On the Feasibility of Creating Morphed Iris-Codes

C. Rathgeb and C. Busch
da/sec – Biometrics and Internet Security Research Group
Hochschule Darmstadt, Germany
{christian.rathgeb,christoph.busch}@h-da.de

Abstract

Morphing techniques can be used to create artificial biometric samples, which resemble the biometric information of two (or more) individuals in image and feature domain. If morphed biometric images or templates are infiltrated to a biometric recognition system the subjects contributing to the morphed image will both (or all) be successfully verified against a single enrolled template. Hence, the unique link between individuals and their biometric reference data is annulled. The vulnerability of face and fingerprint recognition systems to such morphing attacks has been assessed in the recent past.

In this paper we investigate the feasibility of morphing iris-codes. Two relevant attack scenarios are discussed and a scheme for morphing pairs of iris-codes depending on the expected stability of their bits is proposed. Different iris recognition systems, which accept comparison scores at a recommended Hamming distance of 0.32, are shown to be vulnerable to attacks based on the presented morphing technique.

1. Introduction

In past years researchers have pointed out diverse potential vulnerabilities of iris recognition. Proposed attacks, which aim at gaining unauthorized access to the system, can be coarsely categorized into presentation attacks and software-based attacks [20]. Presentation attacks refer to a presentation of an attack instrument to the iris camera with the goal of interfering with the operation of the iris recognition system [15]. Such attacks can be launched by using artificial attack presentation instruments, e.g. print outs, electronic displays, or even textured contact lenses [17]. For a review on iris-based presentation attacks and proposed detection mechanisms the reader is referred to [10]. To launch software attacks an attacker requires knowledge about the inner modules of the biometric system together with access to some of the system components, e.g. database, feature extractor or comparator. Software attacks include masquerade attacks, replay attacks, substitution attacks as well as overriding one of the inner modules of the system. For a review of iris-based software attacks we referred to [12].

More recently, attacks on face and fingerprint recognition systems based on morphed biometric images and templates have been presented. This new type of attack was introduced by Ferrara *et al*. [8]. Motivated by security gaps in the issuance process of electronic travel documents, the authors consider the scenario where an accomplice of a blacklisted subject presents a morphed face image at the time of enrolment. The issued travel document can then be used by both subjects to pass automated border control gates. The authors showed that commercial face recognition software tools are highly vulnerable to such attacks, i.e. different instances of images of either subject are successfully matched against the morphed image. In their experiments decision thresholds yielding a false match rate (FMR) of 0.1% have been used, according to the guidelines provided by the European Agency for the Management of Operational Cooperation at the External Borders (FRONTEX) [2]. In a further study the authors show that morphed face images are realistic enough to fool human examiners [9]. Scherhag *et al*. [22] demonstrated that presentation attack detection schemes employing general purpose texture descriptors used in conjunction with machine learning techniques are not capable of reliably detecting morphed face images. With respect to the above attack scenario it is stressed that a detection of morphed face images becomes even more challenging if images are printed and scanned. Ferrara *et al*. [7] also presented two different methods to morph fingerprints in image and feature domain. For a decision threshold yielding a FMR of 0.1% it is shown that commercial fingerprint recognition software tools are also highly vulnerable to such attacks. Given the fact that in the issuance process of electronic travel documents fingerprint enrolment is usually done live, the authors argue that manufactured fake fingertips may be presented to the scanner. Gomez-Barrero *et al*. [13] proposed a theoretical framework for measuring the vulnerability of biometric systems to attacks based on morphed images or template. The authors identified key factors

which take a major influence on a system's vulnerability to such attacks, e.g. the shape of genuine and impostor score distributions or the FMR the system is operated at. In contrast to the face or fingerprint recognition, iris recognition systems are operated at a much lower FMR (by many orders of magnitude) [5]. Hence, an analysis of the vulnerability of iris recognition systems to morphed iris-based feature vectors, i.e. iris-codes, is of particular interest.

In this work we present a technique to morph two iris-codes, which we referred to as stability-based bit substitution (SBS). The proposed scheme is applied to iris-codes extracted from iris images of the CASIAv4-Interval database using two different feature extraction methods. It is shown that the SBS technique outperforms naïve morphing schemes based on random bit substitution (RBS) and random row substitution (RRS). At a decision criterion recommended for iris-based verification, i.e. a Hamming distance of 0.32, which results in a FMR of about 0.0001% (1 in 1 million) in [4], attacks based on the SBS morphing scheme pose a serious risk. For more liberal decision thresholds resulting in FMRs of 0.001% and 0.01% success chances increase to about 50% and 85%, respectively. Visual inspections of morphed iris-codes suggest that it is hard to distinguish them from real iris-codes.

The remainder of this paper is organized as follows: Sect. 2 summarizes relevant attack scenarios. Two simple morphing schemes and the proposed approach are described in detail in Sect. 3. Experimental evaluations are presented in Sect. 4. Finally, conclusions are drawn in Sect. 5.

2. Attack Scenarios

2.1. Manipulation of Biometric Database

In this attack scenario we assume that an attacker has access to the database of the system in which iris-codes of legitimate subjects are stored. In addition, we assume that the attacker has knowledge about the software components used by the system to extract iris-codes. The attacker first extracts an iris-code from his eye. Then the attacker searches for a suitable counterpart in the database which he morphs with his own iris-code. Finally, he replaces the chosen iris-code by the generated morphed iris-code. From that point onwards, the attacker can gain unnoticed access to the system.

2.2. Presentation Attack at Enrolment

This scenario refers to the one considered by the aforementioned researchers. Again, knowledge about the software components used by the system to extract iris-codes is required. But in contrary to the scenario outlined in the previous subsection no access to the database of the system is required. Two attackers, a blacklisted subject and his accomplice, morph their iris-codes. Subsequently, an iris image is reconstructed from the morphed iris-code. Published works have demonstrated the feasibility of reconstructing realistic iris images from iris-codes [23, 11]. Based on the reconstructed iris image the accomplice launches a presentation attack during enrolment. Such an attack might also be feasible in a supervised enrolment scenario using a printed contact lens. From that point onwards, both subjects can gain access to the system (or share a single electronic travel document). A morphing of iris images might also be employed in this type of attack, which would not require a reconstruction of an image from the morphed iris-code.

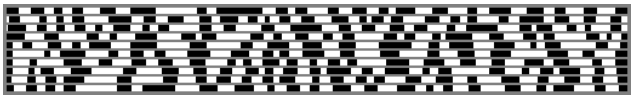

(a) CodeA

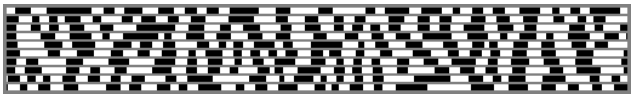

(b) CodeB

Figure 1. Examples of binary iris-codes extracted from the left iris images depicted in the top and bottom row of Fig. 4.

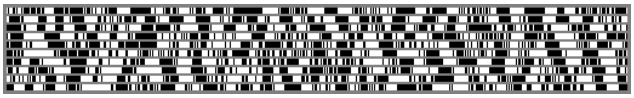

(a) CodeM - random bit substitution

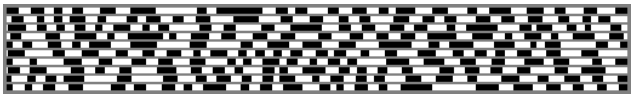

(b) CodeM - random row substitution

Figure 2. Naïve approaches: morphed iris-codes produced from the pair of iris-codes shown in Fig. 1.

3. Morphing Iris-Codes

Given a pair of iris-codes, CodeA and CodeB, and their corresponding noise masks, MaskA and MaskB, a morphed iris-code CodeM and a noise mask MaskM is created. Sample visualisations of iris-codes produced by different eyes are shown in Fig. 1. Let δ be the Hamming distance used as decision threshold in the attacked iris recognition system. CodeM should be created in a way that different instances of contributing iris-codes, CodeA' and CodeB', are accepted by the system, i.e. $HD(\text{CodeM}, \text{CodeA'}) < \delta$ and $HD(\text{CodeM}, \text{CodeB'}) < \delta$.

Prior to the actual morphing process an initial alignment of CodeA and CodeB is performed. Let $s(\text{CodeB}, k)$ denote CodeB circularly shifted by k bits. Circular bit shifts are applied at K different shifting positions and the optimal initial alignment is estimated as, $\min_{k \in K} HD(\text{CodeA}, s(\text{CodeB}, k))$. The masks are shif-

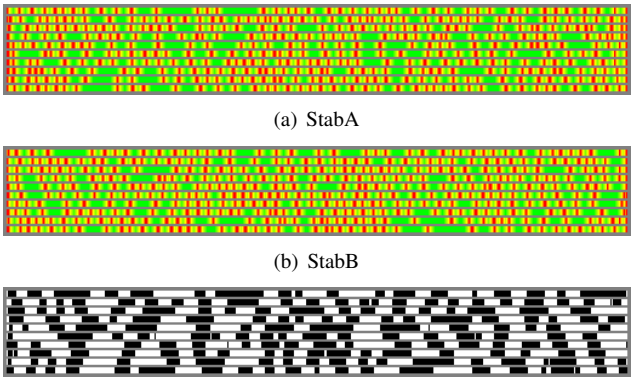

(a) StabA

(b) StabB

(c) CodeM - stability-based bit substitution

Figure 3. Proposed approach: visualisations of stability-codes (expected unstable bits are marked in red, stable bits are marked in green) and the morphed iris-code produced from the pair of iris-codes shown in Fig. 1.

ted accordingly and the morphed mask is constructed as the union of both masks, MaskM $=$ MaskA $\cup$ MaskB. Hence, when comparing CodeA' or CodeB' to CodeM it is ensured that the compared bits have been uncorrupted by eyelids, eyelashes, or other disturbance factors. In the following subsections we describe two naïve morphing schemes and the proposed approach.

3.1. Naïve Approaches

The first simple morphing scheme is referred to as random bit substitution (RBS). Matching bits between CodeA and CodeB are assigned to CodeM. For non-matching bit positions bits are randomly chosen and assigned to CodeM. An example of a resulting morphed iris-code is shown in Fig. 2 (a). Morphed iris-codes generated by this method might be easily detected. Compared to real iris-codes these are expected to exhibit an increased amount of transitions between 1-bit and 0-bit sequences.

The second simple morphing technique is referred to as random row substitution (RRS). In this scheme entire rows are randomly chosen from CodeA or CodeB and assigned to CodeM. It is ensured that the same amount of iris-code rows are chosen from both contributing iris-codes. An example of a resulting morphed iris-code is shown in Fig. 2 (b). A morphed iris-code generated by this method might be detected by analysing the correlation between adjacent rows of CodeM, which will be partly missing.

3.2. Proposed Approach

The proposed morphing approach is referred to as stability-based bit substitution (SBS). The aim of our approach is to assign bits to the morphed iris-code according to the expected stability of contributing iris-code bits. Assume that, at a non-matching bit position, the bit of CodeA expected to be highly stable while the bit of CodeB is expected to be rather unstable. Then it is preferable to assign the bit of CodeA to CodeM at this bit position, since bits at this position will likely flip in other instances of CodeB.

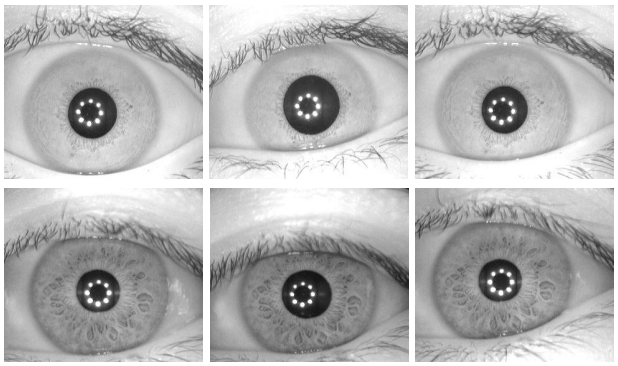

Figure 4. Sample images of two eyes (top and bottom row) of the CASIAv4-Interval iris database [3].

It is well known that iris-code bits are not mutually independent, also see Fig. 1. This is due to the internal spatial correlations within iris textures and nature of employed filters. Daugman [5] has recently shown that rows of iris-codes can be modelled as a "sticky oscillator" Markov process. Moreover, Hollingsworth *et al*. [14] have shown, that for ideal imaging (no eyelash/eyelid occlusions, corneal reflections, etc. on iris textures) so-called "fragile" bits, i.e. bits which exhibit a higher probability than others to flip their value during a genuine comparison, most likely occur between consecutive 1-bit and 0-bit sequences. Since filters employed in the feature extraction stage set iris-code bits by their sign, unstable bits correspond to coefficients close to zero.

Building upon these observations we estimate a stability-code StabA indicating the expected stability of each bit of a given iris-code CodeA. Note that, original filter responses might not be available to the attacker, e.g. in the database manipulation attack scenario. The presented approach bit stabilities are estimated based on a single iris-code. The availability of multiple samples of a pair of iris-codes to be morphed, will a allow a more precise estimation of bit stabilities [14]. The stability-code consists of n sequences of integer values, StabA$=(\mathbf{a}_1, \ldots, \mathbf{a}_n)$, defined by the consecutive 1-bit and 0-bit sequences of CodeA. For a sequence of length l, $\mathbf{a}_i=(a_1, \ldots, a_l)$, the expected stability at position j is estimated as $a_j = \min(j; l+1-j)$, e.g. a 1-bit or 0-bit sequence of length five in CodeA will result in $(1, 2, 3, 2, 1)$ in StabA. Hence, the expected stability decreases with the distance to the transition to the next 1-bit or 0-bit sequence. Example visualisations of stability-codes are depicted in Figs. 3 (a)-(b). Further note that, relative changes in stability across different iris-code regions can be ignored, since corresponding regions of CodeA and CodeB

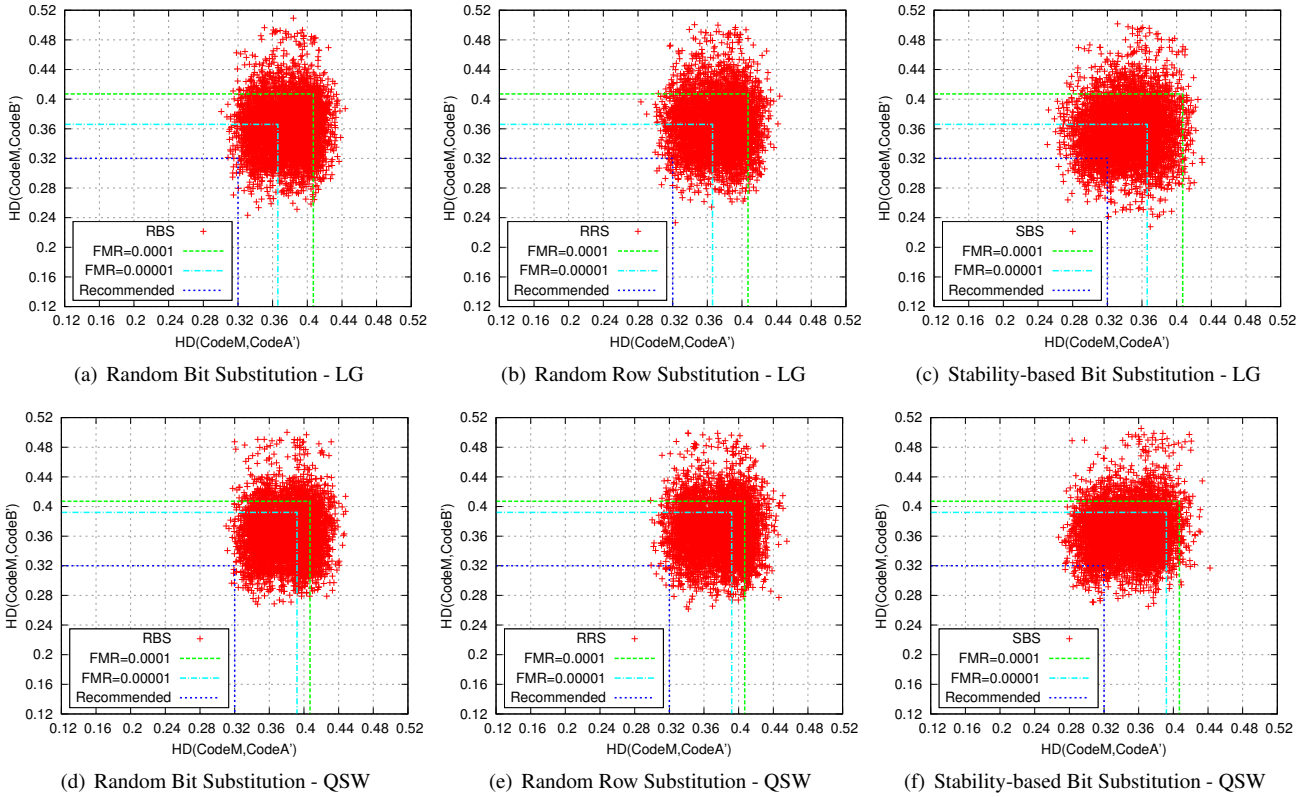

(a) Random Bit Substitution - LG (b) Random Row Substitution - LG (c) Stability-based Bit Substitution - LG

(d) Random Bit Substitution - QSW (e) Random Row Substitution - QSW (f) Stability-based Bit Substitution - QSW

Figure 5. Scatter plots obtained from the first 5,000 attack attempts: x and y values of each point represent two comparisons of different instances of CodeA and CodeB to the corresponding morphed iris-code CodeM, $HD(\text{CodeM}, \text{CodeA'})$ and $HD(\text{CodeM}, \text{CodeB'})$. Points where x and y values are both below a certain Hamming distance are successful attacks, with respect to the specified decision threshold.

are morphed. Matching bits between CodeA and CodeB are assigned to CodeM. At each non-matching bit position m, CodeM is defined as,

$$\text{CodeM}[m] = \begin{cases} \text{CodeA}[m], \text{ if } \text{StabA}[m] > \text{StabB}[m], \\ \text{CodeB}[m], \text{ otherwise.} \end{cases} \quad (1)$$

An example of a resulting morphed iris-code is shown in Fig. 3 (c). It can be observed that, in contrast to the RBS method, the proposed morphing technique produces realistic transitions between 1-bit or 0-bit sequences. Also, compared to the RRS scheme, correlations between iris-code rows are preserved. It is worth noting that, in [14, 5] it is recommended to mask out unstable bits during iris-code comparisons. In this case an attacker could assign 0s in stability codes at corresponding positions and un-mask bits in noise masks accordingly.

4. Experiments

4.1. Experimental Setup

Experimental evaluations are carried out on the CASIAv4-Interval iris database [3]. Sample images of the used dataset are depicted in Figure 4. Based on the observation that noise masks of different eyes tend to be more similar when they originate from the same eye position, we process only images of all 198 left eyes. A total number of 19,503 morphed iris-codes are created from pairs of the first image of each subject. In an attack attempt two different instances of iris-codes contributing to a morphed iris-code are compared against it. An attack attempt is considered successful if the larger of the two obtained Hamming distance scores is below the decision threshold, i.e. $\max(HD(\text{CodeM}, \text{CodeA'}), HD(\text{CodeM}, \text{CodeB'})) < \delta$. We consider decision thresholds at FMR of 0.01% and 0.001%, which are frequently reported by iris recognition researchers. Further, we consider a Hamming distance of 0.32 as decision criterion which was recommended in [4]. Depending on the number of remaining iris images up to five comparisons are performed against each morphed iris-code, resulting in 64,489 attack attempts. The attack success rate (ASR) is measured as the proportion of completed attack attempts that are successful.

In the employed iris recognition systems the iris of a given sample image is detected and transformed to a normalized rectangular texture of 512×64 pixels. In the feature extraction stage two conventional algorithms are employed where normalized enhanced iris textures are divi-

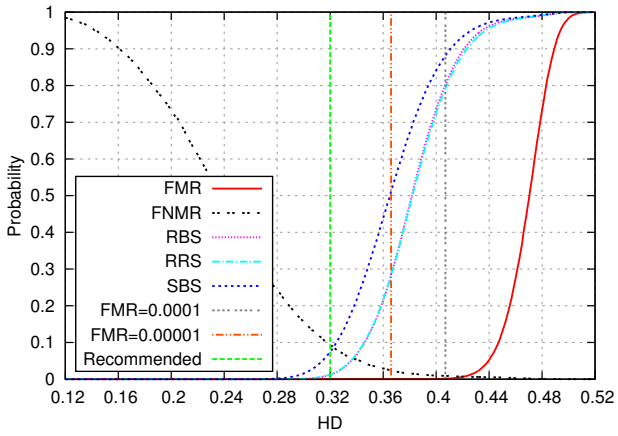

(a) ASR - LG

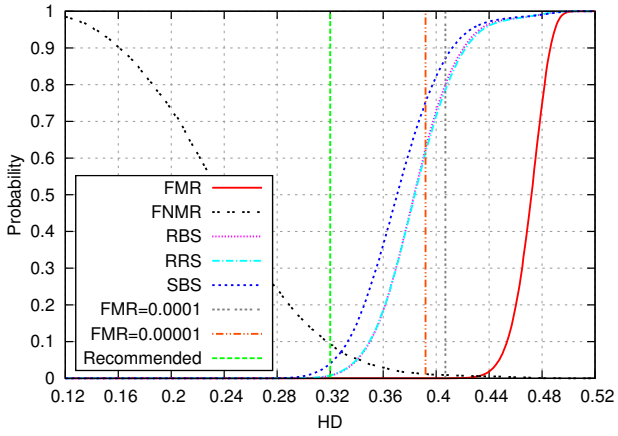

(b) ASR - QSW

Figure 6. Obtained attack success rates for attacks based on different morphing techniques for both iris recognition systems.

ded into stripes to obtain 10 one-dimensional signals, each one averaged from adjacent texture rows. The first feature extraction method is based on 1D-LogGabor wavelet [18] (LG) and the second follows the algorithm proposed by Ma *et al.* [16] (QSW) based on a quadratic spline wavelet transform. Extracted iris-codes are of size 512×10 bits for both algorithms. Custom implementations of employed segmentation and feature extractors are available in [1]. For further details on the employed feature extraction algorithms the reader is referred to [21].

In the initial alignment (prior to generating the morphed iris-code) ±4 bit shifts are applied and at the time of authentication ±12 bit shifts are applied.

4.2. Performance Evaluation

The scatter plots in Fig. 5 show some scores obtained for comparing different iris-code instances of subjects to their corresponding morphed iris-code. It can be observed $HD(\text{CodeM}, \text{CodeA'})$ is generally lower than $HD(\text{CodeM}, \text{CodeB'})$. Due to the fact that CodeB is circularly shifted in the initial alignment prior to creating CodeM, an appropriate alignment between CodeB' and CodeM is often not feasible. Obtained ASRs of attacks based on all morphing methods for both iris recognition systems are plotted in Fig. 6. Compared to the two simple morphing schemes attacks based on the proposed SBS method reveal the highest success chance for all considered thresholds, see Table 1. It is important to note that, in our experiments we generate morphed iris-codes from all possible combinations of iris-codes. In a real world scenario an attacker would search for suitable iris-code (or accomplice) before creating a morphed iris-code. We have identified iris-code pairs which exhibit a rather low Hamming distance as most suitable candidates. Two such examples are depicted in Fig. 7. Considering the size of the database obtained ASRs could be interpreted as alarmingly high even for a conservative decision threshold (depending on the attack scenario).

Table 1. Obtained attack success rates for attacks based on different morphing techniques for both iris recognition systems using different decison thresholds.

Decision threshold	**LG**			**QSW**		
	RBS	**RRS**	**SBS**	**RBS**	**RRS**	**SBS**
FMR $= 0.01\%$	80.20	79.12	88.11	80.01	78.55	87.07
FMR $= 0.001\%$	28.41	28.45	51.19	62.59	61.54	75.27
$HD = 0.32$	1.18	1.28	7.45	0.73	0.77	4.01

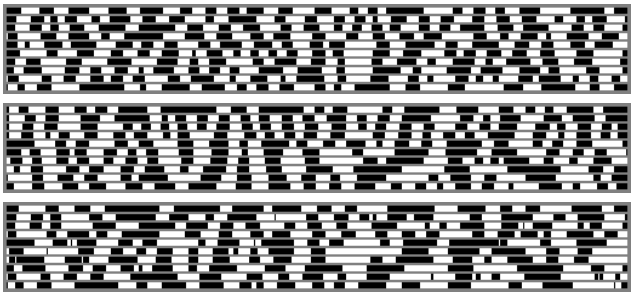

(a) $HD(\text{CodeA}, \text{CodeB}) = 0.471875$

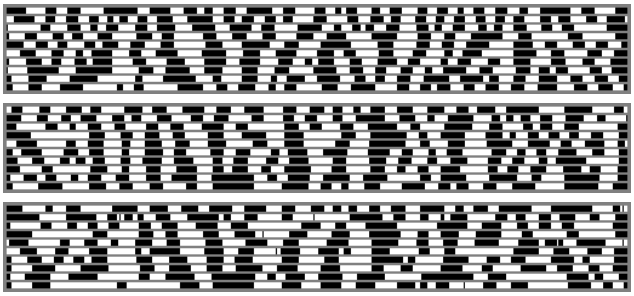

(b) $HD(\text{CodeA}, \text{CodeB}) = 0.455469$

Figure 7. Examples of pairs of original iris-codes (top) and the resulting morphed iris-code (bottom) for which attack attempts achieved scores clearly below a Hamming distance of 0.32. Low initial comparison scores between original iris-codes indicate a high success chance.

5. Conclusion

Operational deployments of iris recognition are operated at extremely low FMRs, e.g. FMR $= 10^{-6}$. Nonetheless, we have shown that iris recognition systems might still be vulnerable to attacks based on morphed iris-codes created by the presented SBS morphing approach. To assess whether a suitable counterpart, which maximizes the success chance of the proposed attack, can be found for any given iris-code, deeper analyses are required. Comparisons of iris-codes extracted from different eyes produce relatively constant dissimilarity scores [6]. Hence, we expect the ASR to be relatively equal for any given iris-code provided that a sufficiently large dataset of iris-codes is available.

The risk of the proposed attacks is expected to become even more serious for iris recognition applied in unconstrained environments or visible wavelengths, where less conservative decision thresholds are needed to achieve acceptable false non-match rates (FNMRs). Depending on the attack scenario such attacks might be prevented by template protection schemes [19], which permanently conceal original iris-codes, and/or robust presentation attack detection techniques [10, 17]. An automated detection of morphed iris-codes or a reconstruction of iris images from them might be subject to future studies.

Acknowledgement

This work was partially supported by the German Federal Ministry of Education and Research (BMBF) as well as by the Hessen State Ministry for Higher Education, Research and the Arts (HMWK) within the Center for Research in Security and Privacy (CRISP).

References

[1] USIT – University of Salzburg iris toolkit. `http://www.wavelab.at/sources/Rathgeb16a`. Version 2.0.x.

[2] FRONTEX – Research and Development Unit: Best practice technical guidelines for automated border control (ABC) systems, 2012. Version 2.0.

[3] Chinese Academy of Sciences' Institute of Automation. CASIA Iris Image Database V4.0 – Interval. `http://biometrics.idealtest.org`, 2010.

[4] J. Daugman. Probing the uniqueness and randomness of iris-codes: Results from 200 billion iris pair comparisons. *Proceedings of the IEEE*, 94(11):1927–1935, 2006.

[5] J. Daugman. Information theory and the iriscode. *Trans. on Information Forensics and Security*, 11(2):400–409, 2016.

[6] J. Daugman and C. Downing. Searching for doppelgängers: assessing the universality of the iriscode impostors distribution. *IET Biometrics*, 5(2):65–75, 2016.

[7] M. Ferrara, R. Cappelli, and D. Maltoni. On the feasibility of creating double-identity fingerprints. *IEEE Transactions on Information Forensics and Security*, 12(4):892–900, 2017.

[8] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. In *Proc. Int'l Joint Conf. on Biometrics (IJCB'14)*, pages 1–7, 2014.

[9] M. Ferrara, A. Franco, and D. Maltoni. On the effects of image alterations on face recognition accuracy. In T. Bourlai, editor, *Face Recognition Across the Imaging Spectrum*, pages 195–222. Springer, 2016.

[10] J. Galbally and M. Gomez-Barrero. A review of iris anti-spoofing. In *Proc. Int'l Workshop on Biometrics and Forensics (IWBF'16)*, pages 1–6, 2016.

[11] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fiérrez, and J. Ortega-Garcia. Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms. *Computer Vision and Image Understanding*, 117(10):1512–1525, 2013.

[12] M. Gomez-Barrero and J. Galbally. Software attacks on iris recognition systems. In C. Rathgeb and C. Busch, editors, *Iris and Periocular biometric recognition*, pages 291–316. IET, 2017.

[13] M. Gomez-Barrero, C. Rathgeb, U. Scherhag, and C. Busch. Is your biometric system robust to morphing attacks? In *Proc. Int'l Workshop on Biometrics and Forensics (IWBF'17)*, pages 1–6, 2016.

[14] K. P. Hollingsworth, K. W. Bowyer, and P. J. Flynn. The best bits in an iris code. *Trans. on Pattern Analysis and Machine Intelligence*, 31(6):964–973, 2009.

[15] ISO/IEC TC JTC1 SC37 Biometrics. *ISO/IEC IS 30107-1. Information Technology – Biometrics presentation attack detection – Part 1: Framework*, Mar. 2016.

[16] L. Ma, T. Tan, Y. Wang, and D. Zhang. Efficient iris recognition by characterizing key local variations. *Trans. on Image Processing*, 13(6):739–750, 2004.

[17] S. Marcel, M. Nixon, and S. Z. Li. *Handbook of Biometric Anti-Spoofing*. Springer-Verlag New York, Inc., 2014.

[18] L. Masek. Recognition of human iris patterns for biometric identification. Master's thesis, University of Western Australia, 2003.

[19] K. Nandakumar and A. K. Jain. Biometric template protection: Bridging the performance gap between theory and practice. *IEEE Signal Processing Magazine - Special Issue on Biometric Security and Privacy*, pages 1–12, 2015.

[20] N. K. Ratha, J. H. Connell, and R. M. Bolle. Enhancing security and privacy in biometrics-based authentication systems. *IBM Systems Journal*, 40(3):614–634, 2001.

[21] C. Rathgeb, A. Uhl, and P. Wild. *Iris Recognition: From Segmentation to Template Security*, volume 59 of *Advances in Information Security*. Springer Verlag, 2013.

[22] U. Scherhag, R. Ramachandra, K. B. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch. On the vulnerability of face recognition systems towards morphed face attacks. In *Proc. Int'l Workshop on Biometrics and Forensics (IWBF'17)*, pages 1–6, 2017.

[23] S. Venugopalan and M. Savvides. How to generate spoofed irises from an iris code template. *IEEE Transactions on Information Forensics and Security*, 6(2):385–395, 2011.