An Overview of Recent Advances in Assessing and Mitigating the Face Morphing Attack

Andrey Makrushin
Otto-von-Guericke University of Magdeburg
Magdeburg, Germany
andrey.makrushin@ovgu.de

Andreas Wolf
Bundesdruckerei GmbH
Berlin, Germany
andreas.wolf@bdr.de

Abstract— **The face morphing attack enables the illegitimate sharing of photo-ID documents intended for identity verification. Multiple users may use the same passport, driver license or health insurance card without being condemned. This paper summarizes recent advances in protecting the photo-ID-based verification from the morphing attack. We explain the attack along with the standard approach of creating morphed face images. We identify research gaps and open challenges by summarizing studies assessing the potential of the morphing attack as well as studies concerned with generating databases of morphed face images and examining the performance of morphing detectors. We discuss new performance metrics looking for conformity with the standard on presentation attack detection. Based on the current advances, we recommend technical and organizational security mechanisms to mitigate or even prevent the morphing attack.**

Keywords—face morphing attack, morphing detection

I. Introduction

The human face as a means of biometric authentication is widely accepted. The number of security systems making use of automated face recognition (AFR) is growing rapidly. Besides many advantages, face-based verification has one substantial disadvantage – it is easy to circumvent. The biometric research community has for a long while been concerned with the face presentation attack [1]. Recently, a novel and even more sophisticated fraud is brought to the forefront – face morphing attack [2]. This attack in its potential consequences by far outperforms face presentation attacks [3].

The face morphing attack has two stages: application for a new document with a morphed photograph by a regular user and a document usage by an illegitimate user. More precisely, a morphed photograph of Alice and Mallory is used by Alice as a part of an application for a new photo-ID document. If an officer accepts the photograph, the issued document will be authentic and perfectly regular and will, therefore, pass all optical and electronic authenticity and integrity checks. Mallory may illegitimately use this document claiming to be Alice without being condemned. A key to successful morphing attack is a morphed photograph. Formally, this is an image that would be generated by a face morphing process while gradually transforming one face image (source) to another (target), if the process stopped in between. Hence, a morphed photograph resembles both source and target faces. Usually, the similarity between the morphed face and the original faces is enough to deceive both humans and machines performing face verification (see Section III). Note that a morphed photograph can be created for more than two faces. For more details on morphing see [4].

The most critical security application menaced by the morphing attack is the Automated Border Control (ABC). Aiming at speeding up border crossing and saving expensive manpower, ABC gates are equipped with AFR systems to capture a live photograph and compare it with a digital passport photograph stored on a chip of an electronic Machine Readable Travel Document (eMRTD). Although the automation is of great benefit for high-throughput applications like an airport passenger control, it opens up the door for wanted criminals to illegally cross a border by practicing the morphing attack [2][5][6]. ABC is one example of an access control scenario that could be compromised by the morphing attack. In fact, every access control scenario based on photo-ID verification is vulnerable even with manual verification. Another prominent example of the morphing attack is an application of a driver license, a health insurance card, or an individual travel card by multiple users. Note that these documents have only printed photographs and foresee a manual identity check. At the moment, there exist no security mechanisms for detecting morphing in a document issuing process in many countries incl. passport application procedures in U.S.A., India and some European countries such as France or Germany [3].

In order to prevent the morphing attack, it is important to understand the morphing process and its eventual traces [3], and integrate morphing detection at both stages: the application for a new document and usage of a document for identity verification. There are three clues to detect morphed images: intrinsic content-independent traces of image manipulation, visual artifacts of blending due to imprecise superimposition of facial components, and suspicious similarity to several gallery images. Detection of intrinsic traces of morphing could be hindered by intentional anti-forensic image manipulations [7]. Moreover, if an application for a new document admits a submission of a printed photograph, intrinsic traces will vanish after re-digitalization. Indeed, the morphed digital image is printed first, then submitted to the document office, then scanned, then rescaled and compressed to meet certain requirements [8] and only then stored on the chip of a document. Re-digitalization may also conceal visual artifacts especially after dramatic downscaling and strong compression.

Hereafter, we summarize studies addressing generation of morphed images in Section II, morphing attack risk assessment in Section III, existing morph detectors in Section IV and performance evaluation metrics in Section V.

ISBN 978-90-827970-1-5 © EURASIP 2018

II. Generation of Morphed Face Images

The accessibility of morphing software and the ease of creating morphed face images makes the face morphing attack extremely dangerous. Currently, even an inexperienced user can create natural-looking high-quality facial morphs. Moreover, the popular smartphone app for image messaging "Snapchat" includes functionality for seamless face swapping [9] that can be seen as an extreme case of morphing with a marginal blending parameter. More sophisticated approaches for seamless face swapping are proposed in [10] and [11].

Early studies on face morphing have roots in late-90s [12] [13]. However, there were no malicious objectives but rather the interest in social psychological face perception and development of animation tools or special effects in the film industry. Nonetheless, the main principles of the morphing process have not been changed since then. A common morphing procedure include four steps: laying a mesh on the faces based on some fixed landmarks, warping corresponding mesh polygons to the same shapes and locations, blending of face regions, and seamless stitching of the blended face into one of the original backgrounds. The proportions of faces in a morph are controlled by the blending factor α. The value ranges from 0 to 1 telling the ratio of a target face. The ratio of a source face is 1-α. According to [14], the 50/50 morphs (α=0.5) are most useful for potential fraudsters, as it gives the highest likelihood of being accepted as a photo-ID for two different people. The success of the morphing attack depending on the blending factor is studied in [15] and [16]. The quality of morphed photographs is assessed regarding three aspects [17]: (i) visual quality indicating the absence of visible artifacts and visual similarity to the faces involved, (ii) biometric quality reflecting the successful verification by an AFR system against the faces involved, and (iii) forensic quality indicating the absence of intrinsic traces of image manipulation.

First dedicated attempts to create facial morphs for performing the morphing attack are done manually using the software tool GIMP and its plug-in GAP [2][5]. Based on the AR face database [18], 80 morphs were generated and further used in the FVC-onGoing Face Morphing Challenge [19]. In [20], GIMP/GAP is used to generate 450 morphs. Note that face databases are selected to comply with the ICAO portrait quality standard for eMRTD [8]. In [21], the authors implement two techniques for blending faces: "face morphing" and "face averaging" and apply these to generate 1423 morphs of each type from the FRGC face database [22]. In [23], the result of warping and blending of whole images is called a "complete morph". In contrast, the authors propose a "splicing morph" resulting from warping and blending of only the face regions with the subsequent stitching into one of the original backgrounds. Splicing keeps images clear of ghosting artifacts. There are 1326 complete and 2614 splicing morphs generated from the Utrecht ECVP face database [24] and 500 splicing morphs from the FEI face database [25]. A similar approach is reported in [15] improving face stitching by means of Poisson blending. The authors generate 7260 morphs from the CMU Multi-PIE face database [26]. Seamless face stitching with an optimal face cutting path and local Poisson blending is proposed in [27]. Combined morphs, proposed in [17], unite advantages of complete and splicing morphs. In [14], morphs are created from Glasgow Face Matching Test images [28] using Psyhomorph software [29]. In [30], a dataset of morphs created from the FERET dataset [31] is made freely available. Two datasets are presented in [16]. The first contains automatically generated morphs with various blending factors, and the second manual morphs generated using Sqirlz Morph. All datasets are summarized in TABLE I.

TABLE I. Databases of morphed face images

Study	*# Morphs*	*Face database*	*Morphing tool*
Ferrara et al. [5]	80	AR	GIMP/GAP
Raghavendra et al. [20]	450	Proprietary	GIMP/GAP
Robertson et al. [14]	?	GFMT	Psyhomorph
Raghavendra et al. [21]	1423	FRGC v.2	morphing
	1423	FRGC v.2	averaging
Makrushin et al. [23]	1326	Utrecht ECVP	complete
	2614	Utrecht ECVP	splicing
	500	FEI	splicing
Neubert et al. [17]	2652	Utrecht ECVP	combined
Wandzik et al. [15]	7260	MultiPIE	splicing + Poisson blending
Seibold et al. [27]	9000	Proprietary	splicing + optimal cutting edge
Biometix [30]	?	FERET	?
Ferrara et al. [16]	6000	AR + FRGC + FERET	splicing
	100	FRGC + FERET	Sqirlz Morph

The other popular trend is generating morphed faces by inverting deep convolutional neural network (DCNN) features using generative adversorial networks (GAN). The impressive results are reported in [32] using images from the CelebA dataset [33]. The performance of a professional artist to morph faces using Photoshop can be found at www.instagram.com/ gesichtermix.

III. Security Risk of the Morphing Attack

In order to prove that the morphing attack is a serious threat to security systems based on face matching, researchers study performances of both humans (see TABLE II.) and AFR systems (see TABLE III.) to biometrically compare morphed and genuine faces. The matching performance is measured in terms of morph acceptance rate (MAR) and false rejection rate (FRR). MAR is a relative number of falsely accepted morphing trials while FRR - a relative number of falsely rejected genuine trials. For humans, "accepted" means that, for a pair of images, an observer decides that images depict the same person, and "rejected" that the images depict different persons. For AFR systems, "accepted" means that the similarity value returned by the system exceeds or equals a certain threshold, otherwise "rejected". The decision threshold is selected to comply to the FRONTEX prescription for ABC [34] so that false acceptance rate (FAR) does not exceed 0.1%.

A. Humans

Simulating ABC scenario, in [5], the face matching is performed by 44 border guards and 543 laymen. Border guards accepted more morphing trials, but rejected less genuine trials, both demonstrating a very poor matching performance. The study in [14] examines whether human performance increases

TABLE II. HUMAN PERFORMANCE TO BIOMETRICALLY COMPARE MORPHED AND GENUINE FACES

Study	*Test persons*	*#Morph / #Gen*	*MAR*	*FRR*
Ferrara et al. [5]	44 border guards	8 / 5	~74.92%	~8.33%
	543 laymen	20 / 10	~57.55%	~12.25%
Robertson et al. [14]	28 laymen	7 / 7	68%	9%
	42 laymen after coaching	7 / 7	21%	~13%

after coaching. In the first experiment, humans performed a standard test and in the second experiment an additional option was included, namely to directly decide that the document image is morphed resulting in a "no match" outcome. Since many morphed images have visual faults, the performance of humans to make correct decisions drastically improved. The difference in acceptance rates in both studies can be explained by the different quality of morphed images. Nonetheless, even the MAR of 21% is unacceptably high for security applications such as an ABC scenario.

B. *AFR systems*

Considering the trend of replacing human experts by AFR systems, researchers tested commercial off-the-shelf (COTS), freeware and open source AFR systems. The COTS systems being tested are VeryLook SDK, Luxand FaceSDK, EyeFace SDK and FaceVACS. Two face-based unlocking apps of Samsung Galaxy S3 and iPhone were examined as well as two open source systems OpenFace and VGG-Face both based on DCNNs. The ability of an AFR system to discern morphed and genuine faces strongly depends on the quality of morphs used in experiments. The MAR_{1000} values measured at 0.1% FAR spread from 33% to 100% revealing a general vulnerability of all kinds of AFR systems to facial morphs.

IV. MORPHING DETECTION

Due to the limited capability of AFR systems and humans to mismatch morphed images, morphing detection is indispensable at both stages: the application for a new document and usage of a document for identity verification. The detection technique is different depending on whether only a document photograph is presented (blind detection) or a document photograph and a live photograph are presented. In the former case, a detection algorithm looks for content-independent anomalies or content-dependent visual artifacts making use of techniques from digital image forensics. In the latter case, face demorphing [16] or biometric thresholding could be applied.

A. *Blind detection*

In [23], human performance to recognize morphed images as such is examined in the experiment with 42 participants. The images were printed on photo paper with passport dimensions of 35x45 mm. The FAR/FRR values of 44.6%/43.64% indicate the inability of humans to reliably perform this task.

The approaches to blind detection of morphed face images are listed in TABLE IV. The very first study [20] examines the texture analysis techniques commonly applied for presentation attack detection. BSIF features classified with a linear SVM are shown to perform best with convincing error rates on the limited set of morphed images. In [6] and [41], the BSIF-based approach is tested with slightly modified datasets resulting in significantly higher error rates. In [41], feature-level fusion of two DCNNs AlexNet and VGG19 trained by transfer learning is shown to outperform BSIF features. In [23], Benford features are utilized to reveal JPEG compression inconsistency for detection of splicing morphs. In [3], the number of keypoints in the face region is counted to disclose the blurring effect after blending operation in morphing. In [42], the images are progressively compressed and the number of keypoints is counted at each stage of compression to measure the loss of details that should be lower in morphed images. In [43], topological data analysis is applied to count the number of connected components based on 2-ones uniform LBP codes. In [27], classification performance of three DCNN architectures in both modes learning from scratch and transfer learning is evaluated for facial morphs. In [17], the feature-level fusion of Benford and keypoint features is done to improve morphing detection.

TABLE III. PERFORMANCE OF AFR SYSTEMS TO BIOMETRICALLY COMPARE MORPHED AND GENUINE FACES

Type	*AFR system*	*Ver.*	*Study*	*MAR_{1000}*	*FRR_{1000}*
Commercial Off-the-Shelf (COTS)	VeryLook SDK [35]	5.4	[2]	100%	?
		5.5	[5]	99.4%	1.1%
		?	[18]	100%	?
		?	[6]	100% 99.19% 95.9%	0%
		9.0	[27]	~95%	?
	Luxand FaceSDK [36]	4.0	[2]	100%	?
		4.0	[5]	70%	38.4%
		6.1	[23]	83.27% 53.77%	0%
		6.1	[17]	72.83%	0%
	EyeFace SDK [37]	3.11.0	[5]	33.1%	77.4%
	FaceVACS [38]	9.1.4	[21]	90.33% 83.62%	?
		?	[9]	90%	?
Freeware	Samsung Galaxy S3, Android's face unlock app	?	[14]	27% (threshold unknown)	?
	iPhone, face unlocking	?	[9]	100%	?
Open source	OpenFace [39]	?	[6]	95.4% 95.9% 95.7%	10.81%
	VGG-Face [40]	?	[17]	51.88% 47.22% 46.91%	0%

B. *Detection in the presence of a live image*

Morphing detection in the presence of a live image is addressed in [16]. In the so-called "demorphing" approach, the live image is subtracted from the document image and the resulting difference image is matched against the live image. A low similarity score indicates a morphed document image.

C. *Detection after anti-forensic image manipulations*

The StirTrace benchmark to simulate common anti-forensic operations is adopted in [7] for face images. It is demonstrated that the detection performance of the morphing detector from [23] drastically drops in case of applying post-processing,

ISBN 978-90-827970-1-5 © EURASIP 2018

TABLE IV. PERFORMANCE OF BLIND MORPH DETECTION ALGORITHMS

Study	*Approach*	*Morphs*	*FAR*	*FRR*
[20]	BSIF features + linear SVM	GIMP/GAP	**3.46%**	**0%**
[6]	BSIF features + linear SVM	from [20]	7.1%	7.1%
[41]	BSIF features + linear SVM	extended set of morphs from [20]	22.7%	22.7%
	Feature-level fusion of AlexNet and VGG19 + P-CRC classifier		8.23%	8.23%
[23]	Benford features + linear SVM	splicing	12.11%	3.82%
		complete	12.97%	
		FEI splicing	**0%**	**3.50%**
[3]	Keypoints + J48 classifier	complete	10.0%	18.7%
		splicing	2.0%	
[42]	Progressive JPEG compression + keypoints + LMT classifier	splicing	3.7%	22.5%
		complete	10.0%	
[43]	Topological data analysis	splicing	0%	2%
		combined	1%	0%
		complete	~40%	?
[27]	AlexNet, from scratch	geometrical alignment + splicing + optimal cutting edge	1.9%	16.2%
	AlexNet, transfer learning		0.9%	11.4%
	GoogLeNet, from scratch		1.8%	10.0%
	GoogLeNet, transfer learning		1.2%	5.6%
	VGG19, from scratch		2.2%	10.9%
	VGG19, transfer learning		**0.8%**	**3.5%**
[17]	Feature-level fusion of Benford and keypoint features + naïve Bayes	complete	13.96%	?
		splicing	0.24%	
		combined	12.41%	

especially for downscaling and noise filters. In [17], the benchmark is extended to evaluate biometric and forensic qualities of morphed face images in terms of performances of selected AFR systems and morphing detectors respectively. To simulate real digital passport photographs, the benchmark contains the so-called "passport scaling 15kB" including cropping to ICAO compatible shape, scaling to 413×531 pixels and JPEG compression to the image size not exceeding 15 kB.

D. Detection after re-digitalization

It is demonstrated in [6] that the print-scan procedure makes the classification with BSIF features from [20] flawed increasing EER form ~7% to more than 20% (see TABLE V.). To support classification of print-scanned images, in [41], two DCNNs AlexNet and VGG19 are trained by transfer learning using small sets of digital as well as print-scanned morphed and genuine images. For both digital and print-scanned images, the morph detection performance is better than that with BSIF features. In [21], print-scanned "morphed faces" and "averaged faces" are tested with LBP features applied in YCbCr and HCV color spaces. The detection performance is better than that with BSIF features. Anyway, the detection error rates are too high to rely on blind detection with re-digitalized passport images.

TABLE V. PERFORMANCE OF MORPH DETECTION ALGORITHMS WITH RE-DIGITALIZED PHOTOGRAPHS

Study	*Approach*	*Images*	*EER*
Scherhag et al. [6]	BSIF-SVM from [20]	digital	**7.1%**
		print-scan (HP)	**20.7%**
		print-scan (Ricon)	**24.8%**
Raghavendra et al. [41]	BSIF-SVM from [20]	digital	22.70%
		print-scan (HP)	26.12%
		print-scan (Ricon)	23.29%
	AlexNet+VGG19 trained by transfer learning	digital	**8.23%**
		print-scan (HP)	**17.64%**
		print-scan (Ricon)	**12.47%**
Raghavendra et al. [21]	BSIF-SVM from [20]	print-scan, morphed	14.65%
		print-scan, averaged	20.51%
	LBP in YCbCr and HCV	print-scan, morphed	**9.48%**
		print-scan, averaged	**2.93%**

V. EVALUATION METHODOLOGY

For vulnerability evaluation of AFR systems under the morphing attack, the biometric performance metrics FAR and FRR are extended by MAR [5] measuring the relative number of successful morphing trials. In the ABC scenario, the morphed persons have different roles: accomplice and criminal [2]. In [16], the "Criminal MAR" (C-MAR) is proposed to measure the relative number of successful morphing trials done by criminals. It is asserted that matching of morphed images against an accomplice is of no interest because an accomplice is the legitimate document owner. C-MAR corresponds to the IAPMR - the metric proposed in the standard on presentation attack detection [44]. For the general case, then the contributing persons interchangeably use the "magic" document, the morphing attack would be considered successful only if all contributing persons were verified. Hence, counting successful morphing trials should be replaced by counting successful morphs. In [23], a morph is considered successful if it matches all original images of the persons involved. It is pointed out that MAR values are pessimistically high proposing "realistic MAR" (rMAR). The same idea is conveyed in [45] proposing the Mated Morph Presentation Match Rate (MMPMR) for one probe image per person and the MinMax- and ProdAvg-MMPMR for several probe images per person.

For morphing detection, the standard metrics are fine, since the task does not differ from the standard two-class problem with morphs as positive examples. FAR and FRR are used as synonyms for FNR (miss rate) and FPR (false alarm rate) respectively. In [44], these metrics are referred to as APCER and BPCER.

VI. CONCLUSION

Due to the ease of creating high-quality facial morphs and the limited capability of automated face recognition (AFR) systems and humans to detect these, the face morphing attack is proven to have potential for compromising photo-ID-based verification and its most critical application – the automated border control (ABC). In order to mitigate the attack, detection of morphed face images is required at both stages: the application for a new document and usage of a document for identity verification. Since morphing detection algorithms still have very high error rates, and the detection performance drastically degrades with re-digitalized and anti-forensically manipulated images, intensive research on the limits of detection techniques is required. Researchers should evaluate the proposed morph detectors in fair tests with large and diverse databases of morphed images which should be made publicly available. Arrangement of evaluation campaigns such as FRVT MORPH [46] with the unified performance metrics would be highly appreciated. If the detection performance is judged insufficient, administrative actions should be taken. The first option could be a requirement to submit high-resolution digital images to enable application of digital image forensics to disclose traces of morphing. Note that as long as persons are allowed to submit printed images to the document issuing office, the morphing attack will remain a serious threat. However, prohibiting an off-site photograph production and refusing printed photographs should be considered as a measure of last resort.

Acknowledgment

The work in this paper has been funded in part by the German Federal Ministry of Education and Science (BMBF) through the research programme under the contract no. FKZ: 16KIS0509K.

References

[1] R. Raghavendra and C. Busch, “Presentation Attack Detection Methods for Face Recognition Systems: A Comprehensive Survey,” ACM Comput. Surv., Vol. 50, No. 1, Article 8, 2017.

[2] M. Ferrara, A. Franco, and D. Maltoni, “The Magic Passport,” Proc. IEEE Int. Joint Conf. on Biometrics, pp. 1–7, 2014.

[3] C. Kraetzer et al., “Modeling Attacks on Photo-ID Documents and Applying Media Forensics for the Detection of Facial Morphing,” Proc. IH&MMSec, pp. 21–32, 2017.

[4] G. Wolberg, “Image morphing: a survey,” Visual Computer, Vol. 14, No. 8, pp. 360–372, 1998.

[5] M. Ferrara, A. Franco, and D. Maltoni, “On the Effects of Image Alterations on Face Recognition Accuracy,” In Bourlai, T. (ed.) Face Recognition Across the Electromagnetic Spectrum, pp. 195–222, 2016.

[6] 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,” Proc. IWBF, 2017.

[7] M. Hildebrandt et al.,“Benchmarking Face Morphing Forgery Detection: Application of StirTrace for Impact - Simulation of Different Processing Steps,” Proc. Int. Workshop on Biometrics and Forensics (IWBF), 2017.

[8] A. Wolf, ICAO: Portrait Quality (Reference Facial Images for MRTD), Version 0.7. Standard. International Civil Aviation Organization, 2016.

[9] A. Agarwal, R. Singh, M. Vatsa, and A. Noore, “SWAPPED! Digital Face Presentation Attack Detection via Weighted Local Magnitude Pattern,” Proc. IEEE Int. Joint Conference on Biometrics, 2017.

[10] I. Kemelmacher-Shlizerman, “Transfiguring portraits,” ACM Trans. on Graphics Vol. 35, No.4, Article 94, 2016.

[11] I. Korshunova, W. Shi, J. Dambre, and L. Theis, “Fast face-swap using convolutional neural networks,” Proc. IEEE Int. Conf. on Computer Vision, pp. 3697–3705, 2017.

[12] T. A. Busey, “Physical and Psychological Representations of Faces - Evidence From Morphing,” Psyhological Science, Vol. 9, No. 6, 1998.

[13] M. Steyvers, “Morphing techniques for manipulating face images,” Behavior Research Methods, Vol. 31, No.2, pp. 359–369, 1999.

[14] D. J. Robertson et al., “Fraudulent ID using face morphs: Experiments on human and automatic recognition,” PloS One, 2017.

[15] L. Wandzik et al., “CNNs under Attack: On the Vulnerability of Deep Neural Networks Based Face Recognition to Image Morphing,” Proc. 16th Int. Workshop on Digital Forensics and Watermarking, 2017.

[16] M. Ferrara, A. Franco, and D. Maltoni, “Face Demorphing,” Trans. on Information Forensics and Security, Vol.13, No.4, pp. 1008–1017, 2018.

[17] T. Neubert et al., “Extended StirTrace Benchmarking of Biometric and Forensic Qualities of Morphed Face Images,” IET Biometrics, 2018.

[18] A. M. Martinez and R. Benavente, “The AR Face Database,” Computer Vision Center, Technical Report #24 , 1998.

[19] FVC onGoing: Face Morphing Challenge, https://biolab.csr.unibo.it/FVCOnGoing/UI/Form/BenchmarkAreas/BenchmarkAreaFMC.aspx, checked 24.2.2018

[20] R. Raghavendra, K. Raja, and C. Busch, “Detecting Morphed Facial Images,” Proc. 8th IEEE Int. Conf. on Biometrics: Theory, Applications and Systems (BTAS), 2016.

[21] R. Raghavendra et al., “Face Morphing Versus Face Averaging: Vulnerability and Detection,” Proc. IJCB, 2017.

[22] P. J. Phillips et al., “Overview of the face recognition grand challenge,” Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 1, pp. 947–954, 2005.

[23] A. Makrushin, T. Neubert and J. Dittmann. “Automatic generation and detection of visually faultless facial morphs,” Proc. 12th Int. Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 6: VISAPP, pp. 39–50, 2017.

[24] Utrecht ECVP face database, http://pics.stir.ac.uk/2D_face_sets.htm, checked 24.2.2018

[25] FEI Face Database, http://fei.edu.br/~cet/facedatabase.html, checked 24.2.2018

[26] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, “Multi-PIE,” Proc. Int. Conf. on Automatic Face and Gesture Recognition, 2008.

[27] C. Seibold et al., “Detection of Face Morphing Attacks by Deep Learning,” Proc. IWDW, 2017.

[28] A. M. Burton, D. White, and A. McNeill, “The Glasgow Face Matching Test,” Behavior Research Methods, Vol. 42, No. 1, pp. 286–291, 2010.

[29] B. Tiddeman, M. Burt, and D. Perrett, “Prototyping and transforming facial textures for perception research,” IEEE Comput. Graph. Appl., Vol. 21, No.5, pp. 42–50, 2001.

[30] Biometix, http://www.biometix.com/2017/09/18/new-face-morphing-dataset-for-vulnerability-research, checked 24.02.2018

[31] FERET face dataset, https://www.nist.gov/itl/iad/image-group/color-feret-database, checked 24.02.2018

[32] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive Growing of GANs for Improved Quality, Stability, and Variation,” Proc.ICLR, 2018.

[33] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep Learning Face Attributes in the Wild,” Proc. International Conference on Computer Vision, 2015.

[34] FRONTEX, “Best Practice Technical Guidelines for Automated Border Control (ABC) Systems,” 2015.

[35] Neurotechnology, VeriLook Face Verification SDK, http://www.neurotechnology.com, checked 24.02.2018

[36] Luxand, Luxand FaceSDK, http://www.luxand.com, checked 24.02.2018

[37] Eyedea Recognition, EyeFace SDK, http://www.eyedea.cz/eyeface-sdk, checked 24.02.2018

[38] Cognitec, FaceVACS, http://www.cognitec.com/technology.html, checked 24.02.2018

[39] B. Amos, B. Ludwiczuk, and M. Satyanarayanan, “Openface: A general-purpose face recognition library with mobile applications,” CMU-CS-16-118, CMU School of Computer Science, Tech. Rep., 2016.

[40] O. M. Parkhi, A. Vedaldi, and A. Zisserman, “Deep face recognition,” Proc. British Machine Vision Conference, 2015.

[41] R. Raghavendra et al., “Transferable Deep-CNN features for detecting digital and print-scanned morphed face images,” Proc. 30th Int. Conf. on Computer Vision and Pattern Recognition Workshop, 2017.

[42] T. Neubert, “Image Degradation: A First Detection Approach for Face Morphing Forgeries,” Proc. IWDW, 2017.

[43] A. Asaad and S Jassim, “Topological Data Analysis for Image Tampering Detection,” Proc. IWDW, 2017.

[44] ISO/IEC 30107-3:2017, Information Technology – Biometric presentation attack detection – Part 3: Testing and reporting, 2017.

[45] U. Scherhag et al., “Biometric Systems under Morphing Attacks: Assessment of Morphing Techniques and Vulnerability Reporting,” Proc. BIOSIG, pp. 149–159, 2017.

[46] NIST, FRVT MORPH, https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt-morph, checked 15.06.2018