Is Your Biometric System Robust to Morphing Attacks?

Marta Gomez-Barrero, Christian Rathgeb, Ulrich Scherhag, Christoph Busch
da/sec - Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany
Email: {marta.gomez-barrero,christian.rathgeb,ulrich.scherhag,christoph.busch}@h-da.de

***Abstract*—The wide deployment of biometric recognition systems has raised several concerns regarding their security. Among other threats, morphing attacks consist of the infiltration of artificial images created using biometric information of two or more subjects. These morphed images are hence positively matched to several subjects. Recent studies have shown that such images pose a concrete threat to civil security: wanted criminal offenders can use an authentic passport to enter a country with a false identity. However, there is still no quantitative manner to analyse this threat. We address this shortcoming by proposing a new framework for the evaluation of the vulnerability of biometric systems to morphing attacks. The experimental analysis on real systems based on face, iris and fingerprint shows that even systems providing high verification accuracy are vulnerable to this kind of attacks, depending on the verification threshold and the shape of the mated and non-mated score distributions.**

I. Introduction

Image morphing has been an active area of research since the 80s [1], [2]. For instance, in the film industry a mesh warping technique designed at Industrial Light & Magic [3] appeared in 1988 on *Willow* and in 1989 on *Indiana Jones and the Last Crusade*. Similarly, in the music industry, morphing techniques were used as early as 1989 in the cover for Queen's album *The Miracle*, and two years later in Michael Jackson's music video *Black or White*. Similarly, a wide deployment of biometric recognition systems has been carried out in the last decade, both for large-scale national and international initiatives (e.g., the Indian Unique ID[1] or the SmartBorders package[2]), and for specific applications such as automatic border crossing[3] or banking[4]. In spite of those facts, it has not been until very recently that the impact on the security of such systems caused by photo alterations, and morphing in particular, has been analysed [4], [5], [6], [7], [8].

Regarding the wider field of research of attacks on biometric systems, within the ISO/IEC IS 30107 on biometric presentation attack detection [9], *Presentation Attacks* are defined as the "presentation to the biometric data capture subsystem with the goal of interfering with the operation of the biometric system". Among other possibilities, an eventual attacker may aim at manipulating the sample presented at enrolment, for instance using a morphed sample which allows the positive verification of two different individuals. Such attacks will be referred to as *morphing attacks* for the remainder of the article. Fig. 1 shows the diagram of such an attack, which is carried out in the following three steps:

- The attacker finds an accomplice, whose biometric characteristic ($\mathbf{M}_{ac}$) is similar enough to his own ($\mathbf{M}_{at}$) to allow a successful morphing and eventual verification of both subjects.
- A morphed sample $\mathbf{M}_{morph}$ is created from the original unaltered samples of the accomplice and the attacker, $\mathbf{M}_{ac}$ and $\mathbf{M}_{at}$.
- The morphed sample $\mathbf{M}_{morph}$ is presented to the system, and its corresponding template $\mathbf{T}_{morph}$ is enroled in the database.

Later on, both the attacker and the accomplice can present their unaltered biometric characteristics to the biometric system, the samples $\mathbf{M}'_{at}$ and $\mathbf{M}'_{ac}$ are captured, which will

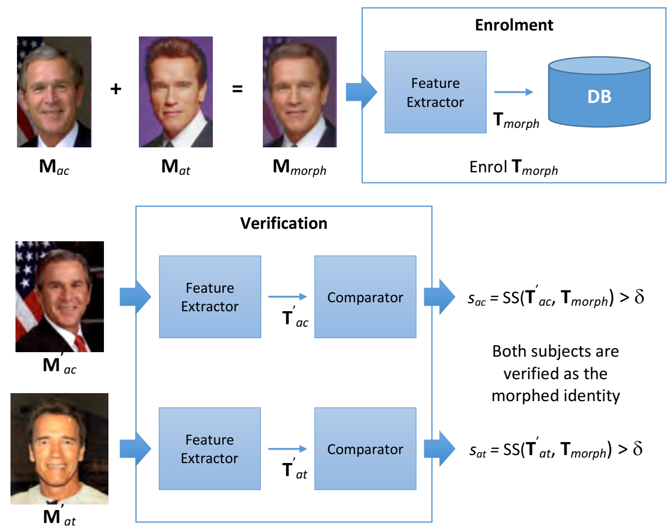

Fig. 1: General diagram of a morphing attack: at enrolment, the attacker and accomplice samples, $\mathbf{M}_{at}$ and $\mathbf{M}_{ac}$, are morphed into a single sample $\mathbf{M}_{morph}$, whose corresponding template $\mathbf{T}_{morph}$ is enroled in the database. During verification, when either probe sample $\mathbf{M}'_{at}$ and $\mathbf{M}'_{ac}$ is presented to the system, their corresponding templates achieve similarity scores with respect to the morphed template, $s_{ac} = SS(\mathbf{T}'_{ac}, \mathbf{T}_{morph})$ (resp. at), higher than the verification threshold δ.

[1] https://uidai.gov.in/

[2] http://ec.europa.eu/dgs/home-affairs/what-we-do/policies/borders-and-visas/smart-borders/index_en.htm

[3] http://www.easypass.de/EasyPass/EN/What_is_EasyPASS/home_node.html

[4] http://www.cnet.com/news/mastercard-app-will-let-you-pay-for-things-with-a-selfie/

yield the corresponding templates $\mathbf{T}'_{at}$ and $\mathbf{T}'_{ac}$. The attack will be successful if both templates obtain similarity scores with respect to the enroled morphed template higher than the verification threshold, δ:

$$s_{at} = SS\left(\mathbf{T}'_{at}, \mathbf{T}_{morph}\right) > \delta \tag{1}$$

$$s_{ac} = SS\left(\mathbf{T}'_{ac}, \mathbf{T}_{morph}\right) > \delta \tag{2}$$

where SS outputs the similarity score between two templates.

Within the recently finished FIDELITY EU project [10], some threats to the concept of secure biometric passports stemming from these attacks have been unveiled. In 2014 it was shown in [5] that morphing attacks on face verification systems are possible due to the similarity of the morphed image to both subjects, as depicted in Fig. 1. Furthermore, it is shown in that work that not only automatic face recognition algorithms, but also human supervisors, can be fooled by morphed facial images. As a consequence, the work developed in [5] results in a concrete threat to civil security: wanted criminal offenders, for example terrorists, can use an authentic passport, complying with all physical safety features, to enter a country with the identity of an accomplice, when performing three basic steps: $i)$ find a rather lookalike accomplice, $ii)$ morph passport photos of both, possibly utilizing free software available on the internet, such as the GNU Image Manipulation Program (GIMP) and the GIMP Animation Package (GAP) tools used in [5], [7], and $iii)$ the accomplice applies for a passport. The passport manufacturer will issue an authentic passport, which can be used to enter a country by both subjects, the accomplice and the criminal offender. We may hence conclude that morphed images pose a real and significant threat to biometric recognition systems, especially for Automated Border Control (ABC) systems.

More recently, in 2016, a detection algorithm for morphed face images based on Binarized Statistical Image Features (BSIF) and Support Vector Machines (SVM) was presented in [7], which is able to achieve an Average Classification Error Rate (ACER) as low as 1.73%.

Even if the aforementioned articles only consider digital morphed images, during the normal procedure for passport issuance, the digital image is printed, presented at the issuance office and scanned. In order to conduct a more realistic analysis, and based on a extended version of the database generated in [7], the vulnerability of both a commercial and a freely available systems to printed and scanned morphed images is analysed in [8]. It is shown that Bona Fide Presentation Classification Error Rate (BPCER) of the scanned images at a fixed Attack Presentation Classification Error Rate (APCER) is increased three to five times with respect to that of the digital samples. As a consequence, more research needs still to be carried out in this direction in order to detect morphed samples not only for face but also for other biometric characteristics.

In fact, probably due to the fact that the face has been selected by the International Civil Aviation Organization (ICAO) as the primary identifier for electronic Machine Readable Travel Documents (eMRTD), so far the impact of morphed samples on biometric systems has been studied only for that characteristic [5], [7]. Therefore, the following questions remain unanswered:

- What about other characteristics, such as fingerprint, also considered in ABC systems [11]? Is it possible to launch similar attacks?
- For a given system, what is the impact of morphing attacks for different operating points (i.e., verification thresholds, δ)?
- How similar should $\mathbf{M}_{at}$ and $\mathbf{M}_{ac}$ be, in order to allow a successful attack?
- What is the relationship between the shape of the mated and non-mated score distributions and the success chances of a morphing attack?
- What is the appropriate value of δ in terms of robustness to morphing attacks and low False Non-Match Rate (FNMR)?

In the present article we aim at answering these questions. To that end, we propose a general framework to assess the feasibility of creating morphed samples and to estimate the success chances of morphing attacks (Fig. 1). This evaluation framework only requires the computations of the mated and non-mated scores of unaltered biometric samples, which is always necessary to fix the verification threshold of the system. In addition, we evaluate three real independent systems, based on different characteristics (i.e., face, iris and fingerprint), to show the generality of the proposed framework and to analyse the impact of morphing attacks on different biometric modalities.

The rest of the paper is organised as follows. Sect. II describes the framework for evaluating the vulnerability of biometric systems to morphing attacks. Then, real empirical examples are given in Sect. III for face, iris and fingerprint, and final conclusions are drawn in Sect. IV.

II. Framework for Measuring the Feasibility of Morphing Attacks

In order to assess the feasibility of carrying out morphing attacks such as the ones described in Sect. I and Fig. 1, we have to answer the following question: what is the probability, denoted as P_{morph}, that the attacker is successful in his attempt? Or in other words, what is the probability of $s_{at} > \delta$? We thus want to compute

$$P_{morph} = p\left(s_{at} > \delta\right) \tag{3}$$

To answer this question and extend formality to the problem being addressed, some notations are introduced in this section. Throughout the article we will use the Harmonized Biometric Vocabulary (HBV) defined in the ISO/IEC 2382-37 [12]. For any clarification on the concepts used, we refer the reader to the mentioned standard. Given that they are often used throughout the article, for the sake of clarity, we will only include here the next definitions:

- *Biometric characteristic*: "biological and behavioural characteristic of an individual from which distinguishing,

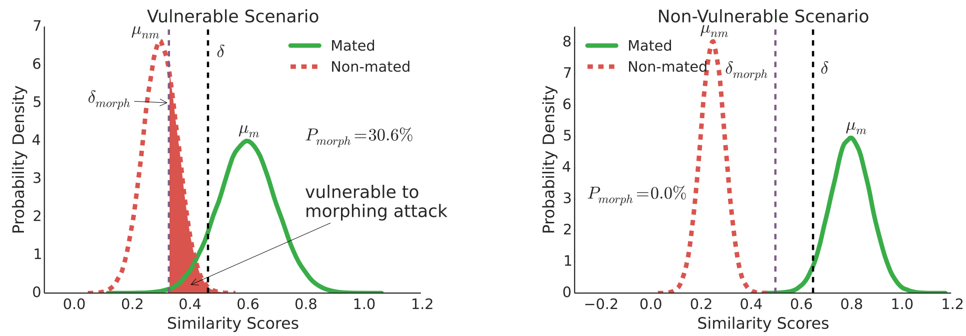

(a) (b)

Fig. 2: Examples of *Mated samples* (solid green) and *Non-mated samples* (dashed red) distibutions for (a) vulnerable and (b) non-vulnerable systems. The probability of success of a morphing attack for the depicted distributions is given by P_{morph}, which corresponds to the shaded area below $p\left(s|H_{nm}\right)$.

repeatable biometric features can be extracted for the purpose of biometric recognition". For example, a fingerprint or an iris are two different biometric characteristics.

- *Biometric instance*: for some characteristics, an individual possesses several instances. For example, the right index fingerprint is a different instance from the left thumb, even if they serve to identify the same person.
- *Mated samples*: "paired biometric probe and biometric reference that are from the same biometric characteristic of the same biometric data subject". For example, two fingerprint samples from the same right index finger.
- *Non-mated samples*: "paired biometric probe and biometric reference that are not from the same biometric instance". For example, two fingerprint samples from different fingers.

In general, depending on the samples compared, two different types of similarity scores are possible within a biometric system: those obtained from the comparison of mated samples, and those yielded by comparisons of non-mated samples. Let us accordingly define the corresponding types of score distributions, where $s = SS\left(\mathbf{T}_1, \mathbf{T}_2\right)$ is the similarity score between two templates, as illustrated in Fig. 1:

- *Mated samples* distribution: scores computed from templates extracted from different samples of a single instance of the same subject. It represents the conditional probability of obtaining a score s knowing that two templates come from mated instances, that is, $p\left(s|H_m\right)$, where $H_m = \{$both templates stem from mated samples$\}$.
- *Non-mated samples* distribution: scores yielded by templates generated from samples of different instances. It represents the conditional probability of obtaining a score s knowing that two templates come from non-mated instances, that is, $p\left(s|H_{nm}\right)$, where $H_{nm} = \{$both templates stem from non-mated samples$\}$.

Two examples of the probability density functions of such distributions are shown in Fig. 2, where the *Non-mated* samples distribution, $p\left(s|H_{nm}\right)$, is depicted in dashed red, and the *Mated* samples distribution, $p\left(s|H_m\right)$, in solid green, their corresponding mean values are denoted as μ_m and μ_{nm}, respectively, and the verification threshold δ is represented with a vertical black dashed line.

In the remainder of the section, we assume that $s_{at} \approx s_{ac}$ (as defined in Eqs. 1 and 2), and refer to any of the scores as s_{at}, which hence denotes the scores obtained from the comparisons of the attacker or the accomplice unaltered samples with the morphed template. It should be thus noted that P_{morph} evaluates the average success chances of the morphing attack, since we have assumed that $s_{ac} \approx s_{at}$. In practice, one of the scores can be higher than the other one (i.e., $s_{at} \neq s_{ac}$), thereby increasing or decreasing the success chances for the attacker.

Now, back to the formal definition of the morphing attack, it should be noted that we are interested on where s_{at} lies with respect to the verification threshold δ. Since it stems from the comparison of samples belonging to non-mated instances (i.e., the attacker or the accomplice, and the morphed sample, which represents a third instance), it will belong to the *Non-mated* samples distribution. However, it is more probable that s_{at} lies on the right tail of the *Non-mated* samples distribution, between the mean values of both score distributions, μ_m and μ_{nm}. The reason behind this fact is that the reference template $\mathbf{T}_{morph}$ is extracted from $\mathbf{M}_{morph}$, which is ultimately a combination of $\mathbf{M}_{at}$ and $\mathbf{M}_{ac}$ that was created to allow a positive verification of both subjects (see Eqs. 1 and 2). Therefore, $\mathbf{T}_{morph}$ lies between both unaltered samples in the n-dimensional space of the biometric templates, and, due to the assumption of $s_{at} \approx s_{ac}$, it is expected to lie on the average of the *Mated* and *Non-mated* scores.

More specifically, for a given accomplice whose characteristic yields a non-mated similarity score s_{nm} with respect to the attacker:

$$s_{nm} = SS\left(\mathbf{T}'_{ac}, \mathbf{T}'_{at}\right) \tag{4}$$

the expected value of s_{at}, denoted μ_{at}, can be estimated as:

$$\mu_{at} = E\left(s_{at}\right) = E\left(\frac{s_{nm} + s_m}{2}\right) = \frac{s_{nm} + \mu_m}{2} \tag{5}$$

where $s_m = SS\left(\mathbf{T}'_{morph}, \mathbf{T}_{morph}\right)$ represents a mated score, and hence has an expected value of μ_m.

In order for the morphing attack to be successful, μ_{at} must lie above the verification threshold δ; otherwise, the identity claim would be rejected and the attacker would have failed in his goal of being recognized with the enrolled morphed template $\mathbf{T}_{morph}$. Therefore, the probability of success of the morphing attack, as defined in Eq. 3, ultimately depends on the chances of obtaining an accomplice for which μ_{at} lies above the verification threshold δ. Which in turn depends on the score yielded by the accomplice with respect to the attacker, s_{nm}:

$$\begin{aligned} P_{morph} &= P\left(\mu_{at} > \delta\right) = P\left(\frac{s_{nm} + \mu_m}{2} > \delta\right) \\ &= P\left(s_{nm} > 2\delta - \mu_m\right) \end{aligned} \tag{6}$$

Denoting

$$\delta_{morph} = 2\delta - \mu_m \tag{7}$$

we can finally compute P_{morph} as follows:

$$P_{morph} = \int_{s \geq \delta_{morph}} p\left(s|H_{nm}\right) \mathrm{d}s \tag{8}$$

In Fig. 2, δ_{morph} is depicted with a purple vertical dashed line, and the area for which $s_{nm} > \delta_{morph}$, thereby granting success in the morphing attack, is shaded in red. This area represents the success probability P_{morph}.

We may observe in Fig. 2 two different scenarios. On the one hand, on Fig. 2a, for the defined threshold δ, we can see that the system is vulnerable to a morphing attack, being $P_{morph} = 30.6\%$. This means that, among all the possible accomplices used to compute the non-mated scores, 30.6% of them will yield morphed samples that allow a positive verification of both the attacker and the accomplice. The attacker would be consequently allowed to succeed in his goal.

On the other hand, on Fig. 2b, the verification threshold δ lies closer to the mean mated score, μ_m, and further from the *Non-mated* samples distribution. Since δ_{morph} only depends on the distance between the verification threshold δ and μ_m (see Eq. 7), and this distance is small compared to the distance between the *Mated* and *Non-mated* samples distributions, in this case δ_{morph} lies to the right of the *Non-mated* samples distribution. As a consequence, none of the non-mated scores s_{nm} is high enough to allow the attacker to succeed, thereby leading to $P_{morph} = 0\%$. In other words, the system is not vulnerable to morphing attacks under the selected verification threshold δ.

III. Experimental Evaluation

To show the generality of the proposed framework, three real systems (i.e., contrary to the simulated distributions plotted in Fig. 2) based on different biometric characteristics, using different features and comparators, will be analysed with the framework proposed in Sect. II:

- **Face verification**: the Log-Gabor Binary Pattern Histograms Sequences algorithm proposed in [13] is used. In particular, experiments are run on a publicly available implementation within the FaceRecLib[5] [14] and the Bob Toolbox. Similarity scores are computed based on histograms intersections. Experiments are carried out on the face subcorpus included in the Desktop Dataset of the Multimodal BioSecure Database[6] [15], which comprises 840 frontal face images from 210 subjects.
- **Iris verification**: we use the implementation of the dyadic wavelet based algorithm proposed by Ma *et al.* [16] within the publicly available University of Salzburg Iris Toolkit v1.0[7] [17]. Similarity scores are computed in terms of the Hamming Distance. Experiments are carried out on the IITD Iris Database version 1.0[8], which comprises 1,120 NIR images from 224 different subjects.
- **Fingerprint verification**: we have selected the Finger-Code scheme presented in [18], in which the final template comprises the standard deviations of the grey values of each sector for a set of Gabor based filters. From the original 640 features, a subset of the best performing 100 has been selected with the method proposed in [19], and similarity scores are computed in terms of the Euclidean distance. Experiments are carried out on fingerprint subcorpus of the BiosecurID database [20], considering only the right index acquired with the optical sensor (6,400 samples from 400 instances).

The corresponding *Mated* (solid green) and *Non-mated* samples distribution for each system, as well as P_{morph} (see Eq. 8), are depicted in Fig. 3. On the top row, the verification threshold δ corresponding to a False Match Rate (FMR) of 0.1% (as recommended by Frontex [11]) is analysed, whereas in the centre other operating points are studied. In both cases, δ is depicted with a black dashed line and the morphing threshold δ_{morph} (see Eq. 7) is plotted in purple. In addition, P_{morph} is plotted against the FNMR in log scale in the bottom row. The numerical analysis of the distributions is included in Table I, together with the success probability of the morphing attack P_{morph}, the difference $|\delta - \mu_{nm}|$ and all the intermediate values required for the computations.

We may observe in Fig. 3 that there is not a direct relationship between P_{morph} and the accuracy of the biometric system. In other words, a higher accuracy does not imply more robustness, nor vice versa. In the systems analysed, for the operating points corresponding to FMR = 0.1%, the most accurate system is the iris based (FNMR = 0.58%, see Table I), then the fingerprint system (FNMR = 7.8%) and the least accurate the face based (FNMR = 19.8%). However, the morphing attack has the highest probability of success for the iris system ($P_{morph} = 99.97\%$) and the lowest for the fingerprint system ($P_{morph} = 2.83\%$), reaching an intermediate value for the face based system ($P_{morph} = 44.56\%$).

On the other hand, for a given system, the higher the FMR, the more vulnerable the system is to morphing attacks. This

[5] https://pypi.python.org/pypi/facereclib
[6] http://biosecure.it-sudparis.eu/AB/
[7] http://www.wavelab.at/sources/
[8] http://www4.comp.polyu.edu.hk/~csajaykr/IITD/Database_Iris.htm

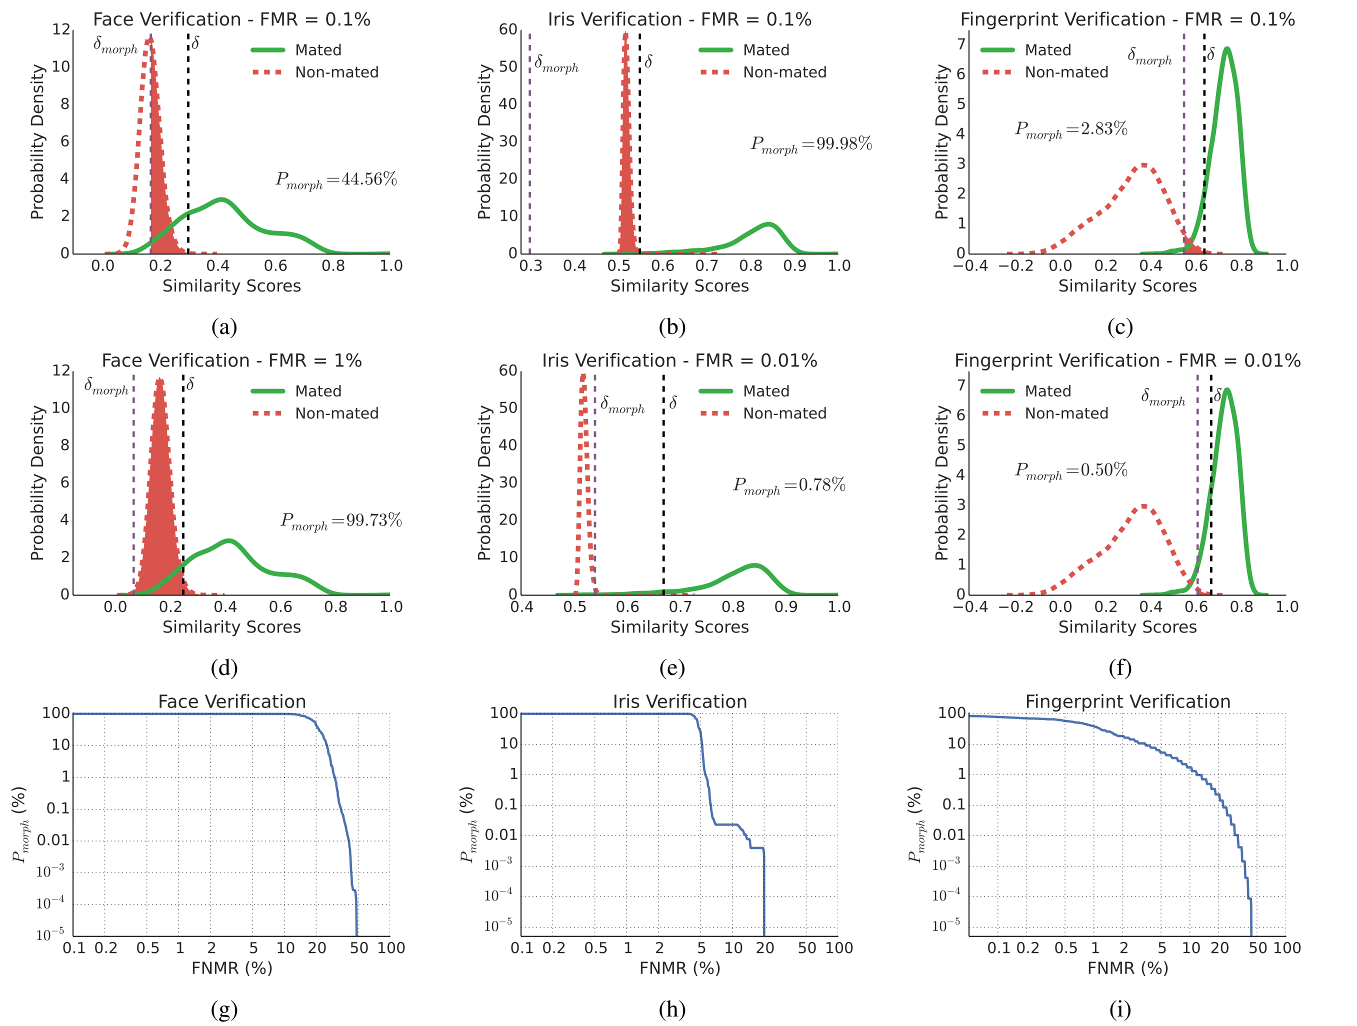

Fig. 3: *Mated samples* (solid green) and *Non-mated samples* (dashed red) distributions for the real systems based on face (left), iris (centre) and fingerprint (right), for different operation points at FMR = 0.1% (top) and other selected points (center). In all cases, the verification threshold δ is depicted with a black dashed line, the morphing threshold δ_{morph} is plotted in purple and the value of P_{morph}, which corresponds to the shaded area below $p(s|H_{nm})$, is also included. In the bottom row P_{morph} vs FNMR is depicted in log scale.

can be observed for the three systems analysed, which show higher values of P_{morph} for the operating point with the highest FMR (Figs. 3d, 3b and 3c). This is due to the fact that, for higher FMRs, δ is lower. As a consequence, it is more probable that that $s_{nm} > \delta$, and hence the attacker will easily find an appropriate accomplice.

Regarding the shape of the distributions, if there is a big overlap between the *Mated* and *Non-mated* distributions (e.g., the face system analysed), low values for P_{morph} are achieved at the cost of high FNMRs: as it may be observed in Fig. 3g, $P_{morph} < 1\%$ leads to FNMR $> 30\%$, whereas for iris it leads to FNMR $> 5\%$ and for fingerprint to FNMR $> 10\%$. An appropriate value of δ should then be chosen in order to minimise P_{morph} and at the same time achieve a low FNMR, which will enhance the usability of the system.

Furthermore, systems are more robust to morphing attacks when there is a big difference between δ and μ_{nm}: $|\delta - \mu_{nm}|$. The reason behind this fact is that, for a large difference $|\delta - \mu_{nm}|$, δ_{morph} will still lie far from μ_{nm}, hence leading to a small number of appropriate accomplices to succeed in the attack. On the other hand, if $|\delta - \mu_{nm}|$ is small, most of the *Non-mated* scores will be higher than δ_{morph}, thereby granting success to the attacker.

In addition, the shape of the *Non-mated* scores distribution also plays an important role: for a small σ_{nm}, and hence a very sharp distribution (e.g., iris system), small changes in δ and δ_{morph} will lead to big changes in P_{morph}, as it may be observed in Fig. 3h for FNMR $\in [5\%, 20\%]$. This is due to the fact that most of the *Non-mated* scores have very similar values, and therefore a small change in δ will lead to either none or most of them being higher than δ_{morph}. On the other hand, for a big σ_{nm} (e.g., fingerprint system), the decrease of P_{morph} is more gradual (see Fig. 3i).

TABLE I: Numerical evaluation of the distributions depicted in Fig. 3, including the corresponding mean (μ) and standard deviation (σ), the operating point analysed in terms of FMR and FNMR, the corresponding verification δ and morphing δ_{morph} (see Eq. 7) thresholds, as well as the probability of success of a morphing attack P_{morph} (see Eq. 8).

	μ_m	σ_m	μ_{nm}	σ_{nm}	FMR	FNMR	δ	$\lvert\delta - \mu_{nm}\rvert$	δ_{morph}	P_{morph}
Face	0.43	0.14	0.17	0.04	1.0%	11.1%	0.25	0.08	0.07	99.73%
					0.1%	19.8%	0.30	0.13	0.17	44.56%
Iris	0.80	0.07	0.52	0.007	0.1%	0.58%	0.55	0.02	0.54	99.97%
					0.01%	5.33%	0.67	0.15	0.30	0.78%
Fingerprint	0.73	0.06	0.32	0.14	0.1%	7.8%	0.64	0.29	0.61	2.83%
					0.01%	14.87%	0.67	0.34	0.55	0.49%

Related to the aforementioned facts, even if a verification system is not robust to morphing attacks for a given operating point (e.g., FMR = 0.1% for the iris system analysed), such robustness can be achieved for lower FMRs (e.g., FMR = 0.01%). In that case, δ is far enough from μ_{nm} with respect to σ_{nm} ($|\delta - \mu_{nm}| = 0.15$ and $\sigma_{nm} = 0.007$ for the iris system, see Table I), and hence δ_{morph} lies to the right and far from μ_{nm} (see Fig. 3e). Consequently, almost none of the samples are close enough to the attacker to allow a positive verification with respect to the morphed template, and the probability of success of the morphing attack drops ($P_{morph} = 0.78\%$).

IV. Conclusions

We have proposed a new framework to evaluate the vulnerability of biometric systems to the so-called morphing attacks, regardless of the biometric characteristic on which they rely for recognition. In order to give an estimation of the success chances of the attack, and accordingly choose an appropriate value for δ, only mated and non-mated scores for the corresponding biometric system need to be computed.

The experimental evaluation carried out on three different real systems, based on face, iris and fingerprint, confirms the fact that not only face based systems are vulnerable to morphing attacks. In fact, even very accurate systems (e.g., iris based) can be fooled with morphed samples if the appropriate verification threshold is not chosen (e.g., for iris $P_{morph} > 99\%$ for FMR = 0.1%, whereas $P_{morph} < 1\%$ for FMR = 0.01%). In particular, we can conclude that two facts play an important role in the evaluation of attacks carried out with morphed images. On the one hand, the decision threshold, δ: assuming the system outputs similarity scores, the lower it is, the higher the success chances of the attacker. On the other hand, the difference between δ and the mean of the *Non-mated* samples distribution, $|\delta - \mu_{nm}|$: the smaller the difference, the most likely it is that the attacker will succeed (i.e., higher P_{morph}).

As a consequence, we need to analyse the score distributions, and their relationship with the verification threshold δ, in order to give an estimation of the vulnerability of a particular system to this kind of attacks and hence choose an appropriate operating point.

Acknowledgments

This work was 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, www.crisp-da.de).

References

[1] G. Wolberg, "Image morphing: a survey," *The visual computer*, vol. 14, no. 8, pp. 360–372, 1998.
[2] A. Patel and P. Lapsiwala, "Image morphing algorithm: A survey," *Int. Journal of Computer App.*, vol. 5, no. 3, pp. 156–160, 2015.
[3] D. B. Smythe, "A two-pass mesh warping algorithm for object transformation and image interpolation," ILM Computer Graphics Department, Lucasfilm, Tech. Rep., 1990.
[4] M. Ferrara, A. Franco *et al.*, "On the impact of alterations on face photo recognition accuracy," in *Proc. ICIAP*. Springer, 2013, pp. 743–751.
[5] M. Ferrara, A. Franco, and D. Maltoni, "The magic passport," in *Proc. IJCB*. IEEE, 2014, pp. 1–7.
[6] ——, *Face Recognition across the imaging spectrum*. Springer, 2016, ch. On the Effects of Image Alterations on Face Recognition Accuracy, pp. 195–222.
[7] R. Ramachandra, K. B. Raja, and C. Busch, "Detecting morphed face images," in *Proc. BTAS*, 2016.
[8] U. Scherhag, M. Gomez-Barrero, C. Rathgeb, and C. Busch, "On the vulnerability of face recognition systems towards morphed face attacks," in *Proc. Int. Workshop on Biometrics and Forensics (IWBF)*, 2017.
[9] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 30107-1:2016, IT – Biometric presentation attack detection – Part 1: Framework*.
[10] FIDELITY, "Fast and trustworthy identity delivery and check with epassports leveraging traveler privacy," 2012.
[11] Frontex, "Best practice technical guidelines for automated border control (abc) systems," 2016.
[12] ISO/IEC TC JTC1 SC37 Biometrics, *ISO/IEC 2382-37:2012 IT – Vocabulary – Part 37: Biometrics*, ISO and IEC, 2012.
[13] W. Zhang, S. Shan *et al.*, "Local gabor binary pattern histogram sequence (LGBPHS): a novel non-statistical model for face representation and recognition," in *Proc. ICCV*, vol. 1, 2005, pp. 786–791.
[14] M. Günther, R. Wallace, and S. Marcel, "An open source framework for standardized comparisons of face recognition algorithms," in *Proc. ECCV*, ser. LNCS, vol. 7585, 2012, pp. 547–556.
[15] J. Ortega-Garcia, J. Fierrez *et al.*, "The multi-scenario multi-environment BioSecure multimodal database (BMDB)," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 32, pp. 1097–1111, 2010.
[16] L. Ma, T. Tan *et al.*, "Personal identification based on iris texture analysis," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, pp. 1519–1533, 2003.
[17] A. Uhl and P. Wild, "Weighted adaptive hough and ellipsopolar transforms for real-time iris segmentation," in *Proc. ICB*, 2012, pp. 1–8.
[18] A. K. Jain, S. Prabhakar *et al.*, "FingerCode: a filterbank for fingerprint representation and matching," in *Proc. CVPR*, 1999.
[19] E. Maiorana, P. Campisi, and A. Neri, "Feature selection and binarization for on-line signature recognition," in *Proc. ICB*, 2009, pp. 1219–1229.
[20] J. Fierrez, J. Galbally *et al.*, "BiosecurID: a multimodal biometric database," *Pattern Analysis and App.*, vol. 13, pp. 235–246, 2009.