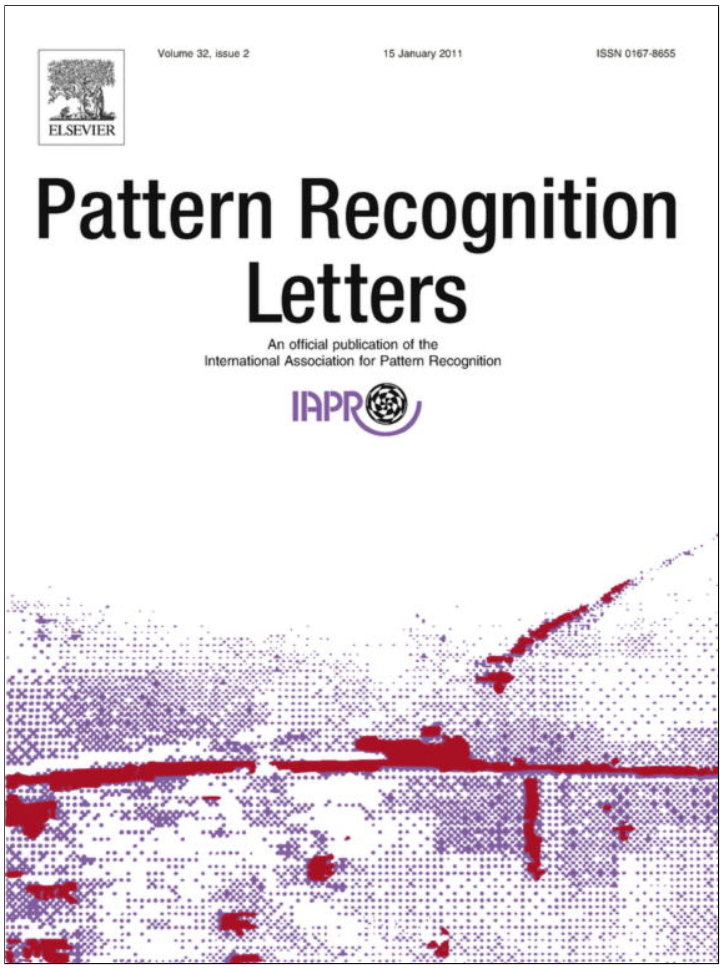

(This is a sample cover image for this issue. The actual cover is not yet available at this time.)

Pattern Recognition Letters 32 (2011) 1643–1651

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

An evaluation of indirect attacks and countermeasures in fingerprint verification systems

Marcos Martinez-Diaz *, Julian Fierrez, Javier Galbally, Javier Ortega-Garcia

Biometric Recognition Group – ATVS, Escuela Politecnica Superior, Universidad Autonoma de Madrid, Campus de Cantoblanco, C/Francisco Tomas y Valiente, 11 28049 Madrid, Spain

ARTICLE INFO

Article history:
Received 28 January 2009
Available online 8 May 2011
Communicated by S. Sarkar

Keywords:
Biometrics
Indirect attacks
Fingerprint verification
Match-on-Card
Vulnerabilities
Hill climbing

ABSTRACT

Biometric recognition systems are vulnerable to numerous security threats. These include direct attacks to the sensor or indirect attacks, which represent the ones aimed towards internal system modules. In this work, indirect attacks against fingerprint verification systems are analyzed in order to better understand how harmful they can be. Software attacks via hill climbing algorithms are implemented and their success rate is studied under different conditions. In a hill climbing attack, a randomly generated synthetic template is presented to the matcher, and is iteratively modified based on the score output until it is accepted as genuine. Countermeasures against such attacks are reviewed and analyzed, focusing on score quantization as a case study. It is found that hill climbing attacks are highly effective in the process of creating synthetic templates that are accepted by the matcher as genuine ones. We also find that score quantization drastically reduces the attack success rate. We analyze the hill climbing approach over two state-of-the-art fingerprint verification systems: the NIST Fingerprint Image Software 2, running on a PC and a prototype system fully embedded in a smart card (Match-on-Card). Results of both systems are obtained using a sub corpus of the publicly available MCYT database.

1. Introduction

Biometric systems are becoming popular in our society, as they provide a convenient method for user authentication and access to secure environments (Jain et al., 2006). The heightened interest in biometrics-based automated personal identification has resulted in the development of several commercial biometric recognition systems. One of their main advantages is that they free the user from passwords that may be stolen or forgotten. Moreover, many of them can be easily embedded in electronic devices (e.g. fingerprint, face or voice recognition systems).

Within biometrics, fingerprints are one of the most commonly used traits due to their widely accepted distinctiveness and the traditional use in forensic environments (Maltoni and Cappelli, 2007). Despite the currently established state of the art in automatic fingerprint recognition techniques, which have reached verification performances adequate for a number of applications, there are still many security concerns which must be faced by system designers (Jain et al., 2008). Biometric template protection techniques try to ensure that the user biometric data stored in a system database cannot be accessed by an eventual attacker. The need for biometric template protection is increased by the fact that biometric traits cannot usually be replaced or changed once stolen.

In the past few years, a considerable effort has been carried out in analyzing, classifying and solving the possible security breaches that biometric verification systems may present. Ratha et al. (2001) have studied and systematically categorized in eight classes the possible types of attacks. These have been depicted in Fig. 1. The attacks can be grouped in *direct* and *indirect* attacks as follows:

Direct attacks correspond to type 1 in Fig. 1. They are aimed directly towards the biometric sensor. Their goal is to present a fake biometric template to the sensor, trying to impersonate a real user. For the case of fingerprints, these can be performed using fake fingerprints, e.g. gummy fingerprints (Galbally et al., 2010a). This kind of attacks don't require any knowledge about the biometric system (e.g. matching algorithm, feature extraction). Countermeasures against such threats are primarily based on liveness detection, and include skin distortion models and skin odor sensors (chemical) among others (Franco and Maltoni, 2007).

Indirect attacks include the rest of types reported by Ratha et al. (2001). Attacks of type 2, 4, 7 and 8 are oriented against communication channels inside the biometric recognition system. On the other hand,

* Corresponding author. Fax: +34 914972235.
E-mail addresses: marcos.martinez@uam.es (M. Martinez-Diaz), julian.fierrez@uam.es (J. Fierrez), javier.galbally@uam.es (J. Galbally), javier.ortega@uam.es (J. Ortega-Garcia).

0167-8655/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2011.04.005

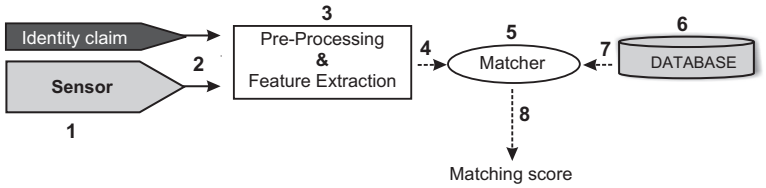

Fig. 1. General architecture of an automatic biometric verification system. Possible attack points are numbered from 1 to 8.

attacks on points 3 and 5 try to bypass or replace the feature extractor and the matcher respectively. Finally, type 6 attacks are aimed directly against the template database. These might try to add, modify or delete user information. Contrary to the case of direct attacks, the attacker must know specific information about the system in order to perform indirect attacks (e.g. template format, communications protocol, etc.), although reverse engineering techniques may also be employed. Moreover, the attacker needs physical or logical access to internal parts of the system, not available to the common user. Countermeasures in this case are usually based on physical and logical security, such as failover systems or encryption. These are described in more detail in Section 4.

Maltoni et al. (2003) have furthermore listed the threats that apply to a generic security system, including those based on automatic fingerprint verification. These threats can be summarized as follows: *Denial of Service* (DoS) attacks try to impede access to the system for legitimate users, *circumvention* refers to accessing the system without authorization, including *contamination or covert acquisition*, where the attacker acquires and uses biometric information from the user (e.g. from latent fingerprints). *Repudiation* considers the cases where a malicious user deliberately denies having accessed the system. Finally in *collusion* attacks the unauthorized user is helped to gain access by another user or system administrator while in *coercion* attacks a legitimate user is forced to help the attacker. The indirect attacks considered in this work, can be generally considered as circumvention attacks.

Most works that have studied indirect attacks against biometric systems in the literature have considered approaches related to hill climbing attacks (Uludag and Jain, 2004; Adler, 2003; Galbally et al., 2007). The term "hill climbing" refers to those attacks in which the similarity score given by the matcher is used to iteratively modify and present to the matcher a synthetically generated biometric template, or group of templates, until the verification threshold is reached.

Match-on-Devices represent a hot topic in biometrics, of which a representative example is Match-on-Card for fingerprint recognition. In Match-on-Card systems, the user information, fingerprint template and matching algorithm are stored in a smart card. Smart cards have integrated circuits or microprocessors that may allow the encryption and protection of stored information and the execution of moderately complex algorithms (Sanchez-Reillo et al., 2003). They allow users to easily carry with them a full biometric verification system. Corroborating the increasing interest in Match-on-Card systems, in the Fingerprint Verification Competition (FVC) 2004 (Cappelli et al., 2006), a special evaluation track was introduced for the case of matching systems with reduced memory and time restrictions. In the 2006 competition, (FVC, 2006), the need for introducing new specific Match-on-Card and Match-on-Device categories was stated. Furthermore, the National Institute of Standards and Technology (NIST) is currently performing the Minutiae Interoperability Exchange (MINEX) II public evaluation of Match-on-Card systems (Grother et al., 2008). The objective of this evaluation is to certify fingerprint Match-on-Card algorithms, required by the US government Personal Identity Verification program for the identification and authentication of Federal employees and contractors. The common approach in these and other related benchmarks in fingerprint recognition (Wilson et al., 2004) is to evaluate competing systems with regard to the verification error rates and other performance measures. While verification performance is key when evaluating security system, we stress the importance of also evaluating the robustness of fingerprint systems against possible attacks.

In this work, we study the feasibility of indirect attacks, and in particular hill climbing attacks, towards two state-of-the-art fingerprint verification systems and assess the performance of proposed countermeasures. The indirect attacks implemented for the evaluation are known as hill climbing attacks (Uludag and Jain, 2004), and are directed to the input of the matcher (point 4 of attack in Fig. 1). The attacks are implemented on both the NIST minutia-based system and a Match-on-Card (MoC) system. In the case of the NIST system, we have full knowledge of the matching algorithm, while on the other hand, only the input and output format is known for the MoC system, since it is a prototype acquired from a vendor.

Our main contribution is the analysis of this type of attacks in real operating conditions, using state-of-the-art fingerprint verification systems and a publicly available fingerprint database. The main research result is therefore providing new insight regarding to what extent are vulnerable fingerprint recognition systems to hill-climbing attacks. Moreover, the effects of countermeasures based on score quantization on the attack success rate and on the verification performance is systematically evaluated for the first time in the literature to the extent of our knowledge.

The paper is structured as follows. Related works are summarized in Section 2. The implemented attack algorithm is described in Section 3. In Section 4 possible countermeasures against indirect attack are summarized. The evaluated systems, experiments and results are presented in Section 5. Conclusions are finally drawn in Section 6.

2. Related works

As has been stated in Section 1, the majority of the works regarding indirect attacks use some type of variant of the hill climbing technique. However, two notable exceptions can be found in the literature. Hill (2001) presents an attack to a biometric system database (type 6 attack in Fig. 1). The templates stored in the database are reverse engineered to obtain a synthetic image of the fingerprint, which is presented to the matcher. Mohanty et al. (2007) propose a type 4 attack where face templates are reconstructed using score information by affine transforms. In that work,

face templates are reconstructed using score information. Their algorithm attempts to model the verification algorithm by a set of different samples to the matcher and modeling the output scores using by affine transformations.

Most hill climbing algorithms studied in the past few years are based on the technique introduced by Soutar et al. (1999). In that preliminary work, a basic hill climbing attack was tested over a simple image recognition system using filter-based correlation. This attack takes advantage of the score given by the matcher to iteratively change a synthetically created template until the score exceeds a fixed decision threshold and the access to the system is granted (this matches the hill climbing attack definition given in Section 1). Thus, depending on whether a synthetic image file is created or the synthetic feature vector is directly generated, these attacks can belong to type 2 or 4, respectively.

When the hill climbing attack is directed to the input of the feature extractor (type 2 attack), no information about the template storage format is required. Only the size and file format presented to the feature extractor is needed. Adler (2003) proposed a type 2 attack with a face recognition system. The input image is iteratively modified until a desired matching score is attained. This work reported results on three commercial recognition systems and showed that after 4,000 iterations, a score corresponding to a very high similarity confidence (99.9%) was reached for all systems tested. Hill climbing attacks applied to signature verification have also been studied by Yamazaki et al. (2005) and Muramatsu (2008). Galbally et al. (2010b) have also analyzed the vulnerability of face verification systems to this type of attacks. Recently, a general hill climbing approach based on Bayesian adaptation was presented by Galbally et al. (2007), testing it against a signature verification system. This algorithm, which can be used against any system working with fixed-length templates, was tested over a signature verification system reaching a success rate over 95%.

The hill climbing attacks performed in the present work are based on those described by Uludag and Jain (2004). In these attacks a synthetic random minutiae template is presented to the input of the matcher (type 4 attack) and, according to the score generated, the random template is iteratively changed until the system returns a positive verification. The minutiae in the template are modified one at a time and the change is only stored if the score returned by the matcher improves the previous one, otherwise it is discarded. Thus, to carry out this type of attack, the following information is needed: (i) the resolution and size of the images captured by the sensor (which is usually a known parameter specified by the vendor), (ii) the template format, and (iii) access to the matcher input (to present the synthetic templates) and output (to get the necessary feedback from the scores). Consequently, in this case we need to know *how* the information is stored, but not *what* the information is.

Interestingly, Maltoni et al. (2003) describe a fast and reliable method to generate realistic synthetic fingerprint images, which is implemented in the software tool SFinGe (Synthetic Fingerprint Generator). Using this application, the previously described type 4 attack (to the input of the matcher) with synthetically generated templates could be easily converted to a type 2 attack (to the input of the feature extractor) using the corresponding synthetic fingerprint images. Thus, the attack would be simplified as the intruder would not need to know the template format used in the system (the attack would work even with non-minutiae based matchers). Furthermore, an algorithm to reconstruct the real fingerprint image from its ISO minutia-based template has been proposed by Cappelli et al. (2007b). In this case, as it was exposed by Cappelli et al. (2007a), if a legitimate user template is compromised it could be employed to carry out a masquerade attack (type 2) against the system (reconstructing the fingerprint image), or even a direct attack with a gummy fingerprint from the image (Galbally et al., 2008).

3. Implementation of hill climbing attacks

The hill climbing attacks studied in this work are implemented as follows. The attacks assume that the user template is stored in the system as a set of minutiae. Minutiae are defined by their position (x,y) and orientation α.

At the beginning of the attack a set of 100 synthetic random minutiae templates is generated. Synthetic templates are divided in 9×9 pixels cells. Each cell can only contain one minutiae, this way we avoid generating minutiae which are closer than the inter-ridge distance. Next, the following steps are followed:

(1) The 100 synthetic minutiae templates are initially sent to the matcher to be compared with the attacked fingerprint.
(2) Out of the 100 synthetic templates, the one that produces the highest score is stored.
(3) The saved template is iteratively modified by means of:
   (a) Changing an existing minutia by moving it to an adjacent cell or by changing its orientation.
   (b) Adding a minutia.
   (c) Replacing a minutia.
   (d) Deleting a minutia from the template.
(4) The four types of iteration mentioned above are executed one at a time and changes are only saved if they cause an improvement in the score.
(5) The algorithm stops either when the decision threshold or the maximum number of iterations allowed is reached.

The performance of these attacks is compared to the one of brute force attacks, in terms of the required attempts to reach the decision threshold. Supposing we have access to an unlimited collection of different fingerprints, the theoretical number of attempts that a brute force attack would need against a verification system is equal to the inverse of the False Acceptance Rate ($1/FAR$).

In the indirect attacks evaluation we study the impact of several parameters, such as the number of initial minutiae or the effectiveness of each type of iteration (a, b, c and d). The effects of the usage of a Region of Interest (ROI) for the placement of synthetic minutiae (i.e. in the generation of the 100 synthetic template set and in step 3 of the algorithm) are also studied. The ROI is defined as the area of the fingerprint images in which most minutiae are found and is obtained heuristically from a fingerprint database as described in Section 5. It can be hypothesized that the generation of synthetic features only in the ROI should improve the algorithm effectiveness, reducing the number of iterations needed.

Once the feasibility of the attack is studied, we analyze the effect of score quantization on the success rate of the best configuration of the attack, and its possible use as a countermeasure against the hill climbing algorithm.

4. Countermeasures against indirect attacks

Effective countermeasures must be implemented in order to reduce the risk and impact of attacks to a biometric recognition system. While liveness detection has been stated as the most popular countermeasure against direct attacks, many alternatives exist for the case of indirect attacks. These measures can be broadly classified in *physical* and *logical* security measures, although specific *algorithmic* countermeasures proposed in the biometric community may increase the overall system security or reduce the impact

of system-specific attacks, such as hill climbing. Countermeasures can be summarized as follows:

Physical countermeasures against indirect attacks protect the system internal modules and data channels. The feature extractor, matcher and template database must be secured, in order to avoid access to them or the introduction of rogue applications and trojan horses in the system. The communication channels must be secured, ensuring that no information can be introduced or extracted from them.

Logical countermeasures provide an additional level of security once physical access to the system has been secured. Data encryption and digital signature techniques can be used to ensure that the data transmitted over the communication channels or in the database cannot be accessed or modified and to avoid reverse engineering. In the last few years, template protection schemes have been proposed as a feasible approach to biometric system security (Jain et al., 2008). In these schemes, the biometric templates are protected using cryptographic constructions that enable the matching to be performed in the encrypted domain. The main difficulties found in these applications arise from the need of finding a combination of both stable features and a tolerant error correcting algorithm for the matching. This subject has generated a great interest in the last few years, for example in the fields of fingerprint (Uludag et al., 2005), voice (Monrose et al., 2001) and signature verification (Freire-Santos et al., 2006).

Algorithmic countermeasures to indirect attacks against biometric systems are embedded in the feature extraction or the matching algorithms. Some notable examples at the feature extraction phase are non-invertible transforms or cancelable biometrics (Ratha et al., 2007) which may be combined with cryptography or other logical countermeasures. Cancelable biometrics are not directly aimed against indirect attacks, but may mitigate their effect. Their goal is to create a cancelable user biometric template (e.g. with a non-invertible transform) that can be replaced if it is compromised. Consequently, if a user template is successfully attacked via hill climbing, it could be canceled and replaced. A specific design of the matching algorithm aimed against indirect attacks can also be implemented, adding an additional level of security. This is the case of score quantization, which increases the system robustness against hill climbing attacks.

4.1. Score quantization

The BioAPI consortium (Consortium, 2001) recommends that biometric algorithms emit only quantized matching scores in order to prevent eventual hill climbing attacks. Such quantization means that small changes in the randomly generated templates will normally not result in a modification of the matching score so that the attack does not have the necessary feedback from the system to be carried out successfully.

Adler (2003) introduced a modified hill climbing algorithm which was robust to quantized scores. However, this algorithm was applied to the input images of the feature extractor (type 2 attack) and was very specific for face recognition systems, so its application to type 4 attacks over fingerprint minutiae-based systems is at least unclear.

The fingerprint verification systems put to test in this work against the hill climbing attack produce integer quantized scores, in the ranges observed in the score distributions depicted in Fig. 2. In Section 5, we test if these quantization steps are effective against the proposed attacks and analyze the effects of introducing coarser quantization steps.

5. Experiments

5.1. Description of evaluated systems

The vulnerabilities to hill climbing attacks are studied on two different minutiae-based fingerprint verification systems, one running on a PC and one embedded in a smart card (Match-on-Card):

- The minutiae-based NIST Fingerprint Image Software 2 (NFIS2) (Watson and Garris, 2004). It is a PC-based fingerprint processing and recognition system formed of independent software modules. The feature extractor generates a text file containing the location, orientation and quality of each minutia from the fingerprint. The matcher uses this file to generate the score. The matching algorithm is rotation and translation invariant since it computes only relative distances and orientations between groups of minutiae.
- A prototype Match-on-Card system. The system is a prototype from a Match-on-Card vendor, developed in 2006. It is a minutiae-based system with the matching algorithm fully embedded in a smart card. This is a good method to protect the privacy of users (their templates do not leave the card), while providing reasonable performance with current technology. Some related works on fingerprint MoC have been reported by Bistarelli et al. (2006) and Mueller and Martini (2006). In the experiments the NIST software is used in the feature extraction process and the resulting templates are transformed to the MoC system format and sent to the smart card. Except for basic information about the input–output interface of the smart card, the specificities of the matching algorithm are unknown in our analysis, being thus a realistic attack scenario. The MoC system evaluated in our experiments is shown in Fig. 3 (a).

5.2. Database and experimental protocol

The hill climbing attacks have been studied using a subcorpus of the MCYT database (Ortega-Garcia and Fierrez-Aguilar, 2003). The subcorpus comprises 10 impressions of the right and left index fingers of 75 users ($75 \times 2 \times 10 = 1{,}500$ images), captured electronically with an optical sensor UareU from Digital Persona (500 dpi, 256×400 images). Six of the samples of each finger were acquired with a high control level (small rotation or displacement of the finger core from the center of the sensor was permitted), another two with a medium control level, and the remaining two with low control level (see Figs. 4–7 for example fingerprint images).

In Fig. 3 (b) we depict the two dimensional histogram of all the minutiae locations in the subcorpus, together with a rectangle that was heuristically obtained and which contains the majority of the minutiae. This rectangle defines the Region of Interest (ROI) and

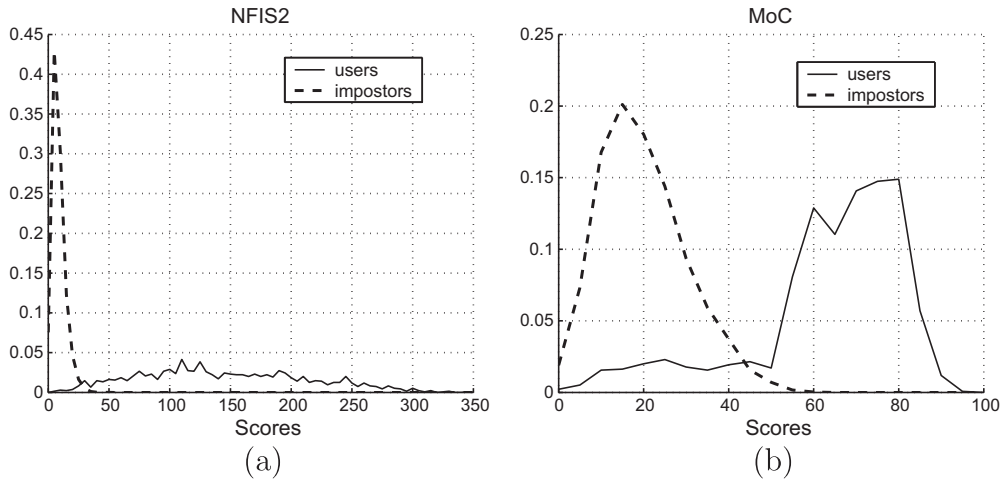

Fig. 2. Score distributions of the NIST (a) and MoC (b) systems.

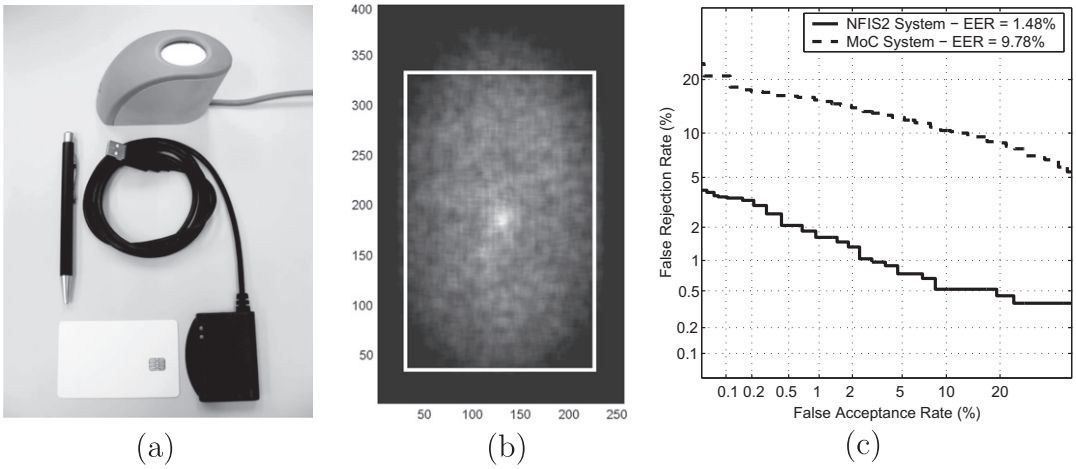

Fig. 3. (a) Top: fingerprint sensor used for acquiring the fingerprints in our experiments. Bottom: MoC system used in our experiments. (b) Histogram of minutiae locations, and Region of Interest (ROI). (c) DET curves corresponding to the NIST2 and MoC systems for the database used in the experiments.

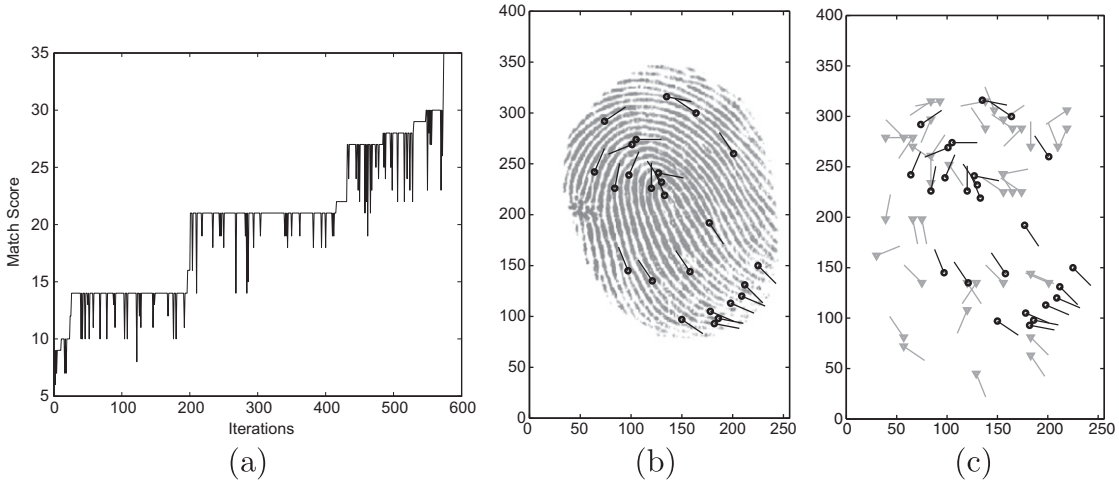

Fig. 4. (a) Score progression, (b) original fingerprint minutiae, and (c) original minutiae (black circles) vs. synthetic minutiae (grey triangles) that achieve a higher score than the decision point on NFIS2 in a relatively short attack.

will be used in the experiments as described in Section 3 to improve the success rate of the attacks. The 1,500 images available in the subcorpus were also used for evaluating the verification error rates of the two studied systems, as depicted in Fig. 3 (c).

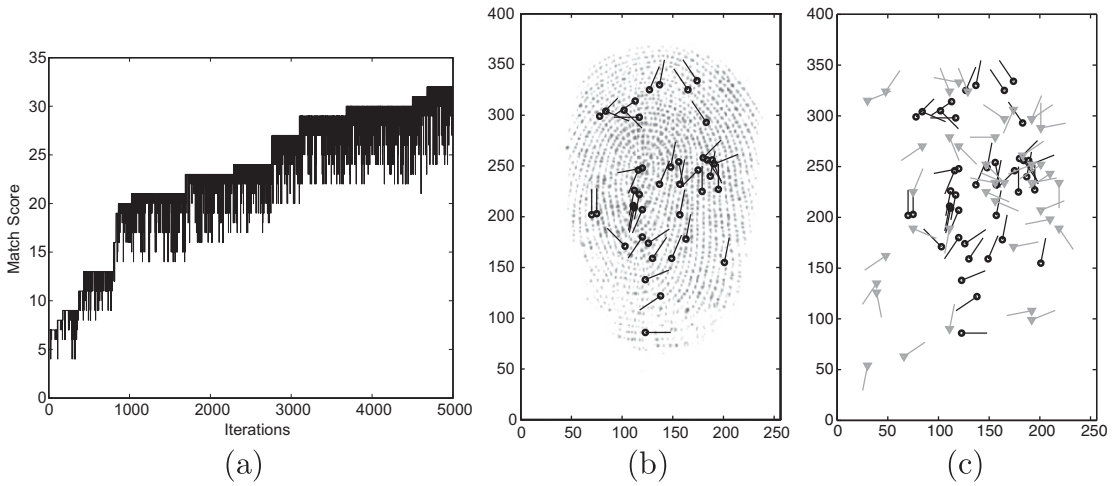

Fig. 5. (a) Score progression, (b) original fingerprint minutiae, and (c) original minutiae (black circles) vs. synthetic minutiae (grey triangles) after 5000 iterations on NFIS2 in an unsuccessful attack.

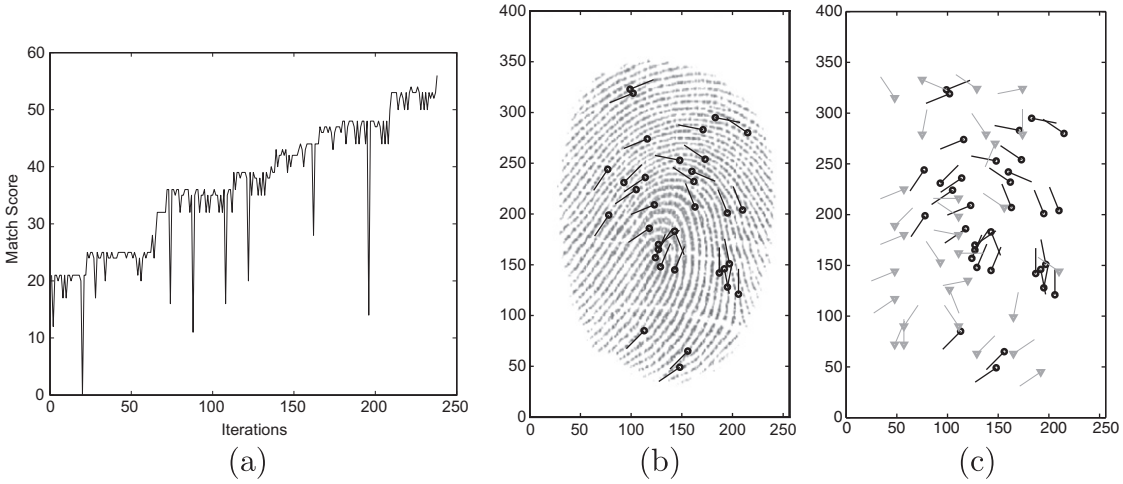

Fig. 6. (a) Score progression, (b) original fingerprint minutiae, and (c) original minutiae (black circles) vs. synthetic minutiae (grey triangles) that achieve a higher score than the decision point on the MoC system in a relatively short attack.

Using one of the impressions of high control level for each fingerprint, the 150 different fingerprints considered in the database were attacked following the algorithm described in Section 3.

For the NFIS2 system, a decision threshold of 35 for the match score is fixed, leading to a 0.10% FAR (False Acceptance Rate) and a 3.33% FRR (False Rejection Rate). This means that a brute force attack would need in average 1/FAR = 1,000 attempts to be successful. For the Match-on-Card system a decision threshold of 55 is selected, resulting in a FAR of 0.16% and a FRR of 17.33%. In this case a brute force attack would need around 630 attempts to break the system. The decision thresholds, FRR and FAR values are obtained from the DET curves of both systems depicted in Fig. 3 (c). The brute force attack number of iterations (1,000 and 630 respectively) will be considered in the experiments in order to evaluate the success rate and speed of the attacks. An attack is considered as successful if it needs less iterations than the ones a brute force would theoretically need. We establish a maximum of 5,000 and 2,000 iterations for the NFIS2 and the MoC system respectively. If the decision threshold is not reached within these limits of iterations, the algorithm ends.

Different configurations are tested, varying the number of initial synthetic minutiae, modifying the iterations of the algorithm or using the previously described ROI. Following the defined protocol, 150 attacks are performed for each possible configuration, that is, each different fingerprint in the database is attacked once per experiment.

5.3. Attack results

Attacks on the NFIS2 system

In the first experiment, the effect of using a ROI is studied. In Table 1 (a) the effect of the ROI when it is included in the configuration of the attack can be seen. The number of hill climbing attacks that need less iterations than an eventual brute force attack raises from 2 to 7 when no synthetic minutiae are allowed to be placed outside the ROI. The number of successful attacks before the maximum number of attempts is reached increases from 64 to 85. This first experiment (Table 1 (a)) also shows that not all the iterations (changing, adding, replacing or deleting a minutia) have the same probability of improving the matching score.

A second experiment is performed to analyze the effectiveness of each type of iteration, defined in Section 3. In Table 1 (b) the effect of eliminating the least effective iterations is studied. The results show that iterations *a* and *d* (changing and deleting a

Table 1
Hill climbing results on NFIS2.

ROI	Iterations	Initial minutiae	Mean score raises				Success before 1000 iterations	Success before 5000 iterations
			a	*b*	*c*	*d*		
(a) Hill climbing statistics using all iterations with and without ROI								
No	*a*, *b*, *c*, *d*	38	1.87	5.16	6.13	0.90	2/150	64/150
Yes	*a*, *b*, *c*, *d*	38	2.41	4.93	5.60	1.35	7/150	85/150
(b) Hill climbing statistics deleting low performing iterations								
Yes	*a*, *b*, *c*, *d*	38	2.41	4.93	5.60	1.35	7/150	85/150
Yes	*a*, *b*, *c*	38	3.18	7.70	7.91	–	28/150	145/150
Yes	*b*, *c*	38	–	9.25	9.76	–	40/150	143/150
(c) Hill climbing statistics using different amounts of initial minutiae								
Yes	*b*, *c*	25	–	10.85	8.95	–	28/150	136/150
Yes	*b*, *c*	38	–	9.25	9.76	–	40/150	143/150
Yes	*b*, *c*	55	–	5.68	13.67	–	12/150	132/150

Table 2
Hill climbing results on the Match-on-Card system.

ROI	Iterations	Initial minutiae	Mean score raises				Success before 630 iterations	Success before 2000 iterations
			a	*b*	*c*	*d*		
(a) Hill climbing statistics using different amounts of initial minutiae								
Yes	*b*, *c*	10	–	7.70	5.30	–	65/150	133/150
Yes	*b*, *c*	25	–	5.53	10.08	–	123/150	146/150
Yes	*b*, *c*	38	–	3.55	13.27	–	78/150	139/150
(b) Hill climbing statistics deleting low performing iterations								
Yes	*a*, *b*, *c*, *d*	25	1.22	4.60	5.71	4.68	52/150	132/150
Yes	*b*, *c*, *d*	25	–	5.24	5.98	5.03	79/150	138/150
Yes	*b*, *c*	25	–	5.53	10.08	–	123/150	146/150
(c) Hill climbing statistics with and without rectangular ROI								
Yes	*b*, *c*	25	–	5.53	10.08	–	123/150	146/150
No	*b*, *c*	25	–	6.13	9.15	–	91/150	148/150

minutiae respectively) have barely any impact in the success rate of the attacks. Actually, when they are not performed, the number of broken fingerprints increased from 85 to 143.

In the third experiment we analyze the impact of the initial number of minutiae in the synthetic fingerprints, using the best configuration so far, i.e., using the ROI and performing iterations *b* and *c*. The NFIS2 system extracts an average of 38 minutiae points from the fingerprints in the database considered. We can see in Table 1 (c) that the success rate of the attacks improves when the initial number of minutiae approaches 38.

In Fig. 4 we show the minutiae maps and the evolution of the matching score in a successful attack against the NIST system. Fig. 5 shows the same data for an unsuccessful attack. In the first case around 580 iterations are needed to reach the desired matching score (35), while in the failed attack the maximum allowed number of iterations is reached before the algorithm reaches the positive verification score.

Attacks on the MoC system

The experiments for the MoC system follow an inverse order than the ones for the NIST system. Based on the best configuration of the attack for the NFIS2 system, we first study the influence of the initial number of minutiae over the final success rate in the MoC system. In this case we find that better results are achieved using 25 initial minutiae, instead of the 38 used in the NFIS2 system. In Table 2 (a) we can see that the number of fingerprints cracked before a brute force attack increases from 78 to 123 when the initial number of minutiae is reduced from 38 to 25.

The contribution of each type of iteration is then analyzed. In Table 2 (b) the effect of each of the iterations over the match score can be observed. As happened in the NFIS2 system, the most effective iterations are *b* and *c*, so *a* and *d* can be again discarded.

In the last experiment we focus on the impact of the ROI over the number of successful attacks. As can be seen in Table 2 (c) when no minutiae are allowed to be placed outside the ROI, the number of fingerprints cracked before a brute force attack increases from 91 to 123. No significant improvement can be observed in the use of the ROI when the maximum number of iterations is reached.

In Figs. 6 and 7 the minutiae maps and the evolution of the matching score in a successful and an unsuccessful attack are respectively depicted for the MoC system. In the first case the desired matching score of 55 is reached in around 240 iterations, while in the failed attack the maximum number of iterations is reached before the algorithm gets to the positive verification matching score.

The results show that the performance of hill climbing attacks is heavily dependent upon the system under attack and the iterations that are performed. Attacks with reduced number of minutiae are highly successful against the MoC system, while their performance against NFIS2 is very poor, even when using the same minutiae feature extractor from NIST. This is probably due to the limitations of the matcher embedded in the smart card.

It may be derived from the results that, at least in the case of NFIS2, hill climbing attacks are less effective than brute force attacks. This statement must be taken with care, as hill climbing attacks require much less resources than the ones needed by a brute force attack. In fact, to perform an efficient brute force attack, the attacker must have a database of more than a thousand different real fingerprint templates which is not straightforward to obtain, whereas there is no need for real templates in the case of a hill climbing attack.

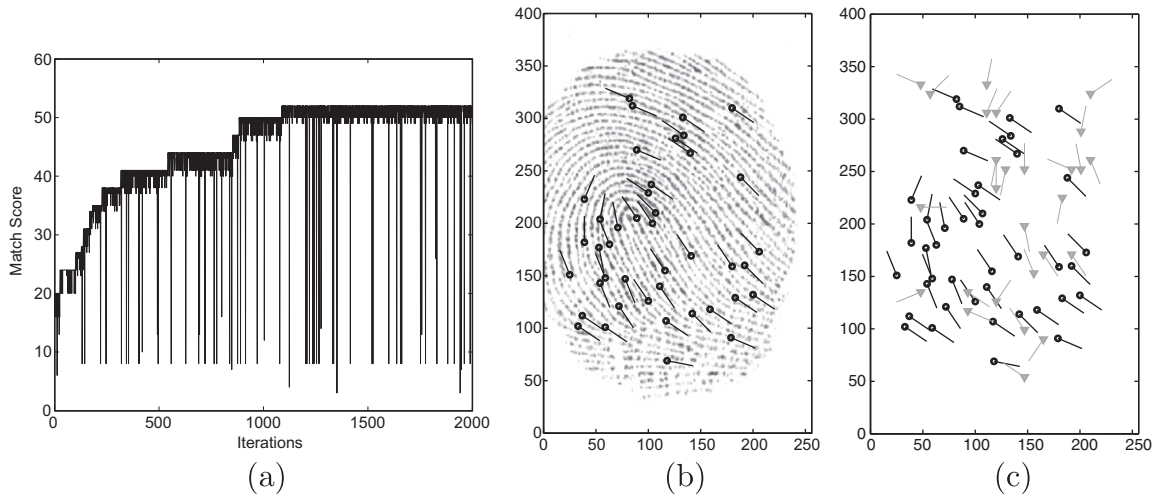

Fig. 7. (a) Score progression, (b) original fingerprint minutiae, and (c) original minutiae (black circles) vs. synthetic minutiae (grey triangles) after 2000 iterations on the MoC system in an unsuccessful attack.

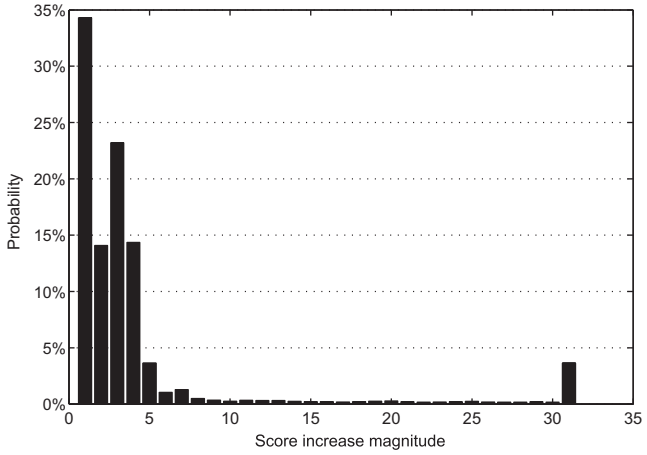

Fig. 8. Distribution of magnitudes corresponding to score increases during an experiment with the MoC system (150 attacks). Match scores are quantized as integer numbers.

Table 3
Evaluation of the hill climbing attack against the NIST and ridge-based systems with score quantization.

	1-unit QS	2-unit QS	5-unit QS
(a) Results for the hill climbing algorithm against the NIST system with different score Quantization Steps (QS)			
Success before 1000 iterations	40/150	3/150	1/150
Success before 5000 iterations	145/150	51/150	1/150
(b) Results for the hill climbing algorithm against the MoC system with different score Quantization Steps (QS)			
Success before 630 iterations	123/150	6/150	0/150
Success before 1000 iterations	146/150	26/150	0/150

5.4. Countermeasure results: score quantization

In our experiments, the four types of iterations described (*a*, *b*, *c* and *d*) may increase or decrease the match score during a hill climbing attack, as seen in Figs. 4-7. It is found that ca. 30% of the total number of iterations produce a score increase. The distribution of the score increase magnitudes during the iterations from the 150 attacks to the MoC system in one of the previous experiments is shown in Fig. 8. As can be seen, in most cases (more than 33% of the score increases), the score increases 1 unit. However, the score is increased in less than 5 points in more than 85% of the cases. It must be taken into account that only score increases are shown in the histogram, as many iterations produce score decreases.

Further experiments are carried out where the similarity scores are forced to follow a 2 and 5 unit quantization step (i.e., scores are quantized to multiples of 2 and 5). Taking into account the distribution shown in Fig. 8, it is expected that this quantization procedure may protect the system against the proposed attacks since most iterations produce score variations which are lower than these quantization steps. In Table 3 we show the performance of the best configuration of the hill climbing attack against the NIST system (a) and the MoC system (b) for different quantization steps (QS).

The experiments show that score quantization is an effective measure in order to prevent the studied hill climbing attack, as the performance of the algorithm drops drastically for just a 2 quantization step in the two systems tested (only 3 and 6 broken accounts respectively). When 5 unit quantization steps are used, the system is nearly invulnerable to the implemented hill climbing attacks. Moreover, verification performance is not significantly decreased with these quantization steps.

6. Conclusions and Future Work

The vulnerabilities to indirect attacks of two representative fingerprint verification systems have been evaluated. Two state-of-the-art fingerprint recognition systems, one running on a PC and the other system fully embedded in a smart card, were evaluated against hill climbing attacks. Experiments were carried out on a sub corpus of the MCYT database. The attacks showed a big dependency on the type of iterations performed and on the system being attacked. For a sufficient number of iterations, success rates of over 90% were reached for both systems, being the PC system the one that needed a higher number of attempts to be cracked. Score quantization has also been studied as a possible countermeasure against hill climbing attacks. It has proved to reduce drastically the attack success rate.

In the systems under analysis, score quantization does not significantly affect the verification performance. Nevertheless score quantization presents some drawbacks, being the most important of them that as the quantization step size grows, the matching

scores decrease their utility for multi-biometric applications (Ross et al., 2006), which typically rely on fusion techniques of real-valued scores (Fierrez-Aguilar et al., 2005).

Interestingly, not all the fingerprints showed the same robustness against this type of attacks, being some of them much more difficult to crack than others. Furthermore, a reduced group of them were not possible to be bypassed using any of the attack configurations tested, which is subject of further research.

Acknowledgments

This work has been supported by Spanish Ministry of Defense, Direccin General de la Guardia Civil, Contexts (S2009/TIC-1485), and Bio-Challenge (TEC2009–11186). J. G. was supported by a FPU Fellowship from the Spanish MEC. J. F. was supported by a Marie Curie Fellowship from the European Comission.

References

Adler, A., 2003. Sample images can be independently restored from face recognition templates. In: Proc. of Canadian Conf. on Electrical and Computer Engineering. Vol. 2. pp. 1163–1166.

BioAPI Consortium, March 2001. BioAPI specification (version 1.1). <www.bioapi.org/Downloads/BioAPI>.

Bistarelli, S., Santini, F., Vaccarelli, A., 2006. An asymmetric fingerprint matching algorithm for java card tm. Pattern Anal. & Appl. 9 (4), 359–376.

Cappelli, R., Lumini, A., Maio, D., Maltoni, D., 2007a. Evaluating minutiae template vulnerability to masquerade attack. In: Proc. IEEE AutoID. pp. 174–179.

Cappelli, R., Maio, D., Lumini, A., Maltoni, D., 2007b. Fingerprint image reconstruction from standard templates. IEEE Trans. Pattern Anal. Machine Intell. 29, 1489–1503.

Cappelli, R., Maio, D., Maltoni, D., Wayman, J.L., Jain, A.K., 2006. Performance evaluation of fingerprint verification systems. IEEE Trans. Pattern Anal. Machine Intell. 28 (1), 3–18.

Fierrez-Aguilar, J., Ortega-Garcia, J., Gonzalez-Rodriguez, J., Bigun, J., 2005. Discriminative multimodal biometric authentication based on quality measures. Pattern Recognit. 38 (5), 777–779.

Franco, A., Maltoni, D., 2007. Advances in Biometrics: Sensors, Systems and Algorithms, Eds. N.K. Ratha and V. Govindaraju. Springer, Ch. Fingerprint Synthesis and Spoof Detection.

Freire-Santos, M., Fierrez-Aguilar, J., Ortega-Garcia, J., 2006. Cryptographic key generation using handwritten signature. In: Proc. SPIE 6202, pp. 225–231.

FVC, 2006. Fingerprint Verification Competition. <http://bias.csr.unibo.it/fvc2006>.

Galbally, J., Cappelli, R., Lumini, A., de Rivera, G.G., Maltoni, D., Fierrez, J., Ortega-Garcia, J., Maio, D., 2010a. An evaluation of direct attacks using fake fingers generated from iso templates. Pattern Recognit. Lett. 31 (8), 725–732.

Galbally, J., Cappelli, R., Lumini, A., Maltoni, D., Fierrez, J., 2008. Fake fingertip generation from a minutiae template. In: Proc. of 19th Internat. Conf. on Pattern Recognition ICPR 2008. pp. 1–4.

Galbally, J., Fierrez, J., Ortega-Garcia, J., 2007. Bayesian hill-climbing attack and its application to signature verification. In: Proc. IAPR Internat. Conf. on Biometrics, ICB. Springer LNCS-4642, pp. 386–395.

Galbally, J., McCool, C., Fierrez, J., Marcel, S., Ortega-Garcia, J., 2010b. On the vulnerability of face verification systems to hill-climbing attacks. Pattern Recognit. 43 (3), 1027–1038.

Grother, P., Salamon, W., Watson, C., Indovina, M., Flanagan, P., February 2008. MINEX II, performance of fingerprint match-on-card algorithms, phase II report, NIST 7477. Tech. rep., Information Access Division, National Institute of Standards and Technology NIST.

Hill, C.J., 2001. Risk of masquerade arising from the storage of Biometrics, B.S. Thesis. Australian National University.

Jain, A.K., Nandakumar, K., Nagar, A., January 2008. Biometric template security. EURASIP J. Advances Signal Process. 2008.

Jain, A.K., Ross, A., Pankanti, S., 2006. Biometrics: A tool for information security. IEEE Trans. Inform. Forensics Security 1 (2), 125–143.

Maltoni, D., Cappelli, R., 2007. Handbook of Biometrics. Springer, pp. 23–42 (Ch. Fingerprint recognition).

Maltoni, D., Maio, D., Jain, A.K., Prabhakar, S., 2003. Handbook of Fingerprint Recognition. Springer.

Mohanty, P., Sarkar, S., Kasturi, R., 2007. From scores to face templates: a model-based spproach. IEEE Trans. Pattern Anal. Machine Intell. 29 (12), 2065–2078.

Monrose, F., Reiter, M.K., Li, Q., Wetzel, S., 2001. Cryptographic key generation from voice. In: Proc. of IEEE Symposium on Security and Privacy. pp. 202–213.

Mueller, R., Martini, U., 2006. Decision level fusion in standardized fingerprint match-on-card. In: Proc. IEEE ICCARV. pp. 185–190.

Muramatsu, D., 2008. Online signature verification algorithm using hill-climbing method. In: Proc. of IEEE/IFIP Internat. Conf. on Embedded and Ubiquitous Computing. pp. 133–138.

Ortega-Garcia, J., Fierrez-Aguilar, et al., 2003. MCYT baseline corpus: a bimodal biometric database. IEE Proc. Vision Image Signal Process. 150 (6), 391–401.

Ratha, N., Chikkerur, S., Connell, J., Bolle, R.M., 2007. Generating cancelable fingerprint templates. IEEE Trans. Pattern Anal. Machine Intell. 29 (4), 561–572.

Ratha, N., Connell, J., Bolle, R., 2001. An analysis of minutiae matching strength. In: Proc. AVBPA. Springer LNCS, pp. 223–228.

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

Sanchez-Reillo, R., Mengihar-Pozo, L., Sanchez-Avila, C., 2003. Microprocessor smart cards with fingerprint user authentication. IEEE AESS Syst. Magazine 18 (3), 22–24.

Soutar, C., Gilroy, R., Stoianov, A., 1999. Biometric system performance and security. In: Proc. of IEEE Automatic Identification Advanced Technologies.

Uludag, U., Jain, A.K., 2004. Attacks on biometric systems: a case study in fingerprints. In: Proc. SPIE-EI 2004, Security, Seganography and Watermarking of Multimedia Contents VI. Vol. 5306. pp. 622–633.

Uludag, U., Pankanti, S., Jain, A.K., 2005. Fuzzy vault for fingerprints. In: Proc. of Audio- and Video-based Biometric Person Authentication AVBPA. pp. 310–319.

Watson, G.I., Garris, M.D., et al., 2004. User's guide to NIST Fingerprint Image Software 2 (NFIS2). National Institute of Standards and Technology.

Wilson, C., Hicklin, R.A., Korves, H., Uler, B., Zoepfl, M., Bone, M., Grother, P., Micheals, R., Otto, S., Watson, C., 2004. Fingerprint vendor technology evaluation 2003: summary of results and analysis report, nistir 7123. Tech. rep., National Institute of Standards and Technology.

Yamazaki, Y., Nakashima, A., Tasaka, K., Komatsu, N., 2005. A study on vulnerability in on-line writer verification system. Internat. Conf. on Document Analysis and Recognition, ICDAR 1, 640–644.