A. Brömme, C. Busch, A. Dantcheva, C. Rathgeb and A. Uhl (Eds.): BIOSIG 2019,
Lecture Notes in Informatics (LNI), Gesellschaft für Informatik, Bonn 2019 1

Decoupling texture blending and shape warping in face morphing

Matteo Ferrara[1], Annalisa Franco[2], Davide Maltoni[3].

Abstract: Automatic face recognition systems (FRS) are quite sensitive to the well-known face morphing attack, as pointed out by several researchers. Considering that, in the perspective of a fraudulent document usage, a criminal would certainly do its best to fool both humans and FRSs, the design of effective countermeasures should consider the trickiest and challenging conditions. Cognitive studies show that facial texture and shape (the two main components modified by face morphing) play a different role in face recognition by humans. This paper aims at understanding the behavior of FRSs with respect to these two factors and to identify the morphing parameters that maximize the probability of a successful attack.

Keywords: Face Morphing Attack, eMRTD, Automated Border Control, Face recognition.

1 Introduction

Face morphing [FFM14] [Sc19] recently emerged as one of the most serious and challenging security threat in the application of automated face recognition systems to Machine Readable Travel Documents (eMRTD); the research community is devoting significant efforts to design effective morphing detection techniques, and NIST recently set up a specific benchmark [Ni19a]. A typical face morphing aims at mixing the identity of two different subjects in a single image by applying two (usually equally weighted) transformations, i.e., shape warping and texture blending [FFM18a] [St99] [Wo98].

Cognitive studies on face recognition [An16] [It14] reveal that, according to human perception, shape and texture play different roles in the recognition process and influence the final decision to a diverse extent. Do the same considerations hold for automated face recognition? What is the behavior of Face Recognition Software (FRS) in relation to blending and warping? Answering these questions can have a relevant impact on the probability of success of the face morphing attack, which involves fooling at the same time the human officer (at document issuance) and the automatic recognition software (at verification stage). Consequently, identifying the most relevant factors in face morphing will allow to increase the robustness of detection algorithms.

[1] Department of Computer Science and Engineering, University of Bologna, via dell'Università, 50 - Cesena - Italy, matteo.ferrara@unibo.it
[2] Department of Computer Science and Engineering, University of Bologna, via dell'Università, 50 - Cesena - Italy, annalisa.franco@unibo.it
[3] Department of Computer Science and Engineering, University of Bologna, via dell'Università, 50 - Cesena - Italy, davide.maltoni@unibo.it

Raghavendra et al. [Ra17] made a preliminary step towards this kind of evaluation, analyzing the vulnerability of FRS to averaging face attacks. However, to the best of our knowledge, a systematic evaluation of the influence of shape and texture on morphing attack has not been carried out before. In addition, the datasets used for the evaluation of morphing detection techniques are all generated applying the same amount of shape warping and image blending, which is not necessarily the optimal choice to improve the probability of attack success.

This paper formalizes a morphing process where different weighting factors are applied for blending and warping (Section 2); an extensive evaluation is then carried out to evaluate the robustness of commercial face verification software on morphed images generated varying the two parameters. The database generated for testing is described in Section 3 while Section 4 describes the experiments carried out and summarizes the results obtained. Finally, Section 5 presents concluding remarks and identifies possible future research directions.

2 Face Morphing Process

Given two images I_0 and I_1 and the corresponding relevant face landmarks (eye corners, nose tip, etc.) $P_0 = \{\mathbf{u}_i, i = 1,..,N\}$ and $P_1 = \{\mathbf{v}_i,\ i = 1,..,N\}$ respectively, the morphed image is generally obtained [FFM18a] as the composition of:

- a warping function $w_{B\to A}$, representing the geometric transformation needed to align the set of points *B* to the set of points *A*;
- an image blending, simply obtained as a weighted average of the pixel intensity of the two images.

Existing morphing techniques usually adopt a single weighting factor α (morphing factor) for both transformations, so that the intensity of each pixel $\mathbf{p}$ of the morphed image I_α can be computed as:

$$I_\alpha(\mathbf{p}) = (1-\alpha)\cdot I_0\left(w_{P_\alpha\to P_0}(\mathbf{p})\right) + \alpha\cdot I_1\left(w_{P_\alpha\to P_1}(\mathbf{p})\right) \quad (1)$$

In eq. (1) P_α represents the intermediate landmark positions obtained from P_0 and P_1 according to α as follows:

$$P_\alpha = \{\mathbf{r}_i | \mathbf{r}_i = (1-\alpha)\cdot\mathbf{u}_i + \alpha\cdot\mathbf{v}_i, \mathbf{u}_i \in P_0, \mathbf{v}_i \in P_1\} \quad (2)$$

To evaluate the importance of geometric warping and image blending separately, Eq. (1) can be modified as follows:

$$I_{\alpha_B,\alpha_W}(\mathbf{p}) = (1-\alpha_B)\cdot I_0\left(w_{P_{\alpha_W}\to P_0}(\mathbf{p})\right) + \alpha_B\cdot I_1\left(w_{P_{\alpha_W}\to P_1}(\mathbf{p})\right) \quad (3)$$

where α_B and α_W are the blending and warping factors, respectively.

As to function $w_{B\to A}$, several warping techniques have been proposed in literature [Wo94]. In this work, the two set of points (A and B) are represented by means of triangular meshes, and the warping function is obtained as the local spatial transformations that map each couple of corresponding triangles [RA89].

The different effects of blending and warping are shown in Fig. 3 where two very different subjects have been selected (see Fig. 1) to highlight the influence of α_B and α_W, separately.

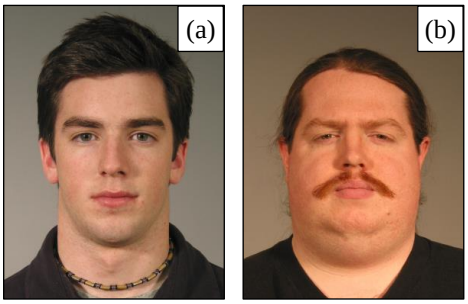

Fig. 1: Images I_0 and I_1 used to generate the morphed images in Fig. 3.

The morphed image is usually pretty clear and smooth in the central face region, while the area surrounding the face exhibits visible artifacts (the contributing images have different hair style, background, etc.). To make the morphed images more realistic, an automatic retouching procedure is adopted; the background region surrounding the face is replaced with the corresponding region of one of the contributing images (the one with the highest blending factor), after a proper alignment. Further details can be found in [FFM18a]. An additional step aimed at equalizing the skin color is here adopted before background substitution. In fact, due to different illumination conditions or skin color between the two face images, the retouching result can be not satisfactory (see Fig. 2.c). To overcome this issue, the histogram matching method described in [GW17] is applied to equalize the face region of the image with smaller blending factor to the face region of the other one (see Fig. 2.d).

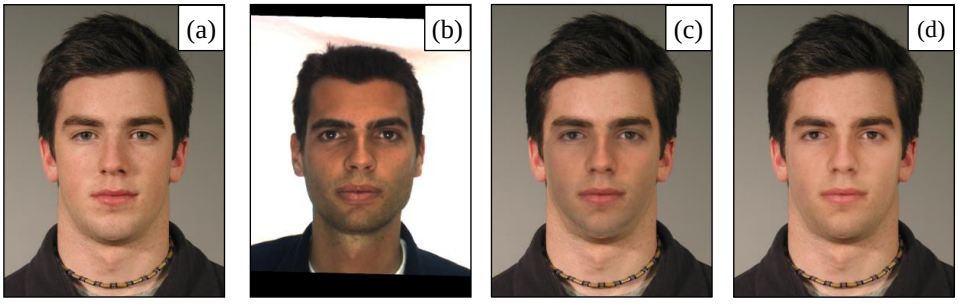

Fig. 2: Example of the automatic face morphing process. (a) (b) the original images I_0 and I_1, and the morphed images obtained with Eq. (3): automatic retouching (c) without and (d) with color equalization. The value of α_B and α_W to obtain (c) and (d) are here 0.5 and 0.3, respectively.

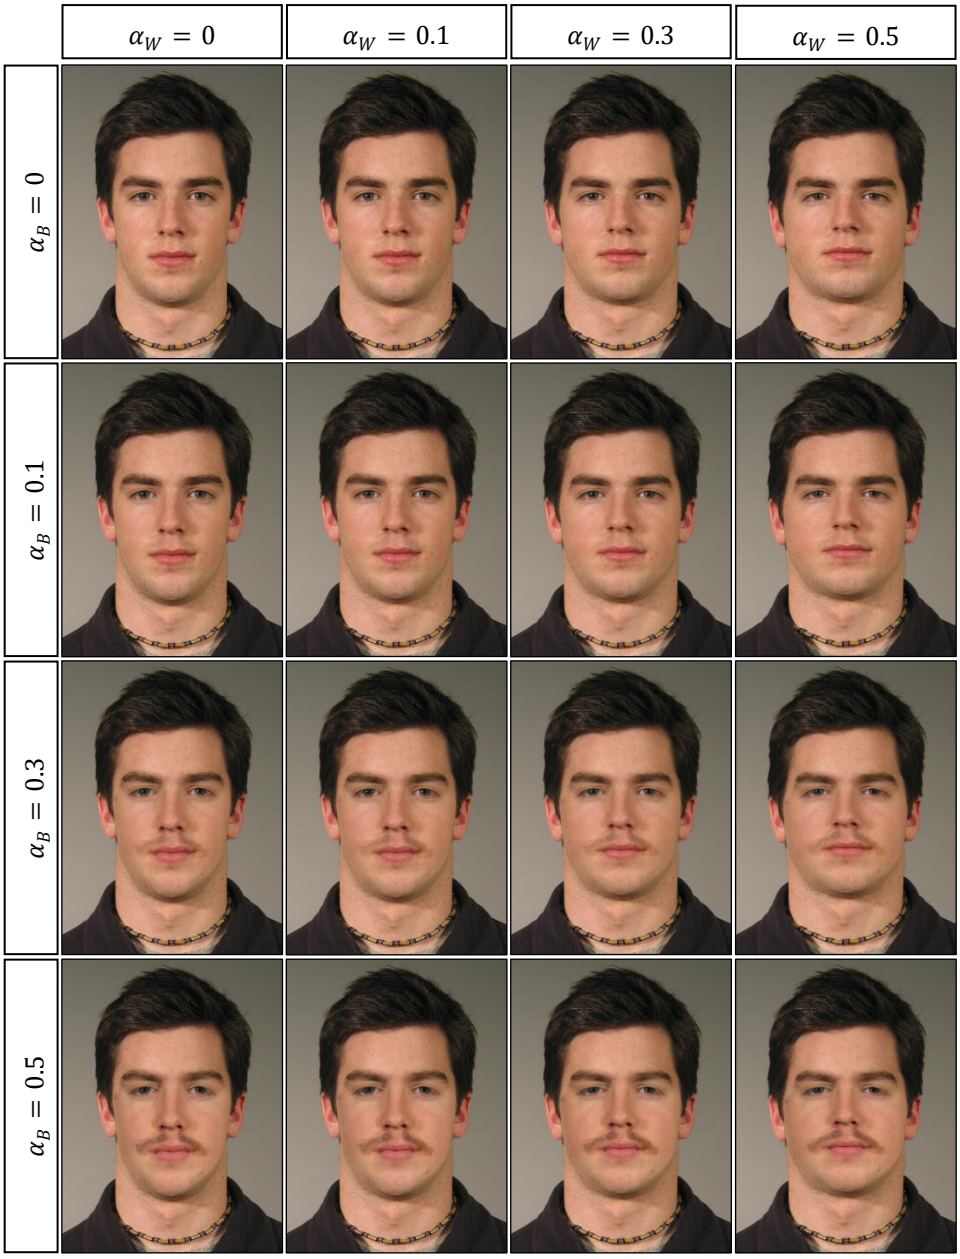

Fig. 3: Morphed images obtained with different blending and warping factors by combining Fig. 1.a (I_0) and Fig. 1.b (I_1).

3 Database

The database used in the experimentation was automatically generated starting from existing face images taken from AR [MB98], FRGC [Pi05] and Color Feret [Pi98] [Pi00] databases. The selected images have been manually checked to ensure they fulfil ISO/ICAO specifications [Is11]. Moreover, the subjects wearing glasses have been excluded since the resulting morphing could be affected by visible artifacts. The final number of subjects is 280: 80 for AR (34 males and 46 females) and 100 for both FRGC and Color Feret (50 males and 50 females each). Two images of each subject have been selected: the former is used for morphing generation and the latter for testing.

Thanks to the fully automatic generation, we can produce a large number of samples and at the same time precisely control the blending and warping factors α_B and α_W. The database consists of a collection of sub-datasets, each containing a specific set of morphed images M_{α_B,α_W} obtained applying the morphing process described in Section 2 with blending factor α_B and warping factor α_W:

$$\alpha_B, \alpha_W \in \{0,0.1,0.2,0.3,0.4,0.5\}.$$

To simulate a scenario where a criminal tries to find an accomplice with no criminal records to apply for an eMRTD by presenting a morphed photo, analogously to [FFM18a], the selection of candidate images to produce morphing cases was performed as follows:

1) the first image of each subject (i.e., the criminal) is compared with the first image of other k subjects of the same gender (i.e., possible accomplices) randomly chosen from the same source database (i.e., AR, FRGC or Color Feret). The subject presenting the maximum similarity with the criminal is selected as the optimal accomplice for morphing. The comparisons have been performed using the commercial face recognition software Neurotechnology VeriLook SDK 10.0 [Ne19]. This choice is aimed at maximizing the probability of fooling the face verification software at the gate.
2) The first image of the criminal and of the optimal accomplice are morphed following the procedure described in Section 2. The result is a set of morphed images, one for each combination of α_B and α_W.

The above procedure is repeated $t = 2$ times for each subject obtaining 560 morphed images for each combination of α_B and α_W. Each time the optimal accomplices used in the previous iterations are excluded from the random selection. The database has been generated with $k = 10$ to simulate a realistic scenario where a criminal can find 10 very good friends with no criminal records who accept to play the role of possible accomplices.

Note that in [FFM18a] at the end of the generation process, a quality evaluation step is applied to discard low quality morphed images: the generated images are compared against the original ones of both subjects and in case of non-match they are rejected. Since the aim of this work is to study the importance of blending and warping factors in generating good morphed images, this step has not been applied.

4 Experimental evaluation

Similar to [FFM18b] the experiments have been carried out using two commercial face recognition SDKs (referred to as SDK_1 and SDK_2) which provided top performance in the recent Face Recognition Vendor Test (FRVT) Ongoing [Ni19b][GNH18]; the names of the SDKs cannot be disclosed and the results will be therefore presented in anonymous form. All the SDKs fulfill the operational conditions suggested by Frontex for ABC gates (a maximum False Rejection Rate of 5% at a False Acceptance Rate of 0.1%) [Fr12]. During the experimentation, for each SDK, the security threshold indicated in the corresponding documentation to achieve FAR=0.1% has been used. Since in this paper we focus on morphing attacks, the performance is evaluated in terms of Mated Morph Presentation Match Rate (MMPMR) [Sc17] with the aim to quantify the percentage of morphing attacks able to fool the SDKs. To this purpose the MMPMR for all SDKs have been measured by comparing morphed face images against the test face image of the criminal subject.

Tab. 1 reports the MMPMR of SDK_1 and SDK_2 for different combinations of α_B and α_W. For both SDKs warping and blending have a different impact on the probability of success of the attack. While geometric modifications obtained increasing the warping factor α_W do not heavily affect recognition accuracy (see ranges $\alpha_B \in [0; 0.1], \alpha_W \in [0.4; 0.5]$), an opposite behavior is observed for the blending factor α_B ($\alpha_B \in [0.4; 0.5], \alpha_W \in [0; 0.1]$). Hence, for a criminal it would be much more convenient to create a morphed image with $\alpha_B = 0.5$ and $\alpha_W \in [0; 0.1]$ instead of using a balanced morphing factor in the range [0.2; 0.3] as stated in [FFM18a][RKB17]. This choice would increases the chances of successful attack at the border (from about 6-46% to 45-81%) keeping unaltered the chances of fooling the human officer during the document issuing process. In fact, a visual inspection of several generated morphs reveals that the difference between the two images is imperceptible, in particular when look-alike subjects are involved (see the example of Fig. 4). Moreover, we should always consider that human recognition capabilities are surprisingly error-prone in front of unfamiliar faces [YB17] and small appearance variations would probably be neglected. Finally, it is important to note that the MMPMR values could be even higher because, in a real scenario, a criminal would try to produce high quality morphed images, discarding the morphs with a low probability of success and applying manual retouching to remove unrealistic artifacts.

5 Conclusions

In the context of face recognition, humans are more sensitive to texture than to geometry; this study reveals that the same holds for FRSs. Assigning different weighting factors to texture blending and geometry warping during the face morphing process significantly increases the chances of success, especially in the presence of look-alike subjects. With respect to the optimal morphing factor identified by previous works in the range [0.2; 0.3], a more tricky setting is represented by images generated with blending factor in the range

$[0.4; 0.5]$ and warping factor in $[0; 0.1]$.

Even if this study has been carried out using a specific morphing algorithm, we think that similar results can be obtained with other morphing techniques since they are very similar; minor differences could be related to landmark detection or the specific warping technique adopted. Moreover, the commercial SDKs used in the evaluation are among the top-performing software at NIST FRVT and they are currently installed in real airport ABC gates. We are therefore convinced that the behavior highlighted in this work is worth of attention, and that uneven weighting of blending and warping should become one of the testing cases to consider for performance assessment of morphing detection mechanisms.

α_B \ α_W	**0**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**
0	2.0%	2.3%	2.7%	2.7%	3.6%	3.8%
0.1	4.8%	5.4%	6.8%	7.9%	8.9%	10.9%
0.2	10.5%	13.6%	17.7%	21.8%	25.9%	28.9%
0.3	30.0%	34.8%	41.3%	46.4%	52.9%	57.5%
0.4	54.1%	62.9%	68.6%	74.5%	77.9%	81.3%
0.5	74.8%	81.1%	86.3%	89.8%	92.9%	95.0%

α_B \ α_W	**0**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**
0	0.0%	0.0%	0.0%	0.0%	0.5%	0.9%
0.1	1.1%	1.1%	1.3%	2.3%	3.2%	4.1%
0.2	3.2%	5.2%	6.4%	8.0%	9.5%	10.4%
0.3	10.5%	13.0%	14.6%	17.7%	21.8%	23.9%
0.4	22.3%	29.5%	34.6%	39.5%	43.4%	48.8%
0.5	45.2%	52.0%	58.6%	64.6%	70.4%	73.8%

Tab. 1 MMPMR of SDK_1 (left) and SDK_2 (right) for each combination of α_B and α_W. Different values are represented by different blue levels (the darker, the greater). Parameter values denoted by the red frame are suggested in [FFM18a] to create effective face morphing. Here we argue that moving to the green region increases the change to successfully perpetrate the attack.

Acknowledgment

The work leading to these results has received funding from the European Union's Internal Security Fund — Borders and Visa under grant agreement n° 842250 - SOTAMD.

The content of this work represents the views of the authors only and is their sole responsibility.

The European Commission does not accept any responsibility for use that may be made of the information it contains

Fig. 4: Example images from the database used for the experimentation. The morphed images are obtained combining the two images I_0 and I_1 with different blending (α_B) and warping (α_W) factors.

References

[An16] T. J. Andrews et al., "Contributions of feature shapes and surface cues to the recognition and neural representation of facial identity," Cortex, vol. 83, pp. 280-291, 2016.

[FFM14] M. Ferrara, A. Franco, and D. Maltoni, "The Magic Passport," in IEEE International Joint Conference on Biometrics (IJCB), Clearwater, Florida, USA, 2014, pp. 1-7.

[FFM18a] M. Ferrara, A. Franco, and D. Maltoni, "Face Demorphing, "IEEE Transactions on Information Forensics and Security, vol. 13, no. 4, pp. 1008-1017, 2018.

[FFM18b] M. Ferrara, A. Franco, and D. Maltoni, "Face Demorphing in the Presence of Facial Appearance Variations," in European Signal Processing Conference (EUSIPCO), 2018.

[Fr12] FRONTEX - R&D Unit, "Best Practice Technical Guidelines for Automated Border Control (ABC) Systems - v2.0," FRONTEX, 2012.

[GNH18] P. Grother, M. Ngan, and K. Hanaoka, "Ongoing Face Recognition - Part 1: Verification," NIST, Gaithersburg, MD, 2018.

[GW17] R. C. Gonzalez and R. E. Woods, Digital Image Processing, 4th ed.: Pearson, 2017.

[Is11] ISO/IEC 19794-5, Information technology - Biometric data interchange formats - Part 5: Face image data, 2011.

[It14] M. L. Itz et al., "Neural correlates of facilitations in face learning by selective caricaturing

of facial shape or reflectance," NeuroImage, vol. 102, pp. 736-747, 2014.

[MB98] A. M. Martinez and R. Benavente, "The AR face database," Computer Vision Center, CVC Technical Report 1998.

[Ne19] Neurotechnology Inc. (2019, June) Neurotechnology Web Site. [Online]. http://www.neurotechnology.com/

[Ni19a] NIST. (2019, June) Face Recognition Vendor Test (FRVT) MORPH. [Online]. https://www.nist.gov/programs-projects/frvt-morph

[Ni19b] NIST. (2019, June) Face Recognition Vendor Test (FRVT). [Online]. http://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt

[Pi00] P. J. Phillips et al., "The FERET evaluation methodology for face-recognition algorithms," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 22, no. 10, pp. 1090-1104, 2000.

[Pi05] P. J. Phillips et al., "Overview of the face recognition grand challenge," in proceedings IEEE Computer Vision and Pattern Recognition, vol. 1, 2005, pp. 947-954.

[Pi98] P. J. Phillips et al., "The FERET database and evaluation procedure for face-recognition algorithms," Image and Vision Computing, vol. 16, no. 5, pp. 295-306, 1998.

[Ra17] R. Raghavendra et al., "Face Morphing Versus Face Averaging: Vulnerability and Detection," in IEEE International Joint Conference on Biometrics (IJCB), 2017.

[RA89] D. F. Rogers and J. A. Adams, Mathematical Elements for Computer Graphics, 2nd ed.: McGraw-Hill Higher Education, 1989.

[RKB17] D. J. Robertson, R. S. S. Kramer, and A. M. Burton, "Fraudulent ID using face morphs: Experiments on human and automatic recognition," PLoS ONE, vol. 12, no. 3, 2017.

[Sc17] U. Scherhag et al., "Biometric Systems under Morphing Attacks: Assessment of Morphing Techniques and Vulnerability Reporting," in International Conference of the Biometrics Special Interest Group (BIOSIG), 2017.

[Sc19] U. Scherhag et al., "Face Recognition Systems Under Morphing Attacks: A Survey," IEEE Access, pp. 23012-23026, 2019.

[St99] M. Steyvers, "Morphing techniques for manipulating face images," Behavior Research Methods, Instruments, & Computers, vol. 31, no. 2, pp. 359-369, 1999.

[Wo94] G. Wolberg, Digital Image Warping, 1st ed. Los Alamitos, CA, USA: IEEE Computer Society Press, 1994.

[Wo98] G. Wolberg, "Image morphing: a survey," The Visual Computer, vol. 14 (8), pp. 360-372, 1998.

[YB17] A. W. Young and A. M. Burton, "Recognizing Faces," Current Directions in Psychological Science, vol. 26, no. 3, pp. 212-217, 2017.