especially with variations in body mass The impact of these factors is difficult to mitigate and, therefore, evaluation of gait recognition algorithms has been predominantly conducted in controlled environments. This has prevented the incorporation of gait recognition in commercial biometric systems.

3. Hand Geometry: Hand geometry, as the name suggests, refers to the geometric structure of the hand (Jain et al., 1999; Duta, 2009). This structure includes width of the fingers at various locations, width of the palm, thickness of the palm, length of the fingers, contour of the palm, etc. Although these metrics do not vary significantly across the population, they can still be used to verify the identity of an individual. Hand geometry measurement is non-intrusive and the verification involves a simple processing of the resulting features. Unlike palmprint (Kong et al., 2009), this method does not involve extraction of detailed features of the hand (for example, wrinkles on the skin).

   Hand geometry-based verification systems have been commercially available since the early 1970s. The earliest literature on the hand geometry biometric is in the form of patents or application-oriented description. Hand geometry systems have been successfully deployed in several applications including nuclear power plants, border control systems (e.g., Ben Gurion airport in Tel Aviv), recreational centers and time-and-attendance systems. In these applications, the biometric system typically operates in the *verification* mode. Since the hand geometry of subsets of individuals can be similar, the *identification* accuracy due to this biometric modality can be low. Further, the shape of an individual's hand can change with time - a factor that is especially pronounced in young children. More recent research has explored the use of hand geometry in conjunction with fingerprints and low-resolution palmprints in a multibiometric configuration for improved accuracy.

4. Periocular: The periocular region represents the region around the eyes. It predominantly consists of the skin, eyebrow, and eye. The use of the periocular region as a biometric cue represents a good trade-off between using the entire face region or using only the iris for recognition (Park et al., 2011). When the entire face is imaged from a distance, the iris information is typically of low resolution; this means the matching performance due to the iris modality will be poor. On the other hand, when the iris is imaged at small standoff (typically, 1 meter), the entire face may not be available, thereby forcing the recognition system to rely only on the iris. However, the periocular biometric can be used for a wide range of distances. Periocular images can also be captured in the NIR spectrum to minimize illumination variation compared to visible spectrum.

   Some of the other benefits in using the periocular biometric trait are as follows: 1) In images where the iris cannot be reliably obtained (or used), the surrounding skin region may be used to either confirm or refute an identity. Blinking or off-angle poses are common sources of noise during iris image acquisition. 2) The periocular region can offer information about eye shape that may be useful as a soft biometric. 3) When portions of the face pertaining to the mouth and nose are occluded, the periocular region may be used to determine the identity. 4) The design of a newer sensor is not necessary as both periocular and face regions can be obtained using a single sensor.

   Recent studies on person identification using periocular traits, both in visible and NIR spectra, show modest identification accuracies (over 80%). However, such an accuracy is possible only when the images are of good quality and exhibit low intra-class variations. It has also been shown that the periocular trait can assist person identification when the face is occluded.

4. State-of-The-Art in Biometric Recognition

The evolution of biometric recognition is an on-going process and biometric systems are experiencing continuous improvements in performance and usability. However, a clear assessment of the present state-of-the-art is required to appreciate the progress made thus far and set the baseline for future improvements. Many independent third-party technology evaluations have been conducted primarily by NIST over the last 20 years for fingerprint, face, iris, and voice modalities. These NIST evaluations serve as an excellent resource to benchmark the current recognition performance of various biometric systems. In general, the error rates of a biometric system depend on a number of test conditions. Consequently, the NIST evaluations tend to be quite extensive and include results obtained under a variety of test conditions. An in-depth discussion of these results is beyond the scope of this paper and we restrict ourselves to highlighting only a few key results.

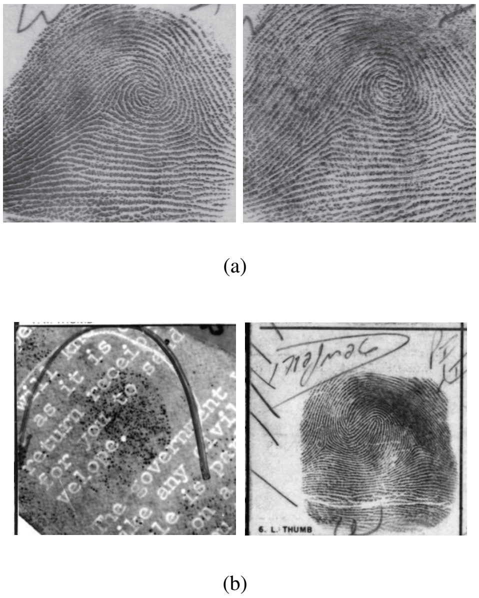

(a)

(b)

Fig. 12. The current state of fingerprint recognition. Recognition based on two rolled or plain fingerprints captured using live-scan sensors (as shown in (a)) can be considered as an almost solved problem as demonstrated by the results of NIST FpVTE 2003 and FVC evaluations. However, the results of NIST ELFT evaluation indicate that fully automated latent identification (i.e., latent-to-rolled print matching), as shown in (b), is still an open problem.

In the case of fingerprint recognition, the challenges vary depending on the type of fingerprint images. Therefore, NIST evaluations are also conducted separately for matching between plain/rolled prints and latent-to-rolled print matching (see Figure 12). The Fingerprint Vendor Technology Evaluation (FpVTE) conducted by NIST over a decade back (in 2003) (Wilson, 2004) shows that the best commercial fingerprint recognition system can achieve a True Acceptance Rate (TAR) of 99.4% at a False Acceptance Rate (FAR) of 0.01% for plain-to-plain matching based on fingerprint data collected from various government sources in the United States. The results of FpVTE 2012 have not yet been released by NIST at the time of writing this paper.

Multiple editions of the Fingerprint Verification Competition (FVC) (University of Bologna, 2006) have also been conducted by the University of Bologna since 2000 to benchmark the performance of different fingerprint recognition algorithms for plain-to-plain fingerprint matching. The results of FpVTE 2003 and the various editions of FVC indicate that the technology for plain-to-plain (as well as rolled-to-rolled) fingerprint matching is fairly mature and very high accuracy can be obtained under typical conditions. However, there may be some scope for improvement in accuracy when the user is uncooperative and provides distorted or partial fingerprint images or if the image quality is very poor due to finger skin conditions.

The results of different phases of Evaluation of Latent Fingerprint Technologies (ELFT) conducted by NIST confirm that the problem of latent-to-rolled print matching is inherently more challenging compared to plain-to-plain matching. The best rank-1 accuracy obtained in ELFT-EFS Phase 2 was only

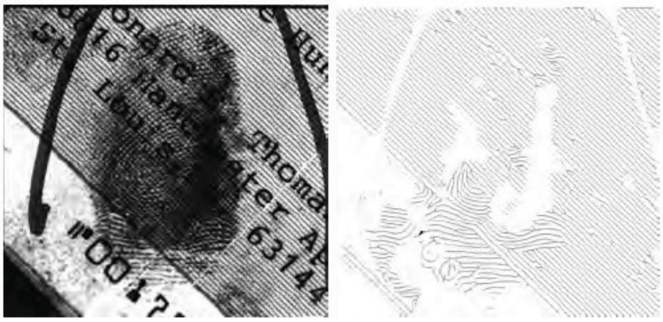

Fig. 13. Poor quality of the latent fingerprint images makes it difficult to reliably extract features from latents. The image on the left is a poor quality latent in which the ridge pattern of interest is smudged and occluded by the presence of structured noise (text, lines, etc.). On the right is shown the ridge skeleton extracted by a commercial fingerprint SDK from the image on the left. Since the SDK used for extracting the above ridge skeleton is not specifically designed for processing latent prints, it fails to extract the correct fingerprint ridges from the latent.

63.4% (M. Indovina *et al.*, 2009). The largest public-domain latent fingerprint database is the NIST Special Database-27 (NIST SD-27) and the best reported rank-1 accuracy on this database is 72%. These numbers clearly show that the problem of fully automated processing and matching of latent prints to rolled impressions or other latent prints is still far from being solved.

The major reason for the deterioration in identification accuracy from plain prints to latent prints is the poor quality of the latent fingerprint images. Latent fingerprints typically contain ridge information from only a partial area of a finger. While a typical rolled fingerprint has around 106 minutiae, a latent print may contain only 21 usable minutiae.[14] Even this partial ridge information in a latent is often smudged, blurred, or occluded by background text and markings, and may exhibit large nonlinear distortion due to pressure variations. When the fingerprint quality is very poor, it becomes very difficult to reliably extract the minutiae and ridge features as shown in Figure 13.

Evaluating the state-of-the-art in face recognition is more complex because of the large scope for variability in face images due to a number of factors such as aging, pose, expression, and illumination. The NIST Face Recognition Vendor Test (FRVT) 2012[15] indicates that face recognition systems can achieve a TAR of approximately 96% at a FAR of 0.1% when matching mugshots of the face (frontal face images obtained under a controlled environment at the time a suspect is booked at the police station). When presented with face images obtained during the visa application process, the TAR improves to nearly 99% at the same FAR of 0.1%. This is because face images for visa processing have more stringent guidelines on illumination, background, and occlusion. While the above results are impressive, they are applicable only to a small number of applications where such good quality face images can be captured from cooperative subjects.

Table 2. Summary of True Accept Rate (TAR) at 0.1% False Accept Rate (FAR) when different face recognition algorithms were evaluated on the NIST Special Database-32, which is also known as the Multiple Encounter Dataset (MEDS II).

Algorithm	TAR at 0.1% FAR
Eigenfaces	9%
Fisherfaces	35%
LBP	34%
COTS-A	58%
COTS-B	88%
COTS-C	97%

As pointed out in section 2, the key differentiator for face recognition compared to fingerprint and iris is the ability to capture face images covertly. However, the face images captured covertly tend to exhibit more intra-class variations. A reasonable indicator of the performance under mildly challenging conditions is the accuracy of various face recognition algorithms on the NIST Special Database-32, which is also known as the Multiple Encounter Dataset (MEDS) (National Institute of Standards and Technology - Information Technology Laboratory, 2010). This database contains face images exhibiting relatively large intra-class variations such as pose and illumination changes, compared to mugshots, as shown in Figure 14. Some well-known face recognition algorithms such as Eigenfaces, Fisherfaces, Local Binary Patterns (LBP), as well as three commercial-off-the-shelf (COTS) face matchers were evaluated on the MEDS II database and the TAR at 0.1% FAR is summarized in Table 2.

The results in Table 2 show a wide gap between the performance of the best COTS matcher and some of the most popular algorithms that are often considered in academic research. The difference in performance is also rather high among three of the best COTS matchers. This indicates the need to carefully choose the baseline system when developing new algorithms for face recognition. It may be very easy to demonstrate an improvement in performance by choosing an outdated baseline (say Eigenfaces). A new face recognition algorithm cannot be considered as an advancement of the state-of-the-art unless one can demonstrate better performance compared to either state-of-the-art COTS face matchers or the best performing algorithms published in the literature.

The Labeled Faces in the Wild (LFW) database contains face photographs for studying the problem of unconstrained face recognition. This dataset contains more than 13,000 images of 5,749 subjects. A number of researchers have reported performance of their algorithms on this dataset. Results on a specific protocol used on this dataset can be seen in Table 3.

One of the largest independent technology evaluations of iris recognition is the NIST IREX III evaluation (Grother et al., 2012). The database used in this test contains approximately 6.1 million iris images acquired from nearly 4.3 million eyes. Among the 95 different algorithms considered in this evaluation, the best algorithm had a false negative identification rate of approximately 2.5% when a single eye was used per person and the threshold was set such that there were no more than 25

[14] Based on images in the NIST SD27 database

[15] `http://www.nist.gov/itl/iad/ig/frgc.cfm`

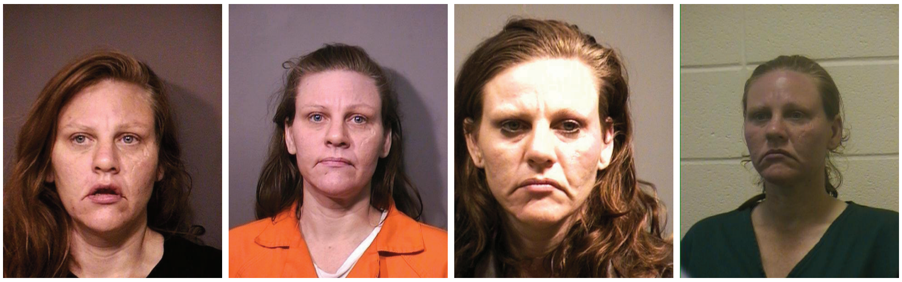

Fig. 14. Sample face images of one subject from the MEDS II database illustrating intra-user variations due to factors such as illumination, pose, expression, and aging.

Table 3. Mean classification accuracy, along with standard error, of some face recognition algorithms that were evaluated on the Labeled Faced in the Wild (LFW) dataset under the "Image-Restricted, No Outside Data Results" protocol. For an explanation of the acronyms, please see http://vis-www.cs.umass.edu/lfw/results.html

Algorithm	Accuracy
Eigenfaces	0.6002 ± 0.0079
Fisher vector faces	0.8747 ± 0.0149
MRF-Fusion-CSKDA	0.9589 ± 0.0194

false positives in every 10^{13} iris comparisons. This ability to operate at a very low probability of a false match is one of the key advantages of iris recognition. It was observed that pupil dilation and constriction has a significant impact on the recognition accuracy, and the size of the iris templates ranged from 1 kilobyte (KB) to 20 KB. Some examples of iris images that could not be recognized correctly during IREX III evaluation are shown in Figure 15. These iris images are of very poor quality, mainly because the users did not interact correctly with the iris sensor.

Apart from fingerprint, face, and iris, significant progress has also been achieved in the case of voice biometrics (also known as speaker verification) over the last two decades. The results of 2012 NIST Speaker Recognition Evaluation (SRE) (Greenberg et al., 2012) show a TAR of approximately 93% at a FAR of 0.1%. This high level of accuracy was achieved despite the challenging nature of the NIST SRE 2012 evaluation, which required the algorithms to detect if a target speaker had spoken in a given test speech segment with significant background noise.

A closer look at the performance of state-of-the-art fingerprint, face, iris, and voice biometric systems indicates that it is possible to achieve very low error rates when the respective biometric samples are acquired under controlled conditions with the cooperation of the user. From the accuracy perspective[16], biometric recognition can be considered as a solved problem in applications where the acquisition of good quality biometric samples from cooperative users is not an issue. However, the fundamental problems in biometrics, namely, feature extraction and matching, become more challenging when the biometric samples are not captured in a controlled environment or if the user is non-cooperative. As shown in Figure 16, there is a huge gap between the accuracy of biometric systems evaluated on good quality biometric samples (bottom left of Figure 16) and that of systems evaluated on poor quality samples (top right of Figure 16). This suggests that the development of better feature extraction and matching algorithms to handle poor quality biometric samples is a fertile ground for research.

5. Unsolved Problems

The unsolved problems in biometric recognition can be divided into two categories: (i) problems that involve fundamental issues related to design of recognition systems and (ii) problems that are specific to applications that will use biometric recognition. As discussed in section 2, questions about the distinctiveness and permanence of a biometric trait have not been adequately addressed by the biometrics research community. Moreover, feature extraction and matching schemes that can handle poor quality biometric samples (e.g., face images from a surveillance video or latent fingerprint images) need to further developed. In the case of application-specific problems, the two main unresolved issues are (i) techniques to shield a biometric system from adversarial attacks/threats and provide assurances on user privacy, and (ii) techniques to assess usability of a biometric system and estimate the return on investment. Finding viable solutions to these unresolved problems will not only strengthen the case for biometrics in existing applications, but also open up new applications for biometric recognition.

5.1. Distinctiveness of Biometric Traits

The concept of quantifying the distinctiveness of a biometric trait (in other words, estimating the individuality of a biometric trait) can be understood from the following simple analogy. Suppose that users in a person recognition system are identified based on a 10-digit personal identification number (PIN). The theoretical limit on the number of users who can be uniquely identified by such a system is 10 billion. In other words, we can say the probability that any two users in such a person recognition system will have the same PIN is 1 in 10 billion. However,

[16] Note that even if a high matching accuracy can be achieved in a technological test, other requirements such as throughput, cost, and usability may need to be satisfied before a biometric system becomes suitable for a particular application. Moreover, a biometric system can still be very useful in an application, even though "zero error rate" may never be achieved under operational settings.

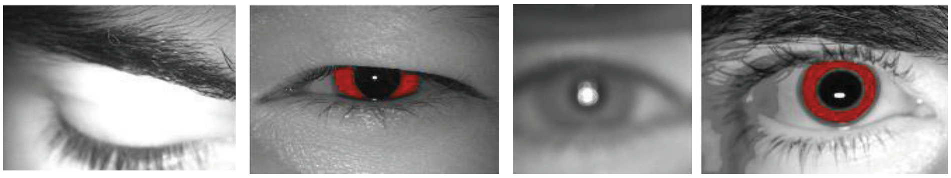

Fig. 15. Examples of iris images from the NIST IREX III evaluation that could not be successfully recognized. It is difficult to extract reliable iris texture features from these images because the eyelids are fully or partially closed (first two images from the left), the images exhibit excessive blur (third image from the left), or the images are highly quantized (rightmost image).

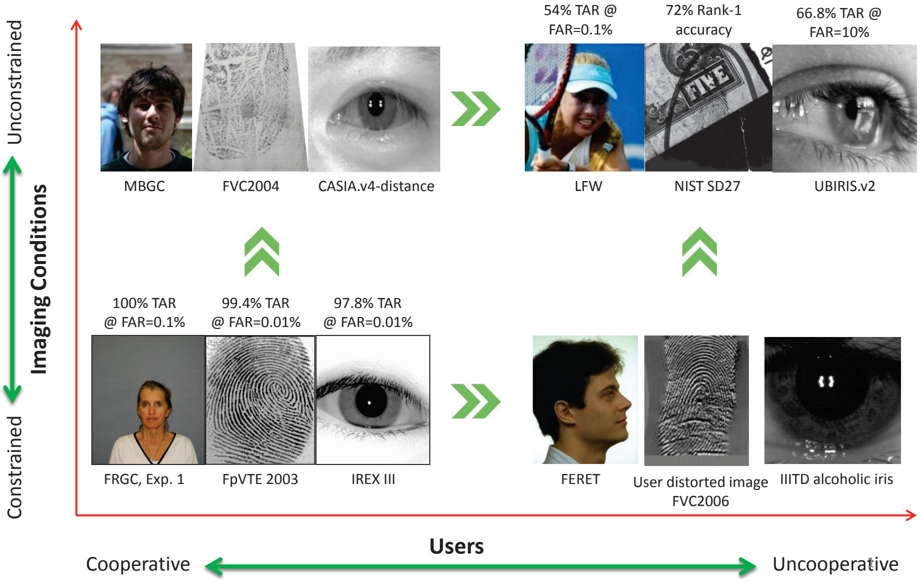

Fig. 16. Degradation in the recognition accuracy of biometric recognition algorithms as the quality of the biometric samples decreases due to unconstrained sensing conditions and/or uncooperative subjects. When good quality samples can be acquired from cooperative users under controlled conditions (as shown in the bottom left of the graph), biometric recognition can be considered as a solved problem from the accuracy perspective. However, biometric recognition is far from being a solved problem when poor quality samples such as those shown on the top right of this image are presented as inputs to the recognition system.

it is not feasible to achieve this theoretical limit in practice because the users seldom choose a random PIN.

Similar to the case of PIN, it is important to know how many users can be uniquely identified by a biometric system based on a specific biometric trait (e.g., right index fingerprint). This information is needed when designing large-scale biometric identification systems involving tens or hundreds of million users in a biometric database. Understanding the distinctiveness of traits can help in the design of biometric templates with sufficient capacity to distinguish between these users. If a single biometric trait is not sufficient to meet the desired accuracy, it is essential to know how many traits (multiple fingerprints, multiple irides, face, etc.) would be required to identify all individuals in a target population with the desired accuracy. For example, consider the Aadhaar system in India where the goal is to achieve de-duplication of more than one billion individuals. This system currently uses ten fingerprints and two irides to perform de-duplication. However, there is no rigorous scientific basis to either claim that these twelve traits are sufficient to achieve complete de-duplication of a billion identities or the same purpose can be met with a smaller number of traits. Further, in forensic applications it is necessary to provide an estimate of the probability that any two or more individuals may have sufficiently similar biometric samples in a given target population. This is needed, for example, to provide credence to latent fingerprint evidence.

One of the basic issues in estimating the individuality of a biometric trait is determining the information level at which the individuality should be measured. It is possible to define individuality based on (i) the biological trait ($I(Y;\mathbf{B})$), (ii) the sensed samples/images recorded from the body trait ($I(Y;\mathbf{M})$), and (iii) the features extracted from the sensed samples ($I(Y;\mathbf{X})$). For example, one can analyze the distinctiveness of the 3-dimensional fingerprint pattern at the tip of a finger, the 2-dimensional flat/plain fingerprint image obtained by pressing the finger against a fingerprint scanner, or (iii) the minutiae features extracted from the 2-D fingerprint image.

In general, it is very difficult to directly measure the individuality of the biological trait because only the sensed samples are available for analysis. Since the sensed samples include different types of noise in addition to the biometric information, estimating the individuality based on raw samples is also very challenging. Furthermore, individuality at the trait or sample level may be of little use except as an upper bound on the individuality of a biometric system. This is because the recognition will be eventually based on the features extracted from the sensed images. For instance, while human faces may be highly distinctive when observed at sufficient detail (e.g., 3-D shape, scars, marks, etc.), individuality of appearance-based 2-D face recognition systems may be limited by the proportion of identical twins[17] in the target population. Consequently, research on the individuality of biometric traits has primarily focused on estimating the individuality based on the extracted features.

The primary difficulty in estimating individuality of a biometric trait based on its feature representation is the lack of robust statistical models to accurately characterize the intra- and inter-subject variations. Consequently, estimating the entropy functions $H(\mathbf{X})$, $H(\mathbf{X}|Y)$, or $H(Y|\mathbf{X})$ becomes a challenging task. Most of the attempts made thus far to estimate the individuality of biometric traits had to make simplifying assumptions in order to keep the problem tractable (Zhu et al., 2007).

Alternatively, attempts have been made to abandon the idea of accurately modeling the features and indirectly estimate individuality based on the match score distributions (Neumann et al., 2007). The basic assumption underlying this approach is that $H(Y|\mathbf{S})$, the conditional entropy of Y given the match score data $\mathbf{S}$, is a good upper bound for $H(Y|\mathbf{X})$. The main limitation of this approach is the need to account for the tails of the match score distributions, which in turn requires a very large number of biometric samples. This is often infeasible due to time and cost considerations.

A good example of analyzing biometric features based on match score distributions is the analysis of impostor score distribution using IrisCodes extracted from $632,500$ different iris images (Daugman, 2006). In (Daugman, 2003), it was estimated that a $2,048$ bit IrisCode representation contains approximately 249 degrees of freedom. However, this result is based on a simple matching model that ignores the need to test multiple relative rotations of the IrisCode. Therefore, one cannot directly conclude that the entropy of an IrisCode template is 249 bits. Moreover, it is not straightforward to obtain a precise estimate of individuality of the IrisCode representation using the above result because it fails to take into account the genuine score distribution (consequently, intra-subject variations are not modeled).

Finally, one can argue that the ability of a biometric system to achieve very low error rates can be considered as evidence of high individuality of the underlying biometric trait. This is because $H(Y|\hat{Y})$ can be considered as a upper bound on $H(Y|\mathbf{X})$, where $H(Y|\hat{Y})$ is a function of the error rates of a biometric system. Estimating the individuality based on empirical error rates has two main limitations: (i) since the error rates are database-dependent, it is not easy to extrapolate them when the population size increases by orders of magnitude or when the population characteristics change, and (ii) the resulting estimate is only a loose lower bound on the true individuality. For example, if a biometric system is able to achieve an FMR of 1 in a trillion, it implies that the entropy of the biometric template could be approximately 40 bits. This is the equivalent to the guessing entropy of a randomly chosen 6 character password chosen from an alphabet of 94 characters (Burr et al., 2006). Intuitively, one would expect the true individuality of a biometric trait to be much higher than this value.

5.2. *Persistence of Biometric Traits*

Persistence of a biometric trait is related to the notion of aging. Aging refers to changes in a biometric trait or the corresponding template over a time span, which can potentially impact the accuracy of a biometric system. For the sake of clarity, we distinguish between two types of aging: trait aging and

[17] It has been reported that the birth rate of monozygotic (identical) twins is about three in every 1,000 births worldwide and this number is gradually increasing due to the rise in fertility treatments (Wikipedia, 2002)

template aging. Trait aging refers to the biological change in a trait over a person's lifetime. This change is inevitable and, unlike other types of intra-subject variations, cannot be easily controlled by the individual. For example, changes in a person's facial structure and appearance can occur over time due to the effects of biological aging. This can, in turn, impact the accuracy of face matchers as shown in Figure 17.

Template aging, on the other hand, refers to changes in a person's biometric template (i.e., the feature set extracted from the biometric trait) over time. While template aging is certainly related to trait aging, it must be noted that the extraction of invariant features from a biometric trait can mitigate the impact of trait aging on template aging. In the case of fingerprints, it is well known that the friction ridge pattern varies over time due to age-related as well as occupation-related changes in the outer skin, sebaceous gland activity, etc. However, these changes, for the most part, do not significantly impact the distribution of minutiae points in the fingerprint image (see Figure 18). This explains the use of minutiae points in defining fingerprint templates that have been successfully used for over 100 years. Consequently, fingerprint trait aging does not necessarily result in template aging. Furthermore, the persistence of a biometric trait varies from person to person.

Since every biological agent experiences aging, it would be facetious to assume that biometric traits are persistent over time. The question that is yet to be answered by the biometric community is the following: can the degree of permanence of a biometric trait/template be computed? In other words, is it possible to measure and predict the degree of change ($\boldsymbol{\eta}_a$) that a certain trait or template is expected to encounter over an individual's lifetime? An answer to this question would allow for the system to periodically, and systematically, update the biometric template of a user in order to account for age related changes (Uludag et al., 2004b).

The impact of age on the performance of face recognition systems is well documented (Ramanathan et al., 2009). This is primarily because of the availability of datasets such as FG-NET (Crowley, 2004) and MORPH (Ricanek and Tesafaye, 2006) (also see (Hager, 2013) for a list of other face aging databases). Several algorithms have been proposed to handle the issue of age variation in face recognition (Park et al., 2010). Most of the proposed techniques are learning-based schemes in which the pairwise time-lapsed face images of a large number of subjects are used to deduce an aging model from both a texture and geometric perspective. The learned model is then used during the matching phase to account for potential disparity in age between the gallery and the probe images. Accounting for this disparity has resulted in improvement in the matching accuracy of face matchers. Additionally, methods for estimating the age of a face image have also been developed (Fu et al., 2010), thereby substantiating - from a computer vision perspective - the manifestation of age on faces.

While the issue of age disparity has been extensively addressed for face biometrics, the issue has received less attention in the context of fingerprint biometrics. This is because the configuration of epidermal ridges that constitute a fingerprint has been established to be stable in postnatal life (Galton, 1892; Babler, 1991). However, more recently, the issue of fingerprint persistence was systematically studied by Yoon and Jain (Yoon and Jain, 2015). In their study, fingerprint match scores were analyzed using multilevel statistical models. Longitudinal fingerprint records of 15,597 subjects were sampled from an operational fingerprint database such that each individual had at least five 10-print records over a minimum time span of 5 years. Their analysis showed that: (i) genuine match scores tend to significantly decrease when time interval between two fingerprints in comparison increases, whereas the change in impostor match scores is negligible; and (ii) fingerprint recognition accuracy at operational settings, nevertheless, tends to be stable as the time interval increases up to 12 years, the maximum time span in the dataset.

Recent literature in iris recognition has provided support for template aging (Grother et al., 2013). Several researchers have observed a decrease in True Accept Rate (TAR) when iris templates separated over a long period of time (more than 3 years) were compared (Baker et al., 2009). However, none of these studies was able to directly relate the degradation in genuine match scores with explicit changes in the iris texture itself. Consequently, the notion of iris template aging has remained a controversial issue at the time of writing this paper.

5.3. *Unconstrained Biometric Sensing Environment*

There are some person recognition applications where it is very difficult to impose constraints on how the biometric trait should be acquired. One well-known example is latent fingerprints acquired from crime scenes. For iris recognition, one of the major issue has been the usability of iris sensors. Most available iris sensors require the subject's eye to be in close proximity to the camera and expect the subject to remain still during the acquisition process. User acceptance of iris recognition technology can be greatly enhanced if iris sensors can be designed to capture the iris pattern at a distance and when the subject is on the move (e.g., Sarnoff's iris-on-the-move system (SRI International, 2013)). However, the iris images obtained in this scenario are unlikely to record the texture details on the iris surface with high fidelity and may also exhibit large intra-subject variations (e.g., rotation and occlusion). Hence, more robust algorithms are required to process such iris images.

Another classic example of unconstrained sensing environment is video surveillance, where face images are acquired using closed circuit television (CCTV) cameras that monitor public places. Constant video surveillance is deemed to be a successful deterrent against crime and consequently surveillance cameras have rapidly proliferated around the world, especially in urban areas. As an example, it has been estimated that there are more than 1 million CCTV cameras in the city of London alone and around 4.9 million of them are spread across the United Kingdom (Barrett, 2013). Almost all existing CCTV cameras are passive in the sense that they merely record the video and the stored video is analyzed by human operators only after an abnormal incident has taken place and reported. Real-time video processing is seldom carried out to predict or detect an abnormal incident, or to identify a perpetrator.

The primary challenge in automated video surveillance is how to detect "persons of interest" in a video and then iden-

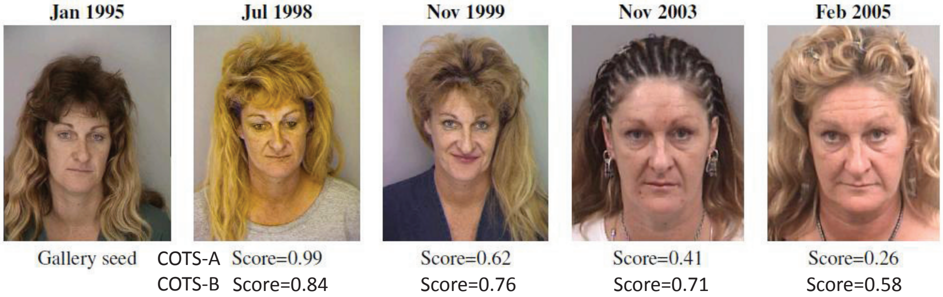

Fig. 17. Degradation in the accuracy of a face recognition system due to trait aging. This figure shows face images of the same person captured over a period of time extracted from a mugshot database provided by the Pinellas County Sheriff's Office (PCSO). Suppose that we consider the first image on the left as the gallery seed and all the other images as probe images. One can easily observe that the match scores output by two state-of-the-art COTS face matchers (denoted as A and B) decrease significantly when the time lapse between the gallery and probe images increases. Note that COTS-B matcher appears to be more robust to aging than COTS-A matcher, indicating that the face template of COTS-B is better than that of COTS-A in compensating for biological aging.

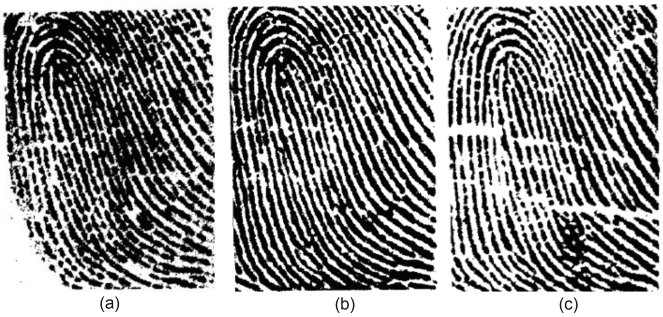

Fig. 18. Herschel's fingerprints at age 7 (a), age 17 (b), and age 40 (c). The pairwise match scores of a state-of-the-art fingerprint matcher for these three fingerprints are: (a) vs. (b) 6,217; (a) vs. (c) 5,032; (b) vs. (c) 5,997; the maximum impostor score of (a) against 10,000 fingerprints in NIST SD4 is 3,325 and that of (b) is 2,935, implying that these three fingerprint images can be claimed to originate from the same finger with high confidence.

tify them using face recognition systems[18]. Here, we focus only on the problem of identifying the "person of interest" in a surveillance video using the face modality[19]. Face recognition in surveillance applications is a very challenging problem due to the following two reasons:

1. The poor quality of face images captured using CCTV cameras. Factors leading to this degradation in quality may include low spatial resolution of the camera, large distance between the subject and the camera, speed at which the subject is moving, illumination variations at the monitored location, and occlusion caused by other objects and people in the scene.
2. Since the subject is not expected to be cooperative (not posing for face capture as in a mugshot scenario), there may be large pose and expression changes as well as occlusion of facial features due to the wearing of accessories like caps and eye-glasses. In some cases, the subject may also intentionally hide his face from the camera to avoid detection.

Apart from the above two issues, surveillance videos typically provide a sequence of face images of the same subject, which needs to be matched against a gallery of still/mugshot images. Generally, it is difficult to establish *a priori* which image in the video sequence is likely to give the correct result. Thus, face recognition in video introduces an additional layer of complexity as well as opportunity because of the availability of evidence provided by multiple probe images that can be combined.

Despite the above challenges, significant progress has been achieved in unconstrained face recognition. This was demonstrated by Klontz and Jain in (Klontz and Jain, 2013), where the authors simulated the scenario of using face recognition to identify the suspects in the Boston marathon bombings (see Figure 19). This was achieved by adding three images each of the two suspects (the Tsarnaev brothers) to a background database of 1 million mugshot images provided by the Pinellas County Sheriff's Office (PCSO). The six images added to the gallery included mugshots as well as face images of the brothers obtained from the social media. The images of the suspects extracted from surveillance cameras and released by the FBI were used as probe images to search the gallery using two state-of-the-art COTS face matchers. It was observed that one of the probe images of the younger brother (Dzhokhar Tsarnaev) matched correctly with his high school graduation photograph included in the gallery (see Figure 19). While this example highlights the potential of automated face recognition technology, it also throws light on the limitations of the state-of-the-art face recognition systems. Firstly, due to issues such as pose, low resolution, and occlusion (e.g., cap and sunglasses), the elder brother (Tamerlan Tsarnaev) could not be successfully identified using both the face matchers. Even in the case of the younger

[18] It is interesting to note that the concept of automatically detecting a person of interest from surveillance video forms the core idea of a television show named *Person of Interest* that is being currently featured on the CBS network.

[19] Note that it is also possible to identify the person using other cues such as gait, ear, and soft biometric traits (e.g., tatoos). Another related problem is person re-identification, where the objective is to track the same person as he/she passes through a network of multiple CCTV cameras.

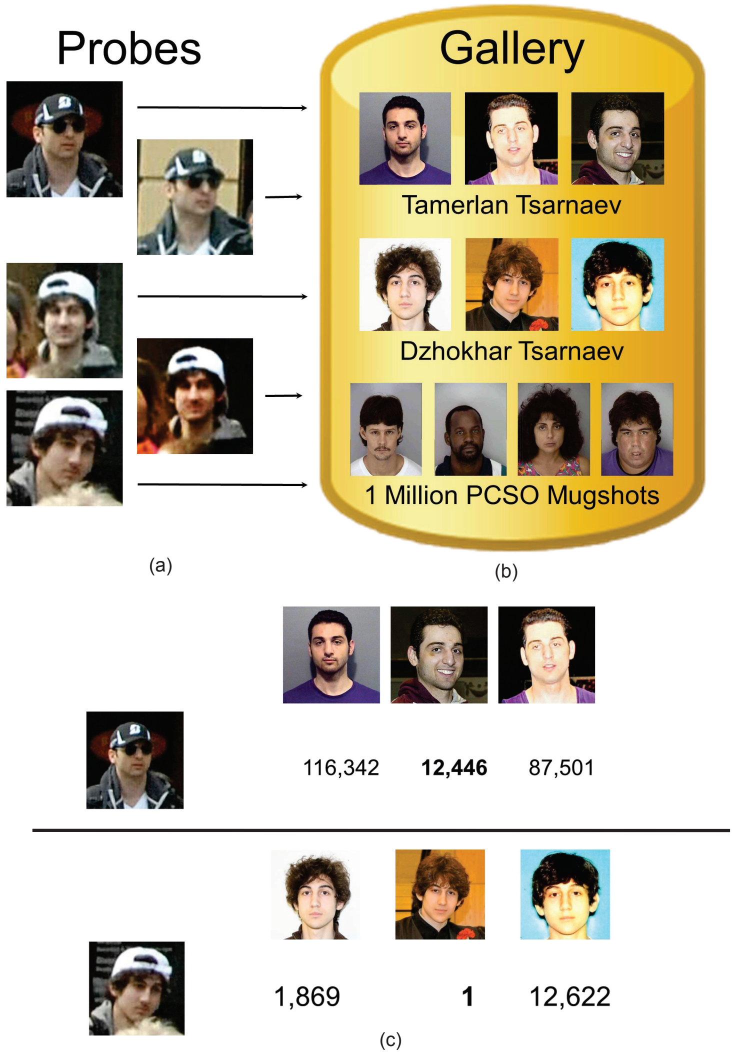

Fig. 19. A simulated example to illustrate how face recognition systems could have been used to identify the suspects in the April 2013 Boston marathon bombings (Wikipedia, 2013). (a) The five face images of the suspects obtained from surveillance videos and released by the FBI. (b) A gallery database constructed by adding three portrait images each of the two suspects (the Tsarnaev brothers) to a background database of 1 million mugshots provided by the Pinellas County Sheriff's Office (PCSO). Note that the six images added to the gallery included mugshots as well as face images of the brothers obtained from the social media. (c) The top retrieval ranks (after demographic filtering) output by a COTS face matcher when the images in (a) are used as probes to search against the gallery in (b). It was observed that one of the probe images of the younger brother (Dzhokhar Tsarnaev) matched correctly (at rank 1) with his high school graduation photograph included in the gallery.

brother, one can argue that the correct match was possible only due to the availability of a graduation photograph with a similar pose. Note that the mugshot image of the younger brother, which is typically the only image available to the law enforcement officials during preliminary investigation, did not result in a successful match. This shows that large improvements in unconstrained face recognition accuracy would be required before fully automated ("lights-out") face recognition systems can be deployed in challenging applications like surveillance.

5.4. System Security & User Privacy

While the main motivation for deploying a biometric system is to protect an application from unauthorized access, there is no guarantee that a biometric system will be completely secure. Just like any other security system, the biometric system may be vulnerable to a number of security threats (see Figure 20), which may eventually affect the security of the end application. These security vulnerabilities may lead to adverse consequences such as denial-of-service to legitimate users, intrusion by unauthorized users, repudiation claims by corrupt users, and erosion of user privacy due to function creep. A number of studies have comprehensively analyzed the security threats faced by a biometric system and suggested remedial measures (Roberts, 2007; Jain et al., 2008).

While many of the adversarial attacks on a biometric system such as Trojan horse, replay, and man-in-the-middle attacks are common to any authentication system and can be addressed by borrowing ideas from secure password-based authentication schemes, there are two vulnerabilities that are more specific to biometric systems. One of them is the problem of spoofing (Marcel et al., 2014), where the biometric sensor is presented with a counterfeit biometric trait (Matsumoto et al., 2002; Chingovska et al., 2012; Zhang et al., 2012). Spoof detection is a critical requirement, especially in unsupervised applications (e.g., authentication on a smartphone) where the presence of a user is not being monitored. The other major threat is the system security and user privacy issues arising from the leakage of biometric template information due to attacks on the template database. Intentional alteration of biometric traits in order to avoid identification (Yoon et al., 2012) is also an emerging threat in some applications (e.g., international border crossing). It must be emphasized that biometric system security and user privacy concerns are important public perception issues, which can potentially derail the success of a biometric system deployment unless they are addressed comprehensively.

Spoof detection typically involves checking for signs of human vitality or liveness (e.g., blood pulse, eye blinking, etc.) and hence, it is also referred to as liveness detection (Marasco and Ross, 2015). To be useful in practice, liveness detection schemes must recognize spoofing attempts in real-time and with high accuracy without causing too much inconvenience to legitimate users. Though spoof detection techniques are generally designed for specific biometric modalities (Parthasaradhi et al., 2005; Antonelli et al., 2006; Nixon and Rowe, 2005; Li et al., 2004; Lee et al., 2006), they can be broadly classified into three main categories. The first approach involves measuring the physiological properties of a live person, which includes blood pulse/pressure, perspiration, spectral/optical properties of the human skin/tissues, electrical/thermal characteristics, and deformation of the muscles/skin. The second approach is based on identifying voluntary or involuntary human behavioral actions like fluctuations in pupil size, blinking, and pupil/eye/head/body movements. The third category is known as the challenge-response mechanism, where the system presents a challenge to the user and measures whether the user responds to the challenge correctly. Examples of challenges include prompting a user to recite a randomly generated phrase/text in speaker verification systems, asking the user to change his or her facial expression (e.g., smile or frown) in face verification systems, or requesting the user to present multiple biometric traits (e.g., different sequence of fingers) in a randomly generated sequence. The key design issues in spoof detection are: (i) how to systematically evaluate the risk of spoofing in a given end application? and (ii) how to select one or more of above approaches to achieve an acceptable tradeoff between spoof detection accuracy and user convenience?

One of the critical steps in minimizing the security and privacy risks associated with biometric systems is to protect the biometric templates stored in the system database. While the risks can be mitigated to some extent by storing the templates in a decentralized fashion (e.g., templates can be stored in individual devices such as smart cards carried by the user), such solutions are not always feasible in many large-scale applications that require a central template database. The ideal solution for biometric template protection is to eliminate the need to store any biometric information in the database. This can be achieved by transforming the biometric trait into a pseudo-random key, which can be regenerated every time a new sample of the same biometric trait is presented (Uludag et al., 2004a). Note that to preserve the recognition accuracy, non-mate samples (from different subjects) must result in different keys. However, the above concept of biometric key generation represents the *holy grail* of biometrics because it requires a representation scheme that is invariant to intra-user variations, but at the same time unique to each user.

A more practical and feasible solution is to transform the raw biometric template into a "secure" template, which satisfies the following three requirements.

1. Non-invertibility: It must be computationally hard to recover the biometric features from the stored template. This prevents the adversary from replaying the biometric features gleaned from the template or creating physical spoofs of the biometric trait. Non-invertibility is quantified by $H(\mathbf{X}_E|\widehat{\mathbf{X}}_E)$, where $\widehat{\mathbf{X}}_E$ is the secure template generated from the raw template $\mathbf{X}_E$.
2. Non-linkability: It should be possible to create multiple secure templates from the same biometric data that are not linkable. This property not only enables the biometric system to revoke and re-issue new biometric templates when the template database is compromised, but also ensures that cross-matching across databases is not possible, thereby preserving the user's privacy. Non-linkability can be measured by $H(Y|\widehat{\mathbf{X}}_E^1, \widehat{\mathbf{X}}_E^2, \cdots)$, where $\widehat{\mathbf{X}}_E^1, \widehat{\mathbf{X}}_E^2, \cdots$ are multiple secure templates generated based on biometric

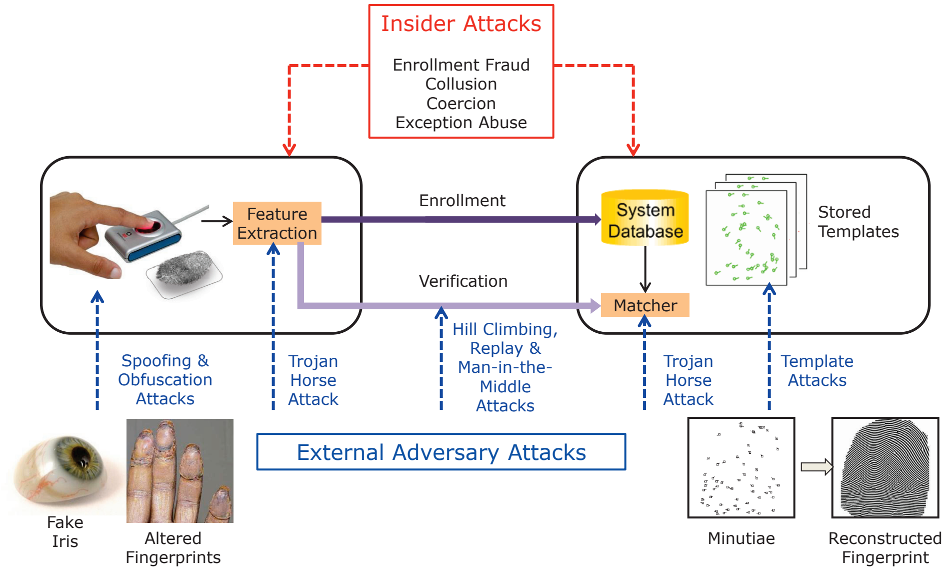

Fig. 20. A summary of possible attacks on a biometric system. While a majority of the security threats are common to any authentication system, the problem of spoofing (presentation of fake biometric traits) and attacks on the template database (either to reverse engineer the original biometric data or perform cross-matching) are specific to biometric systems.

features extracted from the same user Y.

3. Discriminability: The secure template should not degrade the recognition accuracy of the biometric system. In other words, $H(Y|\mathbf{X}_R, \widehat{\mathbf{X}}_E)$ should be as close as possible to $H(Y|\mathbf{X}_R, \mathbf{X}_E)$.

The main challenge in biometric template protection is to design a scheme that generates non-invertible and non-linkable templates without compromising on the matching accuracy. This boils down to designing a template protection scheme such that $H(\mathbf{X}_E|\widehat{\mathbf{X}}_E)$ and $H(Y|\widehat{\mathbf{X}}_E^1, \widehat{\mathbf{X}}_E^2, \cdots)$ are maximized, while minimizing $H(Y|\mathbf{X}_R, \widehat{\mathbf{X}}_E)$. While several approaches such as feature transformation (Ratha et al., 2001) and biometric cryptosystems (Dodis et al., 2006) have been proposed in the literature, the search for a secure biometric template satisfying all the three requirements has proved to be elusive thus far. The emergence of homomorphic encryption technology[20] appears to be promising.

Another issue that has gained considerable attention is the concept of public self-disclosures through online social networks. A large number of face photos are being posted online through social networks such as Facebook. Recent research has established the possibility of deducing potentially sensitive personal data by combining online social network data with off-the-shelf face recognition technology and cloud computing power.[21] In order to address this problem, techniques such as Visual Cryptography (Ross and Othman, 2011) and Privacy-preserving Photo Sharing (Ra et al., 2013) have been proposed. However, this continues to be an important area of research as individuals begin to share large amounts of biometric data (viz., face and voice data) through online social networks.

Finally, the ease with which face and voice data can be surreptitiously recorded using devices such as Google Glass has also raised privacy concerns.[22] The recorded data can potentially be used to deduce an individual's identity and personal information. In order to counter this possibility, researchers have developed 'Anti-Google Glass' technology - a pair of glasses outfitted with LEDs that emit near-infrared light into Google Glass cameras thereby preventing face recognition techniques from detecting a face (Yamada et al., 2013). The use of facial cosmetics has also been shown to degrade the accuracy of face recognition techniques (Dantcheva et al., 2012).

6. The Future of Biometric Recognition

While improvements in biometric algorithms (feature extraction, matching, and security fixes) will continue to play a major role in shaping the future of biometric recognition, it is also important to keep in mind that changes in enabling technologies and products will also have a significant influence on how biometric recognition systems will evolve in the future. For instance, exponential improvements in the performance and cost of processors and memory have already played a dominant role in the development of better biometric sensors. Similarly, advancements in the field of bioelectronics have created new products like lab-on-a-chip. This in turn has enabled rapid DNA analysis and opened up new frontiers for the use of DNA as a biometric identifier. Rapid improvements in communication

[20] http://www.fujitsu.com/global/news/pr/archives/month/2013/20130828-01.html

[21] `http://www.blackhat.com/docs/webcast/acquisti-face-BH-Webinar-2012-out.pdf`

[22] http://www.telegraph.co.uk/technology/google/10494231/The-places-where-Google-Glass-is-banned.html

technologies and proliferation of consumer electronic devices (e.g., smartphones) have also created new avenues for the deployment of biometrics. In applications such as device personalization (e.g., entertainment systems, automobiles), financial transactions (e.g., ATM machines, credit card purchase), facility access (e.g., fitness gyms, private apartments) and online social networks (e.g., messaging over FaceBook), people are likely to avail of biometric technology on a daily basis.

6.1. New Sensors & Computing Platforms

In 1965, Gordon Moore (Moore, 1965) predicted that the number of components (transistors) in an integrated circuit (IC) is likely to double approximately every two years for the next 10 years, while the unit cost of each component is likely to fall. Remarkably, this prediction (also called the Moore's Law) has held true for nearly half-a-century. In the case of biometric recognition, the direct impact of the rapid improvements in ICs is the development of smaller, cheaper, and higher quality biometric sensors as discussed earlier in Section 3. Improvement in sensors have mitigated the intra-subject variations caused to sensor limitations to a large extent. It is expected that the performance of integrated circuits will continue to improve at the same rates in the near future[23]. This will act as a catalyst for the development of novel sensors, which can be expected to further push the limits on quality, usability, and cost. Sensors that can capture multiple biometric traits of the user simultaneously (e.g., all biometric modalities in the human face or human hand) are also likely to developed. It is also necessary to develop a user-friendly ergonomic interface that can still permit the acquisition of repeatable biometric samples from a subject, i.e., reduce the variations caused due to user interactions.

As a direct consequence of this improvement in ICs, the performance of microprocessors has been doubling every two years, while the cost of computing is decreasing at the same rate (see Fig. 21(a)). A similar trend has also been occurring in the case of random access memory (RAM) (see Fig. 21(b)) and other storage devices. These exponential improvements in computing and storage have enabled the deployment of more powerful algorithms to process the captured biometric data. For instance, even though the concept of neural networks had been known for more than 30 years, the availability of powerful processors and the ability to efficiently handle large amounts of data has played a key role in the development of deep learning algorithms, which are powerful tools in many pattern recognition applications.

The availability of cloud computing has also presented new opportunities. Firstly, a cloud architecture can be used to store and access biometric data across different entities (e.g., organizations or countries) under differential policies (e.g., policies defining level of access and data usage). Secondly, a cloud framework can be used by clients to access biometric software development kits (e.g., face matcher) on a need-to-use basis or based on anticipated workload. In such a scenario, biometric recognition can be viewed as a service. Thirdly, cloud-

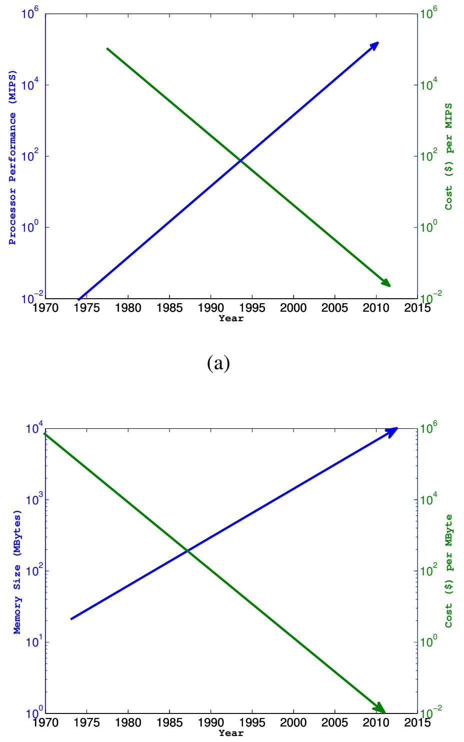

(a)

(b)

Fig. 21. Dramatic improvements in (a) processor performance and (b) random access memory (RAM) capacity due to the doubling of transistors in integrated circuits (ICs) every two years (Moore's Law) (Moravec, 1997; McCallum, 2013). While the processor performance (measured in millions of instructions per second (MIPS)) and RAM capacity (measured in Megabytes) have improved by more than six orders of magnitude (as indicated by the blue lines) over the last 40 years, the per unit costs of these components have been falling exponentially (as indicated by the green lines). These improvements have directly impacted the evolution of biometric sensors enabling the creation of smaller, cheaper, and higher quality biometric sensors.

[23] `http://www.itrs.net/Links/2012ITRS/Home2012.htm`

based biometrics can facilitate rapid analytics (e.g., recognizing a face using a smartphone camera, where the phone accesses the cloud) due to the availability of a large number of parallel nodes (i.e., computational/software resources). However, appropriately harnessing the power of cloud computing, while preserving the privacy and security of biometric data, remains an open-problem in the context of biometrics.

6.2. *Ubiquitous Biometrics*

The notion of ubiquitous biometrics may refer to the identification of an individual at any time and at any place by utilizing all the pieces of information - both biometric and non-biometric - available about the person. An ubiquitous biometric system will exploit other identity cues such as a person's location (de Montjoye et al., 2013), behavior, and recent interaction history (Rashid et al., 2013) in conjunction with the available biometric data (including soft biometric characteristics) to establish the person's identity with a high degree of reliability. This concept can be understood from the following illustrative example. Suppose that a user wishes to perform a banking transaction using his smartphone equipped with a face recognition system. Apart from capturing the person's face, the authentication system can also obtain information about the user's location using the Global Positioning System (GPS) sensors available on the phone. It is also possible to obtain information on the user's recent interactions with the phone (e.g., the applications that were accessed) as well as the transaction history of the user with the bank (e.g., transaction type, amount involved, beneficiaries). All these bits of evidence about the user can be integrated to obtain a strong assurance of identity.

An alternate perspective on ubiquitous biometrics is to venture beyond the task of establishing the identity of a person and gather additional information about the person. For instance, apart from identifying an individual using the face image, the system can also recognize the person's mood based on his facial expressions (Bettadapura, 2012). Moving one step further, it may be possible to find out the person's preferences and behavioral characteristics by mining his social media profile (Rashid et al., 2013). This additional knowledge about the user would be extremely useful in applications that require personalized delivery of services.

It must be emphasized that caution must be exercised when designing such ubiquitous biometric systems. Issues such as application context and user privacy concerns must be carefully assessed and appropriate checks and balances must be in place in order to prevent abuse of biometric recognition systems for unintended purposes. For example, buying a meal from a restaurant should not require the same level of identity assurance as in the case of performing a high-value financial transaction. Similarly, the issues related to the ownership of personal data and appropriate usage rights should be resolved before designing an ubiquitous biometric system that is capable of inferring the complete personality profile of a person.

6.3. *Biometrics For Social Good*

Biometric systems are being increasingly deployed in applications where societal benefits, and not security alone, is the dominant motivating factor. As discussed in Section 1, many national ID systems around the world are basically focused on giving the poor and illiterate people primarily in rural areas an identity, which will allow them to taste the benefits of social welfare schemes and health-care services provided by government and non-profit organizations. The rapid proliferation of mobile phones[24] has also played a major role in accelerating this trend. While mobile phones are generally considered only as a convenient tool for communication and entertainment, they are being increasingly used as a mechanism to deliver services and benefits to segments of the population, who were hitherto unreachable due to lack of physical infrastructure.

A good example of the usage of mobile phones and biometrics for delivering health-care services is the mobile-phone based vaccination registry developed by VaxTrac[25], which is used in African countries like Benin. The primary purpose of this registry is to keep track of the vaccine doses given to children so that redundant doses can be avoided, while simultaneously improving the immunization coverage. Since the children undergoing immunization seldom have proper identity documents and are often known only by their first name, it has been very difficult to keep track of vaccine doses administered to them. The VaxTrac system utilizes fingerprint biometrics to address this problem. An unsolved problem is how to acquire fingerprint images (or for that matter other biometric traits) of newborns and infants that are of sufficient quality for matching over a two-year time span (Jain et al., 2014). Similarly, the youngest age at which an infant's fingerprint can be successfully captured has not been established (Figure 22).

6.4. *Biometrics & Forensics*

Although forensics was one of the earliest applications of biometric recognition, biometric systems are yet to live up to their full potential in solving the problems faced by forensic experts. Biometric recognition can be used in forensics in two distinct ways: (i) as a tool to assist in investigation by identifying suspects and (ii) as an evidence in a court of law. It is worth noting that these two use-cases have very different requirements. In the first case, the key requirements are the speed and accuracy of biometric recognition under challenging imaging conditions. However, errors are tolerable to some extent in this scenario because the investigating officers can make use of other contextual information (e.g., demographic filtering) to eliminate some of the false matches. In the second scenario, the primary requirement is a convincing presentation of biometric evidence with strong statistical basis to the judge and the jury. This in turn involves obtaining a reliable estimate of the individuality of a biometric trait. Based on the discussion in Section 5, it should be clear that both the above requirements (recognition accuracy and individuality estimation) are not fully solved problems. Furthermore, Champod (Champod, 2013) argues that traditional performance metrics like False Match Rate (FMR) and False Non-Match Rate are not suitable for evidence

[24] It is estimated that over 1.7 billion mobile phones were sold worldwide in 2012 alone. Source: `http://www.gartner.com/newsroom/id/2335616`

[25] `http://vaxtrac.com/about`

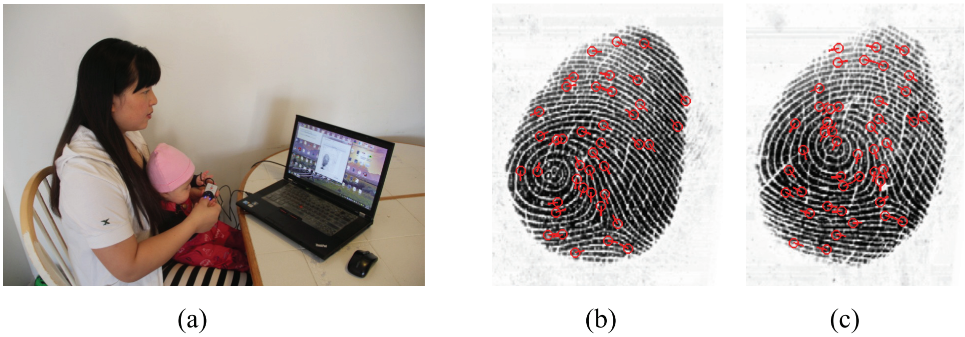

Fig. 22. Capturing an infant's fingerprints using an optical sensor. (a) Image acquisition setup. (b) and (c) Two fingerprint impressions of a four-month-old girl's left thumb; minutiae in these images were extracted using a commercial fingerprint SDK. The match score between these two fingerprints is 216 which indicates a high similarity (the threshold at FAR=0% for this matcher on FVC2002 DB1-A is 51). In this example, the infant's fingerprints were successfully acquired and matched. However, this may not be the case for all infants across different sensors, demographic and age groups.

presentation in a court of law and new metrics such as Rate of Misleading Evidence in favor of the Defense (RMED) and Rate of Misleading Evidence in favor of the Prosecution (RMEP) are needed to describe the performance of biometric systems in the forensics domain.

One of the interesting developments in the intersection of forensics and biometrics is the advancements in matching of DNA samples. The current standard procedures for DNA testing, namely Polymerase Chain Reaction (PCR) and Short Tandem Repeat (STR) analysis, have been in place for around two decades now (New England Innocence Project, 2011). Since these procedures typically involve laboratory analysis by human operators, it may take several hours to several days to obtain an STR profile from a buccal swab. However, prototype devices are now available for rapid DNA analysis (The Federal Bureau of Investigation, 2013b). These devices fully automate the process of developing a STR profile from a reference buccal swab and have a response time of less than two hours. In the near future, it may be possible to further speed up this process to a few minutes, thereby making DNA as a feasible biometric modality even in applications other than forensics. However, one needs to be extremely cautious about the privacy issues associated with DNA-based biometric systems because the DNA samples (or templates) may contain a wealth of personal information (e.g., susceptibility to diseases).

The use of forensic evidence in U.S. Federal courts (and in several State courts) is guided by the Federal Rules of Evidence. In particular, Rule 702 states that testimony provided by an expert witness must be "based on sufficient facts or data. The Daubert standard, which largely stemmed from Rule 702, further defined the criteria for the admissibility of scientific evidence. In Daubert v. Merrell Dow Pharmaceuticals, 509 U.S. 579, the Court ruled that the validity of scientific testimony has to satisfy relevancy and reliability standards, i.e., the experts testimony should be "relevant to the task at hand and should rest "on a reliable foundation. Carefully answering the related questions will play a critical role not only in legal proceedings, but also in bolstering the scientific basis for biometric methods used in forensic investigations (e.g., latent fingerprint matching). In particular, it will be the first step in assuaging criticism leveled by the 2009 National Academy of Sciences' report, *Strengthening Forensic Science in the United States: A Path Forward*, which concluded that claims about the evidential value of forensic data are not supported by rigorous scientific study.

7. Summary

To counter growing security threats and financial fraud, and to facilitate personalization and convenience, the importance of biometrics as a reliable tool for person recognition has been established beyond doubt. It is indeed fascinating that a system can recognize a person with extremely high accuracy within a fraction of a second based on the friction ridge pattern on the tip of his finger, or the textural patterns on the stroma of his iris, using a commodity processor such as a laptop or a mobile phone. This is a significant achievement given that the first paper in automated biometric recognition was published only 50 years ago.

In this paper, we have attempted to summarize the state of the art in biometrics recognition and have identified key challenges that deserve attention. The biometrics community has indeed come a long way over the past 50 years. On one hand, tremendous progress has been made in designing large-scale biometric systems that can rapidly search through biometric databases in order to retrieve a matching identity (e.g., the IrisGuard system deployed in the United Arab Emirates). On the other hand, the advent of smartphones and other consumer devices has led to enhanced interest in designing biometric solutions for resource-constrained devices (e.g., the Touch ID fingerprint system in iPhones). Modern biometric systems are being increasingly tuned to deal with poor quality data, including those encountered in traditional forensics applications. These advancements have been facilitated by attendant progress in computing power, signal processing, computer vision, pattern recognition and machine learning. Despite the challenges that remain, the biometrics community can celebrate its accomplishments over the past 50 years. The technology has indeed redefined the landscape of personal authentication. In order to take biometrics technology to the next level, so that it is pervasive (a la the movie *Minority Report*), biometric researchers need to be aware of the application requirements whilst not ignoring the algorithmic and privacy models necessary to reliably extract and match traits.

Acknowledgment

We are indebted to Jim Wayman for providing historical time-lines and contributors to biometrics literature as well as correct terminology. We would also like to thank Sunpreet Arora, Vincent Bouatou, Kai Cao, John Daugman, Jianjiang Feng, Brendan Klare, Eryun Liu, Xiaoming Liu, Abhishek Nagar, Asem Othman, Sharath Pankanti, Salil Prabhakar, Rob Rowe, Elham Tabassi, Andrew Teoh, Kar-Ann Toh and Soweon Yoon for their careful reading of an earlier draft of this manuscript and providing their valuable suggestions.

References

Abaza, A., Ross, A., Herbert, C., Harrison, M.A.F., Nixon, M., 2013. A survey on ear biometrics. ACM Computing Surveys 45.

Antonelli, A., Cappelli, R., Maio, D., Maltoni, D., 2006. Fake Finger Detection by Skin Distortion Analysis. IEEE Transactions on Information Forensics and Security 1, 360–373.

Ashbaugh, D.R., 1999. Quantitative-Qualitative Friction Ridge Analysis: An Introduction to Basic and Advanced Ridgeology. CRC Press.

Babler, W.J., 1991. Embryologic development of epidermal ridges and their configurations. Birth Defects Original Article Series 27, 95–112.

Baker, S., Bowyer, K.W., Flynn, P.J., 2009. Empirical Evidence for Correct Iris match Score Degradation With Increased Time-Lapse Between Gallery and Probe Matches, in: Proceedings of IEEE International Conf. on Biometrics, Alghero, Italy. pp. 1170–1179.

Barrett, D., 2013. One surveillance camera for every 11 people in britain, says cctv survey. The Telegraph. URL: `http://www.telegraph.co.uk/technology/10172298/One-surveillance-camera-for-every-11-people-in-Britain-says-CCTV-survey.html`.

Belhumeur, P.N., Hespanha, J.P., Kriegman, D.J., 1997. Eigenfaces vs. fisherfaces: recognition using class specific linear projection. IEEE Transactions on Pattern Analysis and Machine Intelligence 19, 711–720.

Bertillon, A., 1896. Signaletic Instructions including the Theory and Practice of Anthropometrical Identification, R.W. McClaughry Translation. The Werner Company.

Bettadapura, V., 2012. Face expression recognition and analysis: the state of the art. arXiv preprint arXiv:1203.6722 .

Beveridge, J.R., Givens, G.H., Phillips, P.J., Draper, B.A., 2009. Factors that influence algorithm performance in the face recognition grand challenge. Computer Vision and Image Understanding 113, 750–762. URL: `http://dx.doi.org/10.1016/j.cviu.2008.12.007`, doi:10.1016/j.cviu.2008.12.007.

Blanz, V., Vetter, T., 2003. Face recognition based on fitting a 3d morphable model. IEEE Transactions on Pattern Analysis and Machine Intelligence 25(9), 1063–1074.

Bledsoe, W.W., 1966. Man-machine Facial Recognition. Technical Report PRI 22. Panoramic Research, Inc.

Burr, W.E., Dodson, D.F., Polk, W.T., 2006. Information Security: Electronic Authentication Guideline. Technical Report Special Report 800-63. NIST.

Champod, C., 2013. Introducing a LR-based identification system in forensic practice: opportunities and challenges, in: Proceedings of International Workshop on Biometrics and Forensics (IWBF), Lisbon, Portugal.

Chingovska, I., Anjos, A., Marcel, S., 2012. On the effectiveness of local binary patterns in face anti-spoofing, in: IEEE BIOSIG 2012.

Crowley, J.L., 2004. The FG-NET Aging Database. URL: `http://www-prima.inrialpes.fr/FGnet/`.

Dantcheva, A., Chen, C., Ross, A., 2012. Can facial cosmetics affect the matching accuracy of face recognition systems?, in: Proc. of 5th IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS), pp. 1–8.

Daugman, J., 2006. Probing the uniqueness and randomness of IrisCodes: Results from 200 billion iris pair comparisons. Proceedings of the IEEE 94, 1927–1935.

Daugman, J.G., 1993. High confidence visual recognition of persons by a test of statistical independence. IEEE Transactions on Pattern Analysis and Machine Intelligence 15, 1148–1160.

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

Department of Homeland Security, 2013. Office of Biometric Identity Management. URL: `http://www.dhs.gov/obim`.

Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A., 2006. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. Technical Report 235. Cryptology ePrint Archive. A preliminary version of this work appeared in EUROCRYPT 2004.

Duta, N., 2009. A survey of biometric technology based on hand shape. Pattern Recognition 42, 2797–2806.

Ernst, R.H., 1971. Hand ID System. United States patent number US 3576537.

Flom, L., Safir, A., 1987. Iris recognition system. United States patent number US 4641349.

Fourre, J.Y., Picard, S., Rieul, F., Fondeur, J.C., 2011. Device for acquiring images of patterns formed by furrows in the skin of fingers or the palm of the hand. United States patent number US 7912250.

Fu, Y., Guo, G., Huang, T.S., 2010. Age Synthesis and Estimation via Faces: A Survey. IEEE Transactions on Pattern Analysis and Machine Intelligence 32, 1955–1976.

Galton, F., 1892. Finger Prints. McMillan & Co, London, UK.

Google, Inc., 2013. Google Glass. URL: `http://www.google.com/glass/start/`.

Greenberg, C., Stanford, V., Martin, A., Yadagiri, M., Doddington, G., 2012. The 2012 NIST Speaker Recognition Evaluation. URL: `http://www.nist.gov/itl/iad/mig/sre12results.cfm`.

Grother, P., Matey, J.R., Tabassi, E., Quinn, G.W., Chumakov, M., 2013. IREX VI: Temporal Stability of Iris Recognition Accuracy. Technical Report NISTIR 7948. NIST.

Grother, P., Quinn, G.W., Matey, J.R., Ngan, M., Salamon, W., Fiumara, G., Watson, C., 2012. IREX III: Performance of Iris Identification Algorithms. Technical Report NISTIR 7836. NIST.

Hager, J.C., 2013. Aging of the Face. URL: `http://www.face-and-emotion.com/dataface/facets/aging.jsp`.

Hawthorne, M.R., 2009. Fingerprints: Analysis and Understanding. CRC Press.

Innocence Project, 2013. DNA Exonerations Nationwide. URL: `http://www.innocenceproject.org/Content/DNA_Exonerations_Nationwide.php`.

Jain, A.K., Cao, K., Arora, S.S., 2014. Recognizing infants and toddlers using fingerprints: Increasing the vaccination coverage, in: Proc. of International Joint Conference on Biometrics.

Jain, A.K., Nandakumar, K., Nagar, A., 2008. Biometric Template Security. EURASIP Journal on Advances in Signal Processing , 1–17.

Jain, A.K., Prabhakar, S., Pankanti, S., 2002. On The Similarity of Identical Twin Fingerprints. Pattern Recognition 35, 2653–2663.

Jain, A.K., Ross, A., Nandakumar, K., 2011. Introduction to Biometrics. Springer.

Jain, A.K., Ross, A., Pankanti, S., 1999. A prototype hand geometry-based verification system, in: 2nd International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA), Washington D.C.. pp. 166–171.

Kanade, T., 1973. Picture Processing System by Computer Complex and Recognition of Human faces. Ph.D. thesis. Kyoto University.

Khoshelham, K., Elberink, S.O., 2012. Accuracy and resolution of kinect depth data for indoor mapping applications. Sensors 12, 1437–1454.

Klare, B., Burge, M., Klontz, J., Vorder Bruegge, R., Jain, A., 2012. Face recognition performance: Role of demographic information. IEEE Transactions on Information Forensics and Security 7, 1789–1801. doi:10.1109/TIFS.2012.2214212.

Klontz, J.C., Jain, A.K., 2013. A Case Study on Unconstrained Facial Recognition Using the Boston Marathon Bombings Suspects. Technical Report MSU-CSE-13-4. Michigan State University.

Kong, A., Zhang, D., Kamel, M., 2009. A survey of palmprint recognition. Pattern Recognition 42, 1408–1418.

Lee, E.C., Park, K.R., Kim, J., 2006. Fake Iris Detection by Using Purkinje Image, in: Proceedings of International Conference on Biometrics, Hong Kong, China. pp. 397–403.

Li, J., Wang, Y., Tan, T., Jain, A., 2004. Live Face Detection Based on the Analysis of Fourier Spectra, in: Proceedings of SPIE Conference on Biometric

Technology for Human Identification, Orlando, USA. pp. 296–303.

Lowe, D., 2004. Distinctive image features from scale-invariant key points. International Journal of Computer Vision 60, 91–110.

Lui, Y.M., Bolme, D., Draper, B., Beveridge, J., Givens, G., Phillips, P., 2009. A meta-analysis of face recognition covariates, in: IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS), pp. 1–8.

M. Indovina *et al.*, 2009. ELFT Phase II - An Evaluation of Automated Latent Fingerprint Identification Technologies. Technical Report NISTIR 7577. NIST. URL: http://fingerprint.nist.gov/latent/NISTIR_7577_ELFT_PhaseII.pdf.

Maltoni, D., Maio, D., Jain, A.K., Prabhakar, S., 2009. Handbook of Fingerprint Recognition (2nd Edition). Springer Verlag.

Marasco, E., Ross, A., 2015. A survey on antispoofing schemes for fingerprint recognition systems. ACM Computing Survey 47.

Marcel, S., Nixon, M., Li, S.Z., 2014. Handbook of Biometric Anti-Spoofing. Springer.

Matsumoto, T., Matsumoto, H., Yamada, K., Hoshino, S., 2002. Impact of Artificial Gummy Fingers on Fingerprint Systems, in: Optical Security and Counterfeit Deterrence Techniques IV, Proceedings of SPIE, San Jose, USA. pp. 275–289.

Mauceri, A.J., 1965. Feasibility Study of Personal Identification by Signature Verification. Technical Report SID65-24. North American Aviation.

McCallum, J.C., 2013. Memory Prices (1957-2013). URL: http://www.jcmit.com/memoryprice.htm.

de Montjoye, Y.A., Hidalgo, C.A., Verleysen, M., Blondel, V.D., 2013. Unique in the Crowd: The privacy bounds of human mobility. Scientific Reports 3.

Moore, G.E., 1965. Cramming more components onto integrated circuits. Electronics 38.

Moravec, H., 1997. List of Microprocessors. URL: http://www.frc.ri.cmu.edu/~hpm/book97/ch3/processor.list.txt.

National Institute of Standards and Technology - Information Technology Laboratory, 2010. NIST Special Database 32 - Multiple Encounter Dataset (MEDS). URL: http://www.nist.gov/itl/iad/ig/sd32.cfm.

Neumann, C., Champod, C., Puch-Solis, R., Egli, N., Anthonioz, A., Bromage-Griffiths, A., 2007. Computation of likelihood ratios in fingerprint identification for configurations of any number of minutiae. Journal of Forensic Sciences 52, 54–63.

New England Innocence Project, 2011. A Brief History of DNA Testing. URL: http://www.newenglandinnocence.org/knowledge-center/resources/dna/.

Nixon, K.A., Rowe, R.K., 2005. Multispectral Fingerprint Imaging for Spoof Detection, in: Proceedings of SPIE Conference on Biometric Technology for Human Identification, Orlando, USA. pp. 214–225.

Nixon, M., Tan, T., Chellappa, R., 2006. Human Identification Based on Gait. Springer.

O'Gorman, L., 2003. Comparing Passwords, Tokens, and Biometrics for User Authentication. Proceedings of the IEEE 91, 2019–2040.

Ojala, T., Pietikainen, M., Maenpaa, T., 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence 24, 971–987.

Park, U., Jillela, R., Ross, A., Jain, A., 2011. Periocular biometrics in the visible spectrum. Information Forensics and Security, IEEE Transactions on 6, 96–106.

Park, U., Tong, Y., Jain, A.K., 2010. Age Invariant Face Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence 32, 947–954.

Parthasaradhi, S., Derakhshani, R., Hornak, L.A., Schuckers, S.A.C., 2005. Time-Series Detection of Perspiration as a Liveness Test in Fingerprint Devices. IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews 35, 335–343.

Parziale, G., Diaz-Santana, E., 2006. The Surround Imager: A Multi-Camera Touchless Device to Acquire 3D Rolled-Equivalent Fingerprints, in: Proceedings of International Conference on Biometrics, pp. 244–250.

Penev, P.S., Atick, J.J., 1996. Local feature analysis: A general statistical theory for object representation. Network: Computation in Neural Systems 7, 477–500.

Planning Commission, Government of India, 2012. Unique Identification Authority of India. URL: http://uidai.gov.in.

Pruzansky, S., 1963. Pattern-Matching Procedure for Automatic Talker Recognition. Journal of the Acoustic Society of America 35, 354–358.

Ra, M.R., Govindan, R., Ortega, A., 2013. P3: Privacy-preserving photo sharing, in: Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation.

Ramanathan, N., Chellappa, R., Biswas, S., 2009. Age progression in human faces: A survey. Journal of Visual Languages and Computing .

Rashid, A., Baron, A., Rayson, P., May-Chahal, C., Greenwood, P., Walkerdine, J., 2013. Who Am I? Analyzing Digital Personas in Cybercrime Investigations. Computer 46, 54–61.

Ratha, N.K., Connell, J.H., Bolle, R.M., 2001. Enhancing Security and Privacy in Biometric Authentication Systems. IBM Systems Journal 40, 614–634.

Ricanek, K., Tesafaye, T., 2006. MORPH Face Database. URL: http://faceaginggroup.com/.

Roberts, C., 2007. Biometric Attack Vectors and Defences. Computers and Security 26, 14–25.

Roberts, L.G., 1963. Machine perception of three-dimensional solids. Ph.D. thesis. Massachusetts Institute of Technology.

Ross, A., Nandakumar, K., Jain, A.K., 2006. Handbook of Multibiometrics. Springer.

Ross, A., Othman, A., 2011. Visual cryptography for biometric privacy. IEEE Transactions on Information Forensics and Security 6, 70 – 81.

Sirovich, L., Kirby, M., 1987. Low-dimensional procedure for the characterization of human faces. Journal of the Optical Society of America 4, 519–524.

Skorka, O., Joseph, D., 2011. Toward a digital camera to rival the human eye. Journal of Electronic Imaging 20, 1–18.

Spearman, E., 1999. Crime and Punishment in England: A Sourcebook. Routledge, London, UK. chapter Identifying Suspects (1894). pp. 256–257.

SRI International, 2013. Iris on the Move Biometric Identification Systems. URL: http://www.sri.com/engage/products-solutions/iris-move-biometric-identification-systems.

Sun, Y., Wang, X., Tang, X., 2014. Deep Learning Face Representation from Predicting 10,000 Classes, in: Proc. of IEEE Conference on Computer Vision and Pattern Recognition, pp. 1891–1898.

Suzuki, T., 2010. Challenges of Image-Sensor Development, in: Proceedings of International Solid-State Circuits Conference, pp. 37–30.

Taigman, Y., Yang, M., Ranzato, M., Wolf, L., 2014. DeepFace: Closing the Gap to Human-level Performance in Face Verification, in: Proc. of IEEE Conference on Computer Vision and Pattern Recognition, pp. 1701–1708.

The Federal Bureau of Investigation, 2013a. CODISNDIS Statistics. URL: http://www.fbi.gov/about-us/lab/biometric-analysis/codis/ndis-statistics.

The Federal Bureau of Investigation, 2013b. Frequently Asked Questions (FAQs) on the CODIS Program and the National DNA Index System. URL: http://www.fbi.gov/about-us/lab/biometric-analysis/codis/codis-and-ndis-fact-sheet.

Trauring, M., 1963. Automatic Comparison of Finger Ridge Patterns. Nature 197, 938–940.

Turk, M., Pentland, A., 1991. Eigenfaces for recognition. Cognitive Neuroscience 3, 72–86.

Uludag, U., Pankanti, S., Prabhakar, S., Jain, A.K., 2004a. Biometric Cryptosystems: Issues and Challenges. Proceedings of the IEEE, Special Issue on Multimedia Security for Digital Rights Management 92, 948–960.

Uludag, U., Ross, A., Jain, A.K., 2004b. Biometric Template Selection and Update: A Case Study in Fingerprints. Pattern Recognition 37, 1533–1542.

University of Bologna, 2006. FVC2006: The Fourth International Fingerprint Verification Competition. URL: http://bias.csr.unibo.it/fvc2006/.

Viola, P.A., Jones, M.J., 2004. Robust real-time face detection. International Journal of Computer Vision 57, 137–154.

Wayman, J.L., 2007. The History of Information Security: A Comprehensive Handbook. Elsevier, Amsterdam. chapter The Scientific Development of Biometrics Over the Last 40 Years.

Wikipedia, 2002. Twin. URL: http://en.wikipedia.org/wiki/Twin.

Wikipedia, 2013. Boston Marathon bombings. URL: http://en.wikipedia.org/wiki/Boston_Marathon_bombings.

Wilson, C.L., 2004. Fingerprint Vendor Technology Evaluation 2003: Summary of Results and Analysis Report. Technical Report NISTIR 7123. NIST.

Wiskott, L., Fellous, J.M., Kuiger, N., von der Malsburg, C., 1997. Face recognition by elastic bunch graph matching. IEEE Transactions on Pattern Analysis and Machine Intelligence 19, 775–779.

Wright, J., Yang, A., Ganesh, A., Sastry, S., Ma, Y., 2009. Robust Face Recognition via Sparse Representation. IEEE Transactions on Pattern Analysis and Machine Intelligence 31, 210–227.

Xia, X., O'Gorman, L., 2003. Innovations in Fingerprint Capture Devices. Pattern Recognition 36, 361–369.

Yamada, T., Gohshi, S., Echizen, I., 2013. Privacy visor: Method for preventing face image detection by using differences in human and device sensitivity, in: Proc. of the 14th Joint IFIP TC6 and TC11 Conference on Communications and Multimedia Security (CMS 2013), pp. 1–10.

Yoon, S., Feng, J., Jain, A.K., 2012. Altered Fingerprints: Analysis and Detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 34, 451–464.

Yoon, S., Jain, A.K., 2015. Longitudinal study of fingerprint recognition. Proc. National Academy of Sciences (PNAS) .

Zhang, Z., Yan, J., Liu, S., Lei, Z., Yi, D., Li, S.Z., 2012. A face antispoofing database with diverse attacks, in: Proceedings of 5th IAPR International Conference on Biometrics, pp. 26–31.

Zhu, Y., Dass, S.C., Jain, 2007. Statistical Models for Assessing the Individuality of Fingerprints. IEEE Transactions on Information Forensics and Security 2, 391–401.

Appendix: Image Sources

The following illustrations in this paper have been generated using images downloaded from the Internet. The corresponding image links are listed below.

- Figure 4
  - http://upload.wikimedia.org/wikipedia/commons/7/74/Bertillon_-_Signalement_Anthropometrique.png
  - http://nutrias.org/~nopl/monthly/sept2002/bcdhowell.jpg
- Figure 7
  - http://galton.org/fingerprints/images/faulds-1920-08-27-thumb.gif
  - http://upload.wikimedia.org/wikipedia/commons/e/ec/Francis_Galton_1850s.jpg
  - http://criminaljustice.state.ny.us/ojis/history/images/henry.jpg
  - http://www.thehindu.com/multimedia/dynamic/00835/IN13_AADHAR_CARD_835657f.jpg
- Figure 10
  - http://groups.csail.mit.edu/medg/people/doyle/gallery/bledsoe/bledsoe.gif
  - http://photodoto.com/wp-content/uploads/2011/10/history-photo-camera-9.jpg
  - http://photodoto.com/wp-content/uploads/2011/10/history-photo-camera-14.jpg
  - http://commons.wikimedia.org/wiki/File:Three_Surveillance_cameras.jpg
  - http://siliconcowboy.files.wordpress.com/2010/11/jphone.jpg
  - http://upload.wikimedia.org/wikipedia/commons/6/67/Xbox-360-Kinect-Standalone.png
  - http://static7.businessinsider.com/image/4d013ea7cadcbb7033010000/looxcie-video-camera.jpg
  - http://cdn.techinasia.com/wp-content/uploads/2013/03/samsung-galaxy-s4-white.jpg
- Figure 11
  - http://www.icb12.iiitd.ac.in/images/daugman.jpg
  - http://upload.wikimedia.org/wikipedia/commons/1/1b/IriScan_model_2100_iris_scanner_1.jpg
  - https://www.fbo.gov/index?s=opportunity&mode=form&tab=core&id=1d31acd78d7a20b1fe598bf4b4661d6b
  - http://www.sri.com/engage/products-solutions/iom-passport-portal-system
  - http://i2.cdn.turner.com/cnn/dam/assets/130411103216-biometric-scanning-iphone-tool-horizontal-gallery.jpg
  - http://www.geekalerts.com/u/fingerprint-mouse.jpg