

BIOMETRIC TEMPLATE SECURITY: CHALLENGES AND SOLUTIONS

Anil K. Jain

Michigan State University
East Lansing, MI, 48824, USA
jain@cse.msu.edu

Arun Ross

West Virginia University
Morgantown, WV, 26506, USA
Arun.Ross@mail.wvu.edu

Umut Uludag

Michigan State University
East Lansing, MI, 48824, USA
uludagum@cse.msu.edu

<http://biometrics.cse.msu.edu>

ABSTRACT

A biometric system is vulnerable to a variety of attacks aimed at undermining the integrity of the authentication process. These attacks are intended to either circumvent the security afforded by the system or to deter the normal functioning of the system. We describe the various threats that can be encountered by a biometric system. We specifically focus on attacks designed to elicit information about the original biometric data of an individual from the stored template. A few algorithms presented in the literature are discussed in this regard. We also examine techniques that can be used to deter or detect these attacks. Furthermore, we provide experimental results pertaining to a hybrid system combining biometrics with cryptography, that converts traditional fingerprint templates into novel cryptographic structures.

1. INTRODUCTION

Establishing the identity of an individual is of paramount importance in several civilian and government applications where errors in recognition can undermine the integrity of the system. Example of such applications include international border control, access to nuclear facilities, airport security, issuance of passports or driver licences, etc. Traditionally, a combination of ID cards (token-based security) and PINs/passwords (knowledge-based security) has been used to validate the identity of an individual. These methods are, however, vulnerable to the wiles of an impostor and cannot be reliably used in large-scale applications such as border control, where the throughput is required to be in the order of thousands of users per day. The advent of biometrics has introduced a secure and efficient alternative to traditional authentication schemes. Biometrics is the science of establishing or determining an identity based on the physiological or behavioral traits of an individual. These traits include fingerprints, facial features, iris, hand geometry, voice, signature, etc. In conjunction with traditional authentication schemes, biometrics is a potent tool for establishing identity [1].

A typical biometric system comprises of several modules. The *sensor module* acquires the raw biometric data of an individual in the form of an image, video, audio or some other signal. The *feature extraction module* operates on the biometric signal and extracts a salient set of features to represent the signal; during user enrolment the extracted feature set, labeled with the user's identity, is stored in the biometric system and is known as a *template*. The *matching module* compares the feature set extracted during authentication with the enrolled template(s) and generates match scores. The *decision module* processes these match scores in order to either determine or verify the identity of an individual. Thus, a bio-

metric system may be viewed as a pattern recognition system whose function is to classify a biometric signal into one of several identities (viz., identification) or into one of two classes - genuine and impostor users (viz., verification).

While a biometric system can enhance user convenience and bolster security, it is also susceptible to various types of threats as discussed below [2, 3].

1. *Circumvention*: An intruder may gain access to the system protected by biometrics and peruse sensitive data such as medical records pertaining to a legitimately enrolled user. Besides violating the privacy of the enrolled user, the impostor can also modify sensitive data.
2. *Repudiation*: A legitimate user may access the facilities offered by an application and then claim that an intruder had circumvented the system. A bank clerk, for example, may modify the financial records of a customer and then deny responsibility by claiming that an intruder could have possibly stolen her biometric data.
3. *Covert acquisition*: An intruder may surreptitiously obtain the raw biometric data of a user to access the system. For example, the latent fingerprints of a user may be lifted from an object by an intruder and later used to construct a digital or physical artefact of that user's finger.
4. *Collusion*: An individual with wide super-user privileges (such as an administrator) may deliberately modify system parameters to permit incursions by an intruder.
5. *Coercion*: An impostor may force a legitimate user (e.g., at gunpoint) to grant him access to the system.
6. *Denial of Service (DoS)*: An attacker may overwhelm the system resources to the point where legitimate users desiring access will be refused service. For example, a server that processes access requests can be flooded with a large number of bogus requests, thereby overloading its computational resources and preventing valid requests from being processed.

Ratha et al. [4] identified several different levels of attacks that can be launched against a biometric system (Figure 1): (i) a fake biometric trait such as an artificial finger may be presented at the sensor, (ii) illegally intercepted data may be resubmitted to the system, (iii) the feature extractor may be replaced by a Trojan horse program that produces pre-determined feature sets, (iv) legitimate feature sets may be replaced with synthetic feature sets, (v) the matcher may be replaced by a Trojan horse program that always outputs high scores thereby defying system security, (vi) the templates stored in the database may be modified or removed, or new templates may be introduced in the database, (vii) the data in the communication channel between various modules of the system may be altered, and (viii) the final decision

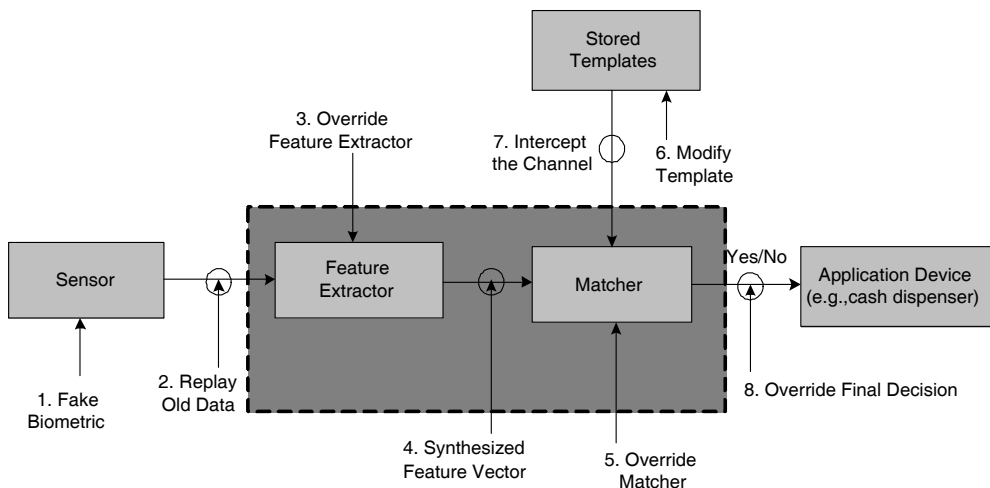


Figure 1: Vulnerabilities in a biometric system (adapted from [4]).

output by the biometric system may be overridden.

The UK Biometric Working Group (UK-BWG) lists several factors that can affect the integrity of the template [5]: (i) accidental template corruption due to a system malfunction such as a hardware failure, (ii) deliberate alteration of an enrolled template by an attacker, and (iii) substitution of a valid template with a bogus template for the purpose of deterring system functionality.

In this paper, we discuss several issues related to template security. Specifically, we examine some of the attacks that can be used to compromise template information. Then, we analyze possible solutions to alleviate this problem.

2. COMPROMISING TEMPLATE INFORMATION

A template represents a set of salient features that summarizes the biometric data (signal) of an individual. Due to its compact nature, it is commonly assumed that the template cannot be used to elicit complete information about the original biometric signal. Furthermore, since the templates are typically stored in an encrypted form, it is substantially difficult to decrypt and determine the contents of the stored template (without the knowledge of correct decrypting keys). Thus, traditionally, template-generating algorithms have been viewed as one-way algorithms. However, in the recent literature there have been techniques presented that contradict these assumptions.

Adler [6] demonstrated that a face image can be regenerated from a face template using a “Hill Climbing Attack” (attack level 2 in Figure 1). He employed an iterative scheme to reconstruct a face image using a face verification system that releases match scores. The algorithm first selects an estimate of the target face from a local database comprising of a few frontal images by observing the match score corresponding to each image. An eigen-face (computed from the local database) scaled by 6 different constants is added to this initial estimate resulting in a set of 6 modified face images which are then presented to the verification system. The image resulting in an improved match score is retained and this process is repeated in an iterative fashion. Within a few thousand iterations, an image that can successfully masquerade as

the target face image is generated. The important feature of this algorithm is that it does not require any knowledge of either the matching technique or the structure of the template used by the authentication system. Furthermore, template encryption does not prevent this algorithm from successfully determining the original face image. The algorithm was able to “break” three commercial face recognition systems.

Uludag and Jain [3] devised a synthetic template generator (STG) that also uses the “Hill Climbing Attack” (attack level 4 in Figure 1) to determine the contents of a target fingerprint template (D_i) for the i^{th} user (see Figure 2). The minutiae template is assumed to be a sequence of (r, c, θ) values representing the location and orientation of component fingerprint minutiae. The STG begins by generating a fixed number of synthetic templates each comprising of randomly generated minutiae points. These templates are compared against the target template in the database (via the matcher) and the synthetic template resulting in the best match score is retained. The retained template is then modified iteratively via the following four operations: (i) the r , c and θ values of an existing minutia are perturbed, (ii) an existing minutia is replaced with a new minutia, (iii) a new minutia is added to the template, and (iv) an existing minutia is deleted. The modified template (T_i^j) is compared against the target template and the match score ($S(D_i, T_i^j)$) computed. This process, viz., modifying the current synthetic template and comparing it against the target template, is repeated until the match score exceeds a pre-determined threshold. The authors used this scheme to break into 160 fingerprint accounts; their algorithm required only 271 iterations, on an average, to exceed the matching threshold for each one of those 160 accounts.

Hill [7] describes a masquerade attack wherein the fingerprint structure is determined using the minutiae template alone (attack level 7 in Figure 1). It is assumed that each minutia point is characterized using its 2D location, orientation and the curvature of the ridge associated with it. Based on minutiae points, the author predicts the shape of the fingerprint (i.e., its class) using a neural network classifier consisting of 23 input neurons, 13 hidden neurons and 4 output neurons (corresponding to 4 fingerprint classes). However,

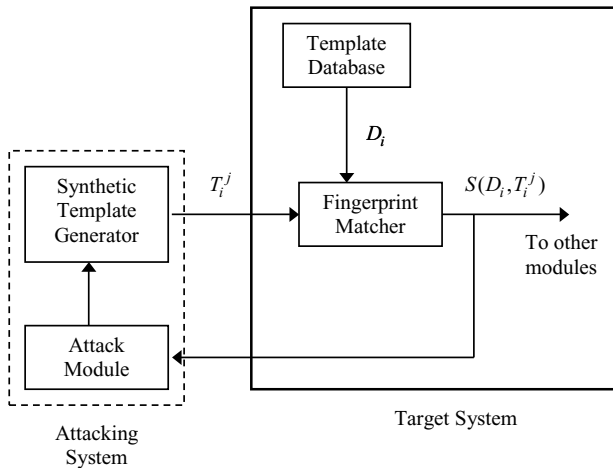


Figure 2: Algorithm to synthesize minutiae templates [3].

the classification performance is rather low (an error rate of 28.9% on a small set of 242 fingerprints). The author then uses a generic orientation map and the minutiae information to generate line drawings that are a digital artefact of the original fingerprint. The proposed technique is observed to work on a database of 25 fingerprints from arch class.

Ross et al. [8] propose another technique to elicit the fingerprint structure from the minutiae template (attack level 7 in Figure 1). Each minutia is assumed to be represented by its 2D spatial location and its local orientation. The authors identify minutia triplets which are used to estimate the underlying orientation map. The estimated orientation map is observed to be remarkably consistent with the flow of ridges in the original (unseen) parent fingerprint. Furthermore, they use a set of 11 features derived from the minutiae points to predict the class of the fingerprint. A 5 Nearest-Neighbor classifier is used to classify the minutiae set of a fingerprint into one of four classes. Their classification experiment conducted on a dataset of 2200 fingerprints exhibits an error rate of 18%. Finally, they use Gabor-like filters (suggested by Cappelli et al. [9]) to generate fingerprints based on the orientation map (Figure 3).




Figure 3: Reconstructing fingerprints [8]: (a) Minutiae distribution of a fingerprint image, (b) predicted orientation map, (c) reconstructed fingerprint.

Besides these types of attacks, an intruder may alter the contents of a template in order to deter a legitimate user from being successfully verified (attack level 6 in Figure 1).

3. PROTECTING BIOMETRIC TEMPLATES

Several methods have been suggested in the literature to protect biometric templates from revealing important information. In order to prevent the Hill-Climbing Attack from successfully converging, Soutar [10] has suggested the use of coarsely quantized match scores by the matcher. However, Adler [11] demonstrated that it is still possible to estimate the unknown enrolled image although the number of iterations required to converge is significantly higher now.

Yeung and Pankanti [12] describe an invisible fragile watermarking technique to detect regions in a fingerprint image that have been tampered by an attacker. In the proposed scheme, a chaotic mixing procedure is employed to transform a visually perceptible watermark to a random-looking textured image in order to make it resilient against attacks. This “mixed” image is then embedded in a fingerprint image. The authors show that the presence of the watermark does not affect the feature extraction process. The use of a watermark also imparts copyright capability by identifying the origin of the raw fingerprint image.

Jain and Uludag [13] suggest the use of steganography principles to hide biometric data (e.g., fingerprint minutiae) in host images (e.g., faces). This is particularly useful in distributed systems where the raw biometric data may have to be transmitted over a non-secure communication channel. Embedding biometric data in an innocuous host image prevents an eavesdropper from accessing sensitive template information. The authors also discuss a novel application wherein the facial features of a user (i.e., eigen-coefficients) are embedded in a host fingerprint image (of the user). In this scenario, the watermarked fingerprint image of a person may be stored in a smart card issued to that person. At an access control site, the fingerprint of the person possessing the card will first be compared with the fingerprint present in the smart card. The eigen-coefficients hidden in the fingerprint image can then be used to reconstruct the user’s face thereby serving as a second source of authentication.

Ferri et al. [14] propose an algorithm to embed dynamic signature features into face images present on ID cards. These features are transformed into a binary stream after compression (used in order to decrease the amount of payload data). A computer-generated hologram converts this stream into the data that is finally embedded in the blue-channel of a face image. During verification, the signature features hidden in the face image are recovered and compared against the signature obtained on-line. Ferri et al. [14] report that any modification of the face image can be detected, thereby disallowing the use of fake ID cards.

Since the biometric trait of a person cannot be easily replaced (unlike passwords and PINs), a compromised template would mean the loss of a user’s identity. Ratha et al. [15] propose the use of distortion functions to generate biometric data that can be *canceled* if necessary. They use a non-invertible transformation function that distorts the input biometric signal (e.g., face image) prior to feature extraction or, alternately, modifies the extracted feature set (e.g., minutiae points) itself. When a stored template is compromised, then the current transformation is replaced with a new function thereby “canceling” the current (compromised) template and generating a new one. This also permits the use of the same biometric trait in several different applications by merely adopting an application-specific transforma-

tion function. However, it is not clear how matching can be accomplished in the transformed domain.

In the realm of template transformation, the so-called *biometric cryptosystems* are gaining popularity (for a survey on existing techniques, see [16]). These systems combine biometrics and cryptography at a level that allows biometric matching to effectively take place in the cryptographic domain, hence exploiting the associated higher security. For example, Uludag et al. [17] convert fingerprint templates (minutiae data) into point lists in 2D space, which implicitly hide a given secret (e.g., a 128-bit key). The list does not reveal the template data, since it is augmented with chaff points to increase security. The template data is identified only when matching minutiae data from an input fingerprint is available. The system is observed to operate at a Genuine Accept Rate (GAR) of 76% with no false accepts on a database comprising of 229 users.

Although several techniques have been proposed to enhance the security of a user's template, government regulations will also have to be established in order to address the issue of template privacy. For example, issues related to the sharing of biometric templates across agencies (e.g., medical companies and law-enforcement agencies) and the inferring of personal information about an enrolled user from biometric data (e.g., "Is this person prone to diabetes?") have to be countered by establishing an appropriate legal framework.

4. SUMMARY AND CONCLUSIONS

We have discussed various types of attacks that can be launched against a biometric system. We have specifically highlighted techniques that can be used to elicit the contents of a biometric template thereby compromising privileged information. We discuss the importance of adopting watermarking and steganography principles to enhance the integrity of biometric templates. Cancelable biometrics may be used to "reset" the biometric template of a user in the event that the user's template is compromised. Also, biometric cryptosystems can contribute to template security by supporting biometric matching in secure cryptographic domains.

Smart cards are gaining popularity as the medium for storing biometric templates. As the amount of available memory increases (e.g., state-of-the-art smart cards have 64-KByte EEPROM), there is a propensity to store more information in the template. This increases the risks associated with template misuse. As a result, the issue of template security and integrity continues to pose several challenges, and it is necessary that further research be conducted in this direction.

REFERENCES

- [1] A. K. Jain, R. Bolle, and S. Pankanti, eds., *Biometrics: Personal Identification in Networked Society*. Kluwer Academic Publishers, 1999.
- [2] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*. Springer-Verlag, 2003.
- [3] U. Uludag and A. K. Jain, "Attacks on biometric systems: a case study in fingerprints," in *Proc. SPIE, Security, Seganography and Watermarking of Multimedia Contents VI*, vol. 5306, pp. 622–633, (San Jose, CA), January 2004.
- [4] N. Ratha, J. H. Connell, and R. M. Bolle, "An analysis of minutiae matching strength," in *Proc. Audio and Video-based Biometric Person Authentication (AVBPA)*, pp. 223–228, (Halmstad, Sweden), June 2001.
- [5] U.K. Biometric Working Group, "Biometric security concerns," Technical Report, CESG, September 2003, <http://www.cesg.gov.uk/site/ast/biometrics/media/BiometricSecurityConcerns.pdf>.
- [6] A. Adler, "Can images be regenerated from biometric templates?," in *Biometrics Consortium Conference*, (Arlington, VA), September 2003.
- [7] C. J. Hill, "Risk of masquerade arising from the storage of biometrics," B.S. Thesis, Australian National University, November 2001, <http://chris.fornax.net/biometrics.html>.
- [8] A. Ross, J. Shah, and A. K. Jain, "Towards reconstructing fingerprints from minutiae points," in *Proc. SPIE, Biometric Technology for Human Identification II*, vol. 5779, pp. 68–80, (Orlando, FL), March 2005.
- [9] R. Cappelli, R. Erol, D. Maio, and D. Maltoni, "Synthetic fingerprint-image generation," in *Proc. Int'l. Conf. Pattern Recognition (ICPR)*, vol. 3, pp. 475–478, (Barcelona, Spain), September 2000.
- [10] C. Soutar, "Biometric system security," White Paper, Bioscrypt, <http://www.bioscrypt.com>.
- [11] A. Adler, "Images can be regenerated from quantized biometric match score data," in *Proc. Canadian Conf. Electrical Computer Eng.*, pp. 469–472, (Niagara Falls, Canada), May 2004.
- [12] M. Yeung and S. Pankanti, "Verification watermarks on fingerprint recognition and retrieval," in *Proc. SPIE, Security and Watermarking of Multimedia Contents*, vol. 3657, pp. 66–78, (San Jose, USA), January 1999.
- [13] A. K. Jain and U. Uludag, "Hiding biometric data," *IEEE Trans. Pattern Anal. Mach. Intelligence*, vol. 25, no. 11, pp. 1493–1498, 2003.
- [14] L. C. Ferri, A. Mayerhofer, M. Frank, C. Vielhauer, and R. Steinmetz, "Biometric authentication for ID cards with hologram watermarks," in *Proc. SPIE, Security and Watermarking of Multimedia Contents IV*, vol. 4675, pp. 629–640, (Bellingham, WA), January 2002.
- [15] N. Ratha, J. Connell, and R. bolle, "Enhancing security and privacy in biometrics-based authentication systems," *IBM Systems Journal*, vol. 40, no. 3, pp. 614–634, 2001.
- [16] U. Uludag, S. Pankanti, S. Prabhakar, and A. K. Jain, "Biometric cryptosystems: issues and challenges," *Proceedings of the IEEE*, vol. 92, no. 6, pp. 948–960, 2004.
- [17] U. Uludag, S. Pankanti, and A. K. Jain, "Fuzzy vault for fingerprints," *To appear in Proc. Audio- and Video-based Biometric Person Authentication (AVBPA)*, (Rye Brook, NY), July 2005.