

Towards detection of morphed face images in electronic travel documents

U. Scherhag, C. Rathgeb and C. Busch
da/sec - Biometrics and Internet Security Research Group
Hochschule Darmstadt, Germany

{ulrich.scherhag, christian.rathgeb, christoph.busch}@h-da.de

Abstract—The vulnerability of face recognition systems to attacks based on morphed biometric samples has been established in the recent past. Such attacks pose a severe security threat to a biometric recognition system in particular within the widely deployed border control applications. However, so far a reliable detection of morphed images has remained an unsolved research challenge.

In this work, automated morph detection algorithms based on general purpose pattern recognition algorithms are benchmarked for two scenarios relevant in the context of fraud detection for electronic travel documents, i.e. single image (no-reference) and image pair (differential) morph detection. In the latter scenario a trusted live capture from an authentication attempt serves as additional source of information and, hence, the difference between features obtained from this face image and a potential morph can be estimated. A dataset of 2,206 ICAO compliant bona fide face images of the FRGCv2 face database is used to automatically generate 4,808 morphs. It is shown that in a differential scenario morph detectors which utilize a score level-based fusion of detection scores obtained from a single image and differences between image pairs generally outperform no-reference morph detectors with regard to the employed algorithms and used parameters. On average a relative improvement of more than 25% in terms of detection equal error rate is achieved.

Keywords-electronic travel documents; biometrics; face recognition; image morphing; fraud detection;

I. INTRODUCTION

Electronic travel documents, e.g. the ePassport, are equipped with biometric information in order to establish a strong link between the document and its holder. Facial images which are printed to the document or integrated via chips are predominantly used to verify this link, e.g. by border control agencies or Automated Border Control (ABC) gates. Automated face recognition represents a longstanding field of research and a variety of methods have been proposed over the past three decades [25], [15]. Generic face recognition systems comprise four major modules: face detection, face alignment, feature extraction, and comparison, where the latter two are generally conceded as key modules. Due to a high intra-class variability in human faces across the validity period of an electronic travel document, face recognition systems at ABC gates are operated at rather high False Match Rates (FMRs) to achieve acceptable False Non-Match Rates (FNMRs) [1], in contrast to other biometric technologies, e.g. iris recognition [4].

Recently, attacks on face recognition systems based on morphed biometric images have been presented [6], [22]. Morphing techniques can be used to create artificial




Figure 1: Examples for bona fide and morphed face images

biometric samples, which resemble the biometric information of two (or more) individuals in image and feature domain. If morphed biometric images are infiltrated to a biometric recognition system, e.g. during the issuance process of an ePassport, the subjects contributing to the morphed image will both (or all) be successfully verified against that single enrolled template. Hence, the unique link between individuals and their biometric reference data is not warranted. Fig. 1 shows an example of morphing two facial images.

Such attacks pose severe security threats to biometric systems, in particular to the issuance and verification process of electronic travel documents [6]. Black-listed criminal offenders can use an authentic passport complying with all physical security features to enter a country with the identity of an accomplice when performing three basic steps: (1) find a rather lookalike accomplice, (2) morph passport face photos of both, possibly utilizing free software available on the internet, and (3) the accomplice applies for an ePassport; the passport manufacturer will then issue an authentic passport equipped with the morphed biometric reference image and other identity attributes of the accomplice, which can be used to enter a country by both subjects. Diverse commercial face recognition systems have been found to be highly vulnerable to this type of attack [6]. Hence, an automated detection of morphed face images is vital to retain the security of operational face recognition systems where two detection scenarios depicted in Fig. 2 can be distinguished [21]:

- *Single image morph detection*: the detector processes a single image, e.g. an off-line authenticity check of an electronic travel document (this scenario is also referred to as no-reference morph detection);
- *Image pair morph detection*: a trusted live capture from an authentication attempt serves as additional

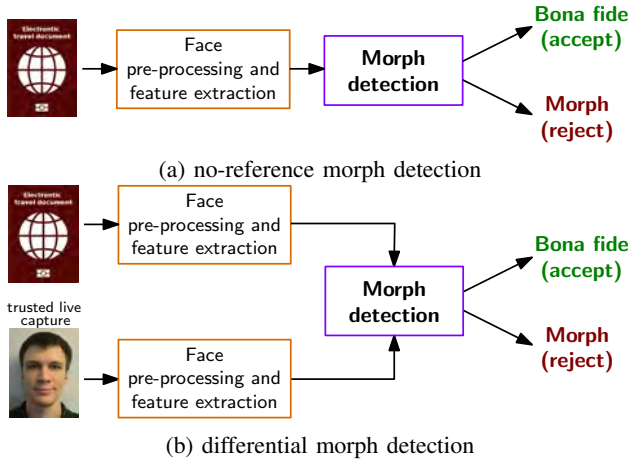


Figure 2: Morph detection scenarios

source of information for the morph detector, e.g. during authentication at an ABC gate (this scenario is also referred to as differential morph detection). Note that all information extracted by no-reference morph detectors might as well be leveraged within this scenario.

In this work, it is demonstrated that a Commercial Off-The-Shelf (COTS) face recognition system is highly vulnerable to the above mentioned attack using the generated morphed face images. The morph detection performance of general pattern recognition algorithms including texture descriptors, keypoint extractors, gradient estimators and a deep learning-based method is benchmarked for both mentioned scenarios. Within the differential scenario two approaches are considered to detect morphed face images, i.e. an analysis of feature vectors' differences obtained from a pair of images as well as a fusion of the resulting score with that obtained in the no-reference scenario. To the authors' knowledge this work represents the first attempt to directly compare the detection performance of morph detectors in a no-reference and a differential scenario on a comprehensive database of bona fide and morphed face images. Hence, this work provides the first quantitative measure of improvement in terms detection performance to be expected in a differential scenario (compared to a no-reference scenario) which is highly relevant to the detection of morphed face images in electronic travel documents.

This paper is organized as follows: related works are briefly summarized in Sect. II. Sect. III describes the different employed detection subsystems and how these are employed in both scenarios. Experimental results are reported and discussed in Sect. IV. Finally, conclusions are drawn in Sect. V.

II. RELATED WORK

Attacks based on morphed biometric samples were first introduced by Ferrara *et al.* [6]. Motivated by security gaps in the issuance process of electronic travel documents, 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 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) [1]. In a further study, the authors show that morphed face images are realistic enough to fool human examiners [7]. Scherhag *et al.* [22] reported moderate detection performance for benchmarking several general purpose texture descriptors used in conjunction with machine learning techniques to detect 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. Further, Hildebrandt *et al.* [9] suggest to employ generic image forgery detection techniques, in particular multi-compression anomaly detection, to reliably detect morphed facial images. Kraetzer *et al.* [14] evaluate the feasibility of detecting facial morphs with keypoint descriptors and edge operators. The benefits of deep neural networks for detecting morphed images has been recently investigated by Ramachandra *et al.* [18] and Seibold *et al.* [23]. In [19] the differences between face morphing and face averaging in the vulnerability of face recognition systems to both types of approaches are elaborated. The vulnerability of biometric systems relying on other biometric characteristics, e.g. fingerprint or iris, has been established, too [5], [20].

Gomez-Barrero *et al.* [8] proposed the first theoretical framework for measuring the vulnerability of biometric systems to attacks. Evaluations are conducted for diverse biometric systems where expected comparison scores of attacks based on morphed images or templates are directly derived from the mated and non-mated distributions of a face, fingerprint and iris recognition system. 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. Since there is no standardised manner to evaluate the vulnerability of biometric systems to attacks based on morphed images or templates, Scherhag *et al.* [21] introduced new metrics for vulnerability reporting (see Sect. IV), which strongly relate to the metrics defined in [11]. In addition, the authors provide recommendations on the assessment of morphing techniques. It is emphasized that unrealistic assumptions with respect to the quality of morphed biometric samples might cloud the picture regarding the performance of detection algorithms. Eventually, it is important to note that so far there is no publicly available database of morphed face images and no publicly available morph detection algorithms.

III. DETECTION SUBSYSTEMS

In the following subsection, we briefly summarize the employed pre-processing and feature extractors. Subsequently, their use in a no-reference and differential morph detection scenario is outlined.



(a) Subject 1 (b) Morph (c) Subject 2

Figure 3: Examples for bona fide and morphed face images



(a) Subject 1 (b) Morph (c) Subject 2

Figure 5: Example for SURF keypoint detection



(a) Subject 1 (b) Morph (c) Subject 2

Figure 4: Example for BSIF responses

A. Pre-processing and feature extraction

In the pre-processing stage the face of a subject is segmented and normalized according to eye coordinates detected by the *dlib* landmark detector [13]. Subsequently, the normalized region is cropped to 320×320 pixels to ensure that the morph detection algorithm is only applied to the facial region. Finally, the cropped face part is converted to a grayscale image. Fig. 3 depicts pre-processed face images of two subjects and their corresponding morph.

At feature extraction the pre-processed face image is optionally divided into 4×4 cells to retain local information. That is, feature extractors are applied separately on texture cells and the final feature vector is formed as a concatenation of feature vectors extracted from each cell. We employ the following four types of feature extraction methods (up to two algorithms are considered per type):

Texture descriptors: Local Binary Patterns (LBP) [16] and Binarized Statistical Image Features (BSIF) [12] are extracted from the cropped face images. For details on these texture descriptors the reader is referred to [16], [12]. While LBP simply processes neighbouring pixel values of each pixel, BSIF utilizes specific filters learned from a set of images. Obtained feature values are stored in a corresponding histograms. The use of these well-established general purpose texture descriptors has shown to be successful in diverse texture classification problems. Focusing on morph detection the process of image morphing is expected to cause changes in textual properties between bona fide and morphed face images. An example of BSIF applied to the images of Fig. 3 is depicted in Fig. 4. By testing different spatial sampling rates the two best configurations for LBP and BSIF were determined, thus, 3×3 and 9×9 LBP-patches and same sized BSIF filter sets extracting 8 bit per pixel for 3×3 and 12 bit per pixel for 9×9 are employed.

Keypoint extractors: Scale Invariant Feature Transform (SIFT) [17] and Speeded Up Robust Features (SURF) [3] extract sets of local keypoints. For details on keypoint



(a) Subject 1 (b) Morph (c) Subject 2

Figure 6: Example for sharpness features (two dimensions)

detection, the extraction of keypoint descriptors and keypoint matching the reader is referred to [17], [3]. Keypoint extractors are employed, since morphed images are expected to contain fewer keypoint locations, which are defined as maxima and minima of the result of difference of Gaussians function. That is, the amount of detected keypoints is used as descriptive feature. Fig. 5 shows an example of SURF keypoints detected in the images of Fig. 3.

Gradient estimators: Histogram of Gradients (HOG) and sharpness features are extracted from the normalized grayscale images. For further details to HOG the reader is referred to [24]. As a sharpness feature the mean of the gradient in two dimensions are calculated. The use of gradient-based methods is motivated by the fact that due to the morphing process high frequency changes are reduced and, hence, the steepness of gradients is decreased. An example of sharpness features extracted from the images of Fig. 3 is depicted in Fig. 6.

Deep learning-based method: we employ the *OpenFace* [2] algorithm in which rescaled images of 96×96 pixels are fed to the default pre-trained Deep Neural Network (DNN) to obtain a 128 dimensional face representation. This algorithm is applied to the pre-processed face image (no division into texture cells is applied). The use of Deep Facial Features (DFF) is motivated by recent advances in face recognition.

B. No-reference vs. differential morph detection

In the training stage feature vectors are extracted for each algorithm and support vector machines (SVMs) with Radial Basis Function (RBF) kernels are trained to distin-

Table I: Training and test set used for experimental evaluations

Gender	Training set					Test set				
	No. of subjects	No. of images	Trusted live capture	Bona fide images	Morphed images	No. of subjects	No. of images	Trusted live capture	Bona fide images	Morphed images
Male	49	2,236	98	485	1,653	58	2,210	499*	499	1,711
Female	35	1,403	101	561	741	39	1,165	462*	462	703
All	84	3,639	199	1,046	2,394	97	3,375	961*	961	2,414

*same images for live capture and bona fide

guish between bona fide and morphed face images using a disjoint training set. The SVMs of each single algorithm generate a normalized attack detection score in the range $[0, 1]$. For the no-reference morph detection approach, the *elements* of feature vectors extracted from a single image are analysed. For differential morph detection a trusted live capture from an authentication attempt of the same subject serves as additional input. This information is utilized by estimating the vector *differences* between feature vectors extracted from processed pairs of images which are used to train separate SVMs. In the differential scenario a *fusion* of attack detection scores produced by both of the above approaches can be employed. The fusion of both scores is expected to achieve competitive results since it combines morph detectors trained on absolute feature values as well as relative differences between bona fide and morphed face images.

IV. EXPERIMENTS

In the following subsections, we describe the experimental setup, conduct a vulnerability assessment of a COTS face recognition system to attacks based on the generated morphed face images and report and discuss the detection performance of the proposed system.

A. Experimental setup

Experiments are performed on a subset of the FRGCv2 face database. A total number of 2,206 frontal faces with neutral expression have been manually chosen and ICAO compliance has been verified, i.e. the distance between the eyes of a face has to be at least 90 pixels [10]. Based on this subset 4,808 morphed faces have been automatically generated for pairs of subjects of same gender using the *OpenCV* library. Further example images of bona fide and morphed face images are shown in Fig. 7 which illustrates the high quality of morphed face images being well in the quality limits set forth by ICAO and ISO/IEC standards. The division of images into training and test sets which has been chosen to obtain a balance between bona fide and morphed images during training is listed in Table I. During training and testing the trusted live capture subset is used as additional input for the differential scenario together with either a morph or a bona fide image. At training disjoint sets (live captures and bona fide images) and at testing all pair-wise combinations are used.

B. Face recognition vulnerability assessment

The vulnerability of a COTS face recognition system to attacks based on the generated morphed face images is assessed according to the metrics specified in [21],



(a) Subject 1 (b) Morph (c) Subject 2

Figure 7: Examples of bona fide and morphed face images of subjects of same gender, ethnicity and age group

in particular, in terms of Mated Morph Presentation Match Rate (MMPMR). This metric is an adaptation of the general Impostor Attack Presentation Match Rate (IAPMR) introduced in ISO/IEC 30107-3 [11] which is defined as the proportion of attack presentations using the same presentation attack instrument species in which the target reference is matched. However, in the adaptation the MMPMR covers the fact that not one target subject (contained in the morphed reference) is matched - but both subjects who earlier contributed to the morphed image are expected to be matched if the morphing attack is considered to be successful.

When employing the default decision threshold of the COTS face recognition system a MMPMR of 1 is obtained. This means all face images of subjects contributing to a morphed face image are successfully matched against it, hence, the attacks reveal a success chance of 100%.

C. Morph detection performance evaluation

The performance of the detection algorithms is reported according to metrics defined in ISO/IEC 30107-3 [11]. The Bona Fide Presentation Classification Error Rate (BPCER) is defined as the proportion of bona fide presentations incorrectly classified as presentation attacks in a specific scenario. The Attack Presentation Classification Error Rate (APCER) is defined as the proportion of attack pre-

Table II: Performance rates of no-reference and differential morph detection methods

Algorithm	D-EER			BPCER-10			BPCER-20		
	No-reference elements	Differential differences	fusion	No-reference elements	Differential differences	fusion	No-reference elements	Differential differences	fusion
LBP _{1,3}	5.1%	3.9%	2.6%	1.7%	0.6%	0.3%	5.3%	2.7%	1.3%
LBP _{4,3}	5.2%	3.9%	2.6%	1.7%	0.6%	0.3%	5.5%	2.9%	1.3%
LBP _{1,9}	13.7%	7.3%	8.5%	21.0%	5.2%	6.0%	35.0%	16.0%	16.2%
LBP _{4,9}	11.9%	7.4%	7.8%	15.0%	4.6%	5.2%	30.1%	15.8%	12.0%
BSIF _{1,3}	2.9%	4.4%	2.4%	1.3%	1.7%	1.0%	1.9%	3.8%	1.5%
BSIF _{4,3}	3.5%	4.7%	2.6%	1.1%	2.2%	0.8%	2.3%	4.5%	1.3%
BSIF _{1,9}	16.5%	9.3%	12.5%	22.9%	8.4%	15.1%	33.2%	23.6%	24.5%
BSIF _{4,9}	10.9%	9.8%	9.5%	11.7%	9.3%	8.9%	27.8%	33.7%	24.1%
SIFT ₁	22.9%	28.8%	18.0%	31.0%	54.0%	28.8%	38.1%	73.6%	38.3%
SIFT ₄	14.6%	31.8%	16.6%	18.5%	58.5%	22.8%	27.4%	72.7%	34.0%
SURF ₁	20.0%	18.1%	16.1%	26.1%	39.8%	25.4%	60.1%	66.1%	32.9%
SURF ₄	19.1%	25.2%	17.9%	27.4%	48.9%	27.7%	41.4%	69.0%	39.6%
Sharp ₁	26.8%	23.5%	22.6%	44.1%	59.8%	39.4%	54.9%	79.9%	50.0%
Sharp ₄	17.9%	6.6%	5.8%	24.2%	2.5%	1.5%	39.8%	9.4%	7.7%
HOG ₁	26.8%	26.6%	21.3%	50.9%	56.7%	31.4%	73.9%	84.9%	42.8%
HOG ₄	24.3%	24.1%	20.4%	64.8%	66.4%	49.3%	81.3%	83.2%	74.4%
DFE	26.1%	24.2%	21.0%	60.4%	67.7%	49.3%	77.2%	83.8%	74.2%

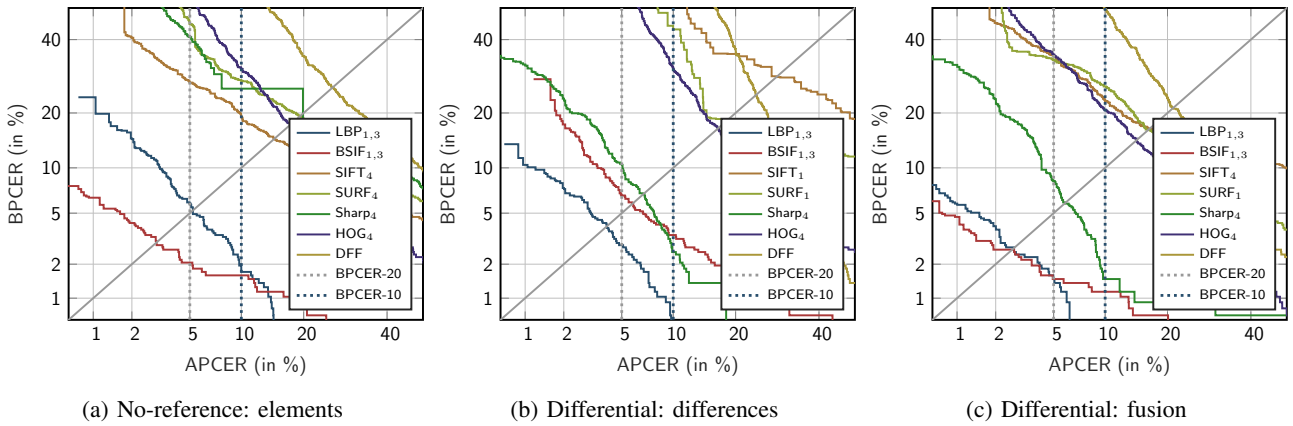


Figure 8: DET-plots of no-reference and differential morph detection methods

presentations using the same presentation attack instrument species incorrectly classified as bona fide presentations in a specific scenario. Further, the BPCER-10 and BPCER-20 represent the operation points yielding an APCER of 10% and 5%, respectively. Additionally, to be comparable to published works, the Detection Equal Error Rate (D-EER) will be reported.

Performance rates of all detection methods and different scenarios are listed in Table II. The corresponding detection error trade-off (DET) curves are depicted in Fig. 8. The division into cells is indicated via sub-indices. For the texture descriptors the patch size is indicated via a second sub-index, e.g. LBP_{4,3} corresponds to a division into 4×4 cells and a LBP-patch size of 3×3 pixels.

Fig. 8 (a) depicts the DET-plots of the best performing configurations for the no-reference approach. Competitive detection rates are achieved for texture descriptors where BSIF_{1,3} achieves the best performance of D-EER=2.9%. Moderate detection accuracy is achieved for keypoint extractors and gradient estimators. Applying the default net designed for recognition purposes, DFE reveals the highest D-EER. However, it is expected that application-specific training will significantly improve deep learning-based approaches with the potential drawback of data-

overfitting.

Within the differential approach which utilizes only vector differences the detection performance is improved to D-EER=3.9% for LBP while most BSIF settings are influenced negatively. The effect on the keypoint extractors, gradient estimators and deep features is mostly minor or negative. The corresponding DET-plot is shown in Fig. 8 (b). However, in the differential scenario where a score level fusion of both approaches is performed detection performance is generally improved compared to the no-reference scenario. Fig. 8 (c) shows the according DET-plot. The lowest D-EER for BSIF is as low as 2.4%. In terms of BPCER-10 and BPCER-20 even more significant performance gains are obtained which is emphasized in the corresponding DET curves. With respect to the performance requirements defined in [1], i.e. FNMR \leq 5% at a FMR of 0.1%, the differential fusion approach reveals practical detection performance.

V. CONCLUSION AND FUTURE WORK

The integration of biometric information to electronic travel documents is vital in order to establish a strong and permanent link between the document and its holder. However, this link is annulled if a morphed face image is

infiltrated to the biometric system, e.g. during the issuance process of an electronic travel document. This vulnerability in the life-cycle of electronic travel documents, e.g. the ePassport, calls for robust and reliable morph detection subsystems to be integrated to identity verification checks based on electronic travel documents.

Research on attacks based on morphed biometric samples and prevention of these is still in statu nascendi. However, at the time of this writing we see an increasing interest in this topic where ongoing activities of different research labs focus on the detection of morphed faces in a no-reference scenario, e.g. [22], [9], [14], [18]. In contrast to this, the presented work additionally considers the scenario in which a bona fide face image of the subject to be verified is available during the morph detection process. With regard to identity verification checks based on face verification and electronic travel documents, e.g. at ABC gates, this differential scenario is of particular interest. For numerous trained morph detectors based on different well-established pattern recognition methods it is shown that within this scenario substantial improvement in terms of detection performance can be achieved. That is, the quantitative difference between detection performance rates obtained in a no-reference and a differential scenario reported in this work will provide a useful basis for system engineers implementing face morph detection algorithms in operational systems.

The creation of a database of printed and scanned (morphed) face images and a corresponding evaluation of the presented morph detection methods in different scenarios is subject to future work. Further, alternative morph creation software could be analysed in 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] FRONTEX – Research and Development Unit: Best practice technical guidelines for automated border control (ABC) systems, 2012. Version 2.0.
- [2] B. Amos, B. Ludwiczuk, and M. Satyanarayanan. OpenFace: A general-purpose face recognition library with mobile applications. Technical report, CMU School of Computer Science, 2016.
- [3] H. Bay, A. Ess, T. Tuytelaars, and L. V. Gool. Speeded-Up Robust Features (SURF). *Computer Vision and Image Understanding*, 110(3):346 – 359, 2008.
- [4] J. Daugman. How iris recognition works. *Trans. on Circuits and Systems for Video Technology*, 14(1):21–30, 2004.
- [5] M. Ferrara, R. Cappelli, and D. Maltoni. On the feasibility of creating double-identity fingerprints. *IEEE Trans. on Information Forensics and Security*, 12(4):892–900, 2017.
- [6] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. In *Proc. Int. Joint Conf. on Biometrics (IJCB)*, pages 1–7, 2014.
- [7] M. Ferrara, A. Franco, and D. Maltoni. On the effects of image alterations on face recognition accuracy. In *Face Recognition Across the Imaging Spectrum*, pages 195–222. Springer International Publishing, 2016.
- [8] M. Gomez-Barrero, C. Rathgeb, U. Scherhag, and C. Busch. Is your biometric system robust to morphing attacks? In *Proc. Int. Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2017.
- [9] M. Hildebrandt, T. Neubert, A. Makrushin, and J. Dittmann. Benchmarking face morphing forgery detection: Application of stirtrace for impact simulation of different processing steps. In *Proc. Int. Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2017.
- [10] International Organization for Standardization. Information technology – Biometric data interchange formats – Part 5: Face image data. ISO/IEC 19794-5:2005 consolidated, JTC 1/SC 37, 2005.
- [11] ISO/IEC JTC1 SC37 Biometrics. *ISO/IEC IS 30107-3:2017, IT – Biometric presentation attack detection – Part 3: Testing and Reporting*.
- [12] J. Kannala and E. Rahtu. BSIF: Binarized statistical image features. In *Proc. Int'l Conf. on Pattern Recognition (ICPR'12)*, pages 1363–1366, 2012.
- [13] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10, 2009.
- [14] C. Kraetzer, A. Makrushin, T. Neubert, M. Hildebrandt, and J. Dittmann. Modeling attacks on photo-ID documents and applying media forensics for the detection of facial morphing. In *Proc. Workshop on Information Hiding and Multimedia Security (IH&MMSec)*, pages 21–32, 2017.
- [15] S. Z. Li and A. K. Jain. *Handbook of Face Recognition (2nd edition)*. Springer, 2011.
- [16] S. Liao, X. Zhu, Z. Lei, L. Zhang, and S. Z. Li. Learning multi-scale block local binary patterns for face recognition. In *Proc. Int'l Conf. on Biometrics (ICB'07)*, pages 828–837, 2007.
- [17] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, 2004.
- [18] R. Ramachandra, K. Raja, S. Venkatesh, and C. Busch. Transferable deep-cnn features for detecting digital and print-scanned morphed face images. In *2017 IEEE Conf. on Computer Vision and Pattern Recognition Workshop (CVPRW)*, July 2017.
- [19] R. Ramachandra, K. Raja, N. Vetrekar, S. Venkatesh, and C. Busch. Face morphing versus face averaging: Vulnerability and detection. In *Proc. Int. Joint Conf. on Biometrics (IJCB)*, pages 1–6, 2017.
- [20] C. Rathgeb and C. Busch. On the feasibility of creating morphed iris-codes. In *Proc. Int. Joint Conf. on Biometrics (IJCB)*, pages 1–6, 2017.
- [21] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. Veldhuis, L. Spreeuwens, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Raghavendra, and C. Busch. Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting. In *Int. Conf. of the Biometrics Special Interest Group (BIOSIG)*, pages 1–12, 2017.
- [22] U. Scherhag, R. Raghavendra, 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. Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2017.
- [23] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert. Detection of face morphing attacks by deep learning. In *Proc. Int. Workshop on Digital Watermarking (IWDW)*, pages 107–120, 2017.
- [24] C. Shu, X. Ding, and C. Fang. Histogram of the oriented gradient for face recognition. *Tsinghua Science & Technology*, 16(2):216–224, 2011.
- [25] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Comput. Surveys*, 35(4):399–458, 2003.