

Fused Classification For Differential Face Morphing Detection

Iurii Medvedev
Institute of Systems
and Robotics,
University of Coimbra,
Coimbra, Portugal
iurii.medvedev@isr.uc.pt

Joana Alves Pimenta
¹Institute of Systems
and Robotics,
University of Coimbra,
Coimbra, Portugal
joana.pimenta@isr.uc.pt

Nuno Gonçalves ^{1,2}
²Portuguese Mint and Official
Printing Office (INCM),
Lisbon, Portugal
nunogon@deec.uc.pt

Abstract

Face morphing, a sophisticated presentation attack technique, poses significant security risks to face recognition systems. Traditional methods struggle to detect morphing attacks, which involve blending multiple face images to create a synthetic image that can match different individuals. In this paper, we focus on the differential detection of face morphing and propose an extended approach based on fused classification method for no-reference scenario. We introduce a public face morphing detection benchmark for the differential scenario and utilize a specific data mining technique to enhance the performance of our approach. Experimental results demonstrate the effectiveness of our method in detecting morphing attacks.

1. Introduction

The development of deep learning techniques in recent years has led to significant progress in the field of face recognition, but sophisticated presentation attack techniques, such as face morphing, continue to pose security risks that require new protection solutions. Face morphing involves merging/blending of two or more digital face images to create a synthetic image that can share the biometric properties of original images and match different individuals. It can be difficult to detect such generated images with use of traditional human or computer-based face recognition methods.

The risks associated with face morphing are not hypothetical; they have been demonstrated through real-world incidents. One notable example occurred in 2018 when a German activist exploited face morphing techniques to issue an authentic German passport using a morphed face image of Federica Mogherini (at that time High Representative of the Union for Foreign Affairs and Security Policy) blended with their own photo [28]. This incident

highlighted the potential for face-morphing attacks to deceive identity verification systems and emphasized the need for effective detection methods. Additionally, face morphs have been occasionally detected during border control procedures, raising concerns about the circulation of morphed documents. Some recent investigations [44] acknowledged the presence of morphing cases, indicating the tangible risks and uncertainties surrounding the prevalence of such documents. These real-world examples underscore the urgency of robust morphing detection approaches to mitigate the security risks associated with face morphing attacks. That is why face morphing and its detection methods have gained interest in both industry [43] and academia [25].

Morphing detection methods in facial biometric systems can be categorized into two pipelines based on the processing scenario. The *no-reference* morphing attack detection algorithm is designed to detect morphing in a single image, with a focus on mitigating the risks of accepting manipulated images during the *enrollment* process, where successful acceptance of forged images can lead to the issuance of an authentic document that could deceive the face recognition system.

On the other hand, the *differential* morphing attack detection algorithms involve acquiring live data from an authentication system to provide reference information for detecting morphing attacks. This usually occurs during *automatic border control* and such approaches aim to identify discrepancies between the presented face and the stored biometric data, enabling the system to detect potential morphing attempts in real time and prevent unauthorized access with malicious ID documents (documents with accepted face morphs).

In this paper, we focus on differential face morphing detection and propose a novel deep learning method that incorporates sophisticated face recognition tasks and employs a fused classification scheme for morphs classification. We follow the no-reference MorDeepy [21] approach and adopt its methodology and data for the differential case.

Additionally, we extend benchmark utilities, which are proposed in [21] with a public face morphing detection benchmark for differential scenario.

2. Related Work

2.1. Face Recognition

Contemporary face recognition methods heavily rely on the utilization of deep learning techniques, particularly convolutional neural networks (CNNs), which have proven to be highly effective in extracting discriminative features from unconstrained facial images. [33]. These networks possess the ability to learn complex patterns and structures, making them well-suited for tackling the challenges associated with facial pattern recognition tasks.

Various deep learning strategies are employed for face recognition, all aimed at extracting low-dimensional facial representations - deep face features with high discriminatory capabilities.

For example, metric learning techniques focus on explicit optimization of the face representation by contrasting pairs of matched and non-matched samples with similarity metric [38]. Achieving reliable convergence with these methods necessitates extensive datasets and advanced sample mining techniques.

Classification-based methods have received major attention and they are better represented in recent academic research. These methods focus on learning face representation implicitly through a closed-set identity classification task [41]. Deep networks in these approaches encapsulate face representation in the last hidden layer and typically employ various softmax-based loss functions [41].

To achieve better discriminate properties of deep facial features various techniques are used. For instance, explicit compacting of intra-class features to their center [46] or several types of marginal restrictions, which address inter-class discrepancy [11,40]. Many recent works were focused on investigating sample-specific learning strategies, which are driven by various characteristics, such as sample quality [22], hardness of classification [15]. Some works use on properties of embedding as a proxy for the image quality (like norm of the features) [19,23], or rely on artificial assignment by known data augmentation [39]. These approaches try to control the feature distribution in the discriminative feature domain.

2.2. Face Morphing

Modern face recognition systems can very accurately match images of individuals however, they are still vulnerable to various malicious presentation attacks. Face morphing allows to design such attacks and drastically increases the probability for the face recognition network to return matched embeddings for unmatched biometric samples, es-

pecially in cases when the thresholds of face recognition systems are not set to support critically low false match rates.

Basic landmark-based face morphs were first investigated by Ferrara *et al.* [12]. Face morphing was performed directly in the image spatial domain by the face landmark alignment, image warping, and blending. Various morphing algorithms mentioned in the literature follow this strategy [3].

The field of face morphing has witnessed significant advancements with recent breakthroughs in Deep Learning techniques, leading to the development of several innovative tools and methodologies. Face morphing has witnessed significant advancements with recent breakthroughs in Deep Learning techniques. Generative Adversarial Networks (GANs) have emerged as a prominent and widely utilized approach in various generative tasks, including face morphing. MorGAN [7] approach pioneered this tool for face morphing generation. The StyleGAN [17] approach introduced a latent domain representation to control various aspects of the generated image, which enables the generation of high-quality face morphs without blending artifacts. The MIPGAN [47] method optimized StyleGAN specifically for face morphing, preserving the identity of the generated morphed face image. The diffusion autoencoders for face morphing were proposed by MorDIFF [6] to generate smooth and high-fidelity face morphing attacks.

2.3. Face Morphing Detection

Initially, the problem of face morphing detection focused on the no-reference scenario, where validation decisions were based on single image presentations. However, considering practical concerns, it became valuable to explore a differential approach that simulates the process of document verification by border control officers.

No-reference face morphing detection algorithms initially relied on analyzing local image characteristics like Binarized Statistical Image Features (BSIF) [29] or sensor noise (Photo Response Non-Uniformity) [9], texture features [32], local features in frequency and spatial image domain [24] or fusion of various features [35]. Deep learning methods for the no-reference case typically involve binary classification of pretrained face recognition features [30], which can be combined with local texture characteristics [45]. Additional pixel-wise supervision [8] or attention mechanism [1] can be applied. MorDeepy method [21] generalized single image morphing detection to unseen attacks by additional feature regularisation with face recognition task.

In contrast to no-reference case, differential face morphing detection is closely correlated with face recognition, as the discriminability of deep face representation usually helps combating attacks in this scenario.

For instance, many differential approaches rely on classification of pretrained deep features for face recognition [36].

Borghi et al. [4] conducted differential morphing detection through the fine-tuning of pretrained networks within a sophisticated framework involving identity verification and artifacts detection modules.

Qin et al. [27] proposed a method for detecting and locating face morphing attacks for both (single-image and differential) scenarios. The authors utilise feature-wise supervision, which provides better characterization of morphing patterns and localization of morphed areas.

Ferrara et al. [13] presented an alternative approach to the differential scenario, which involves reverting morphing by retouching the testing face image using a trusted live capture. This technique aims to unveil the true identity of the legitimate document owner.

In this work we propose to extend the no-reference MorDeepy [21] approach for the differential morphing detection scenario by adopting the methodology and data mining techniques to the differential pipeline.

3. Methodology

From the original work [21] we inherit the S-MAD methodology, which requires several modifications for the differential case. Recall that the face morphing detection here is made by the behavior of deep face features, which is achieved by regularizing the morphing detection with the face recognition task. The definition of the task is motivated by the ubiquity of classifying the face morphs (since they belong to 2 or more identities). This leads to the setup of two separate CNN-based deep networks that treat bona fide samples similarly but handle morphed samples differently. These networks do not share weights and are not trained in a contrastive manner, where positive and negative pairs are matched. Both networks learn high-level features through classification tasks, with each network assigning different identity labels to face morphs. The *First Network* labels them based on the original identity from the first source image, while the *Second Network* labels them according to the second original label.

In comparison to the S-MAD scenario, where the same image is sent to both networks, the D-MAD case implies processing a pair of images. That is why the fused classification strategy is adopted to the D-MAD in the following way (see Figure 1).

We keep the assumption of associating each image with two identity labels y_1 and y_2 , which are defined basing on the image origin. For the Bona Fide samples those labels are the same (copied from the original face image label), when for the Morphs those labels are different and are taken from the source face images. The sampling process for the *First Network* is not changed. For instance the image \hat{I} with

\hat{y}_1 and \hat{y}_2 is sampled for the input. For the input of the *Second Network* the image \check{I} (complementary to \hat{I}) with \check{y}_1 and \check{y}_2 is selected with a condition that $\check{y}_1 = \hat{y}_2$. The loss function components require the following rules. The identity classification for the *First Network* is made by the \hat{y}_1 , and by the \check{y}_2 for the *Second Network*. The ground truth cross label for the morphing binary classification is made by matching \hat{y}_2 and \check{y}_1 .

It is important to note that such formulation allows both images \hat{I} and \check{I} to be Morphs. However to match the D-MAD scenario, where the *Live Enrollment* image is genuine and trusted, we supervise selecting the \check{I} as a Bona Fide sample.

Due to the above modifications the formulation of the identity classification softmax-based loss components is transformed as follows:

$$L_1 = -\frac{1}{N} \sum_i \log \left(\frac{e^{\hat{W}_{\hat{y}_1}^T \hat{f}_i + \hat{b}_{\hat{y}_1}}}{\sum_j^C e^{\hat{f}_i \hat{y}_j}} \right) \quad (1)$$

$$L_2 = -\frac{1}{N} \sum_i \log \left(\frac{e^{\check{W}_{\check{y}_2}^T \check{f}_i + \check{b}_{\check{y}_2}}}{\sum_j^C e^{\check{f}_i \check{y}_j}} \right), \quad (2)$$

where $\{\hat{f}_i, \check{f}_i\}$ denote the deep features of the i -th sample pair, $\{\hat{W}, \check{W}\}$ and $\{\hat{b}, \check{b}\}$ are weights and biases of last fully connected layer (respectively for the $\{First, Second\}$ networks). N is the number of samples in a batch and C is the total number of classes.

For the morph binary classification component only the definition of the cross label is changed:

$$L_3 = -\frac{1}{N} \sum_i t \log \frac{1}{1 + e^{-D}} + (1-t) \log \left(1 - \frac{1}{1 + e^{-D}} \right), \quad (3)$$

where $D = \hat{f} \cdot \check{f}$ is a dot product of high level features extracted by *First* and *Second* backbones and the cross-label $t = \text{abs}(\text{sgn}(\hat{y}_2 - \check{y}_1))$ of the i -th sample pair.

The above strategy implies pushing morph samples towards their original classes differently by *First* and *Second* networks. This allows to increase the distance between the morph samples in the feature domain.

In this work we also consider a modification of the above approach. We propose to allocate separate classes for the morph samples. The formulation of such case means redefining the classification labels for the morph samples and doubling the number of identity classes C : $\hat{y}_1^* = \hat{y}_1 + C$; $\check{y}_2^* = \check{y}_2 + C$.

Such class allocation is made differently by *First* and *Second* in and does not impact the differentiating morphs by these networks. However, since it also pushes the morph samples away from their original classes it can help to increase discriminative power of deep face features. Further

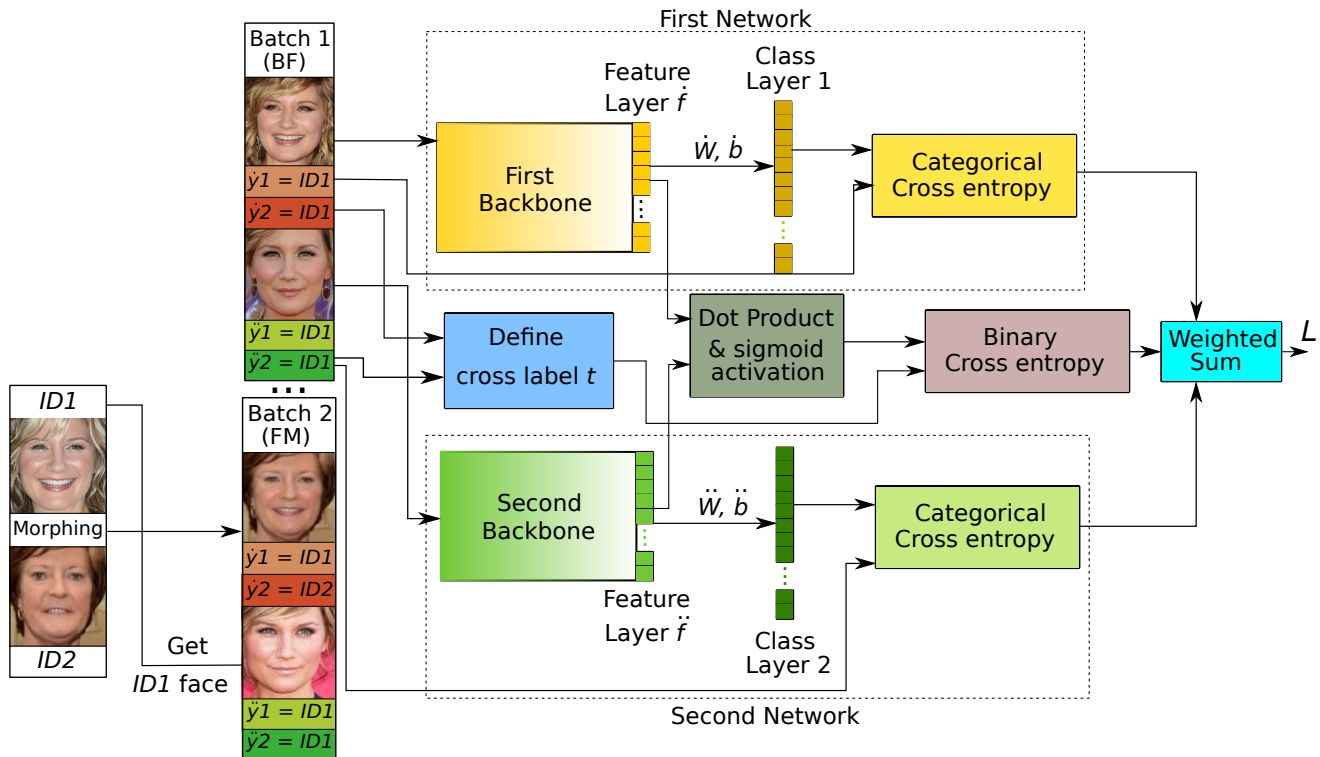


Figure 1. Schematic of the proposed D-MAD method. For simplicity of visualization batch contains a single image pair.

in the work we make experiments with both cases and mark the original strategy with the label $V1$ and the modified strategy with $V2$

4. Data Mining

With the adopted methodology the training can be proceeded on the same data as S-MAD. The only sampling of this data is different. Recall that VGGFace2 dataset [5], is used as a source of original bona fide images. We repeat the quality based filtering of this dataset and generate respective morphs.

4.1. Morph dataset

In this work we utilize two automatic methods for generating morphs. First we use a customized landmark-based morphing approach with blending coefficient 0.5. Second we generate GAN-based morphs with use of the StyleGAN [17] method. To synthesize such morphs for two original images, they are first projected to a latent domain and then their deep representations are interpolated. The resulting morph is generated from such interpolated latent embedding.

To ensure effective learning in the fused classification task (see Figure 1), it is crucial to have unambiguous class labeling in our training dataset of morphs, which are gen-

erated from face images from different classes. To address this, we follow the original S-MAD approach [21] and employ a strategy where the dataset is split into two disjoint parts attributed to the *First* and *Second* networks. When generating face morphs, we randomly pair images from these identity subsets and label the morphed images accordingly for classification by the respective networks. This approach acts as a regularization technique and enhances the performance of morphing detection. By separating the dataset into two disjoint identity sets, we ensure consistent classification of morphed combinations by the networks.

4.2. Selfmorphing

Fully automatic landmark morphing methods often introduce visible artifacts to the generated images, which can bias the learning process towards these artifacts. However, real fraudulent morphs are retouched to remove such perceptual artifacts. To address this, we utilize *selfmorphs*, generated by applying face morphing to images of the same identity. We follow the original S-MAD approach [21] and use *selfmorphs* as bona fide samples to focus on the behavior of deep face features rather than detecting artifacts. We assume that the deep discriminative face features remain intact after *selfmorphing*.

In our work *selfmorphs* are generated for both landmark-

Table 1. Comparison of Fused Classification (FC) with the Binary Classification (BC) by $APCER@BPCER = (0.1, 0.01)$ in several protocols.

Method	$APCER@BPCER = \delta$							
	protocol-asml		protocol-facemorpher		protocol-webmorph		protocol-stylegan	
	$\delta = 0.1$	$\delta = 0.01$	$\delta = 0.1$	$\delta = 0.01$	$\delta = 0.1$	$\delta = 0.01$	$\delta = 0.1$	$\delta = 0.01$
BC	0.315	0.729	0.245	0.649	0.391	0.701	0.913	0.997
FCV1	0.063	0.351	0.066	0.514	0.135	0.529	0.556	0.959
FCV2	0.039	0.275	0.061	0.315	0.102	0.4595	0.501	0.957




Figure 2. DET curves for the D-MAD algorithms by NIST FRVT MORPH benchmark. a) protocol Visa-Border; b) protocol Manual; c) protocol MIPGAN-II; d) protocol Print + Scanned.

based and GAN-based morphing approaches.

4.3. Result dataset

Our resulting dataset consist of $\sim 500k$ original images from VGGFace2, $\sim 250k$ their landmark-based selfmorphs, $\sim 250k$ their GAN-based selfmorphs, $\sim 500k$ landmark-based morphs and $\sim 500k$ GAN-based morphs. The overall dataset is balanced by the amount of bona fides and morphs.

5. Benchmarking.

One commonly used metric for evaluating single image morphing detection is the relationship between the Bona fide Presentation Classification Error Rate (BPCER) and the Attack Presentation Classification Error Rate (APCER) as specified by ISO/IEC 30107-3 [16]. This relationship can be visualized using a Detection Error Trade-off (DET) curve.

For this work we adopt the public Face Morphing Detection benchmark utilities [21]¹ for the differential case. We develop the functionality for generation verification protocols in the differential pipeline and generate several protocols basing on the public data. Bona fide pairs in all the

protocols are combined from the frontal faces of the following public datasets: FRLI Set [10], FEI [2], Aberdeen and Utrecht [37] (~ 500 pairs in total). Morph pairs are combined by pairing images from the morphs from FRLI-Morphs dataset [34] and bona fides from FRLI Set [10]. We propose several protocols for different type of morphs (protocol names correspond to the FRLI-Morph subset names): *protocol-asml* ($\sim 4.5k$ morph pairs); *protocol-facemorpher* ($\sim 2.5k$ morph pairs); *protocol-webmorph* ($\sim 2.5k$ morph pairs); *protocol-stylegan* ($\sim 2.5k$ morph pairs).

Several benchmarks (with restricted data and protocols) available for evaluating the performance of morphing detection or morphing resistant algorithms: The NIST FRVT MORPH [26] and FVC-onGoing MAD [31]. They accept both no-reference and differential morphing algorithms, however they are proprietary and managed by a specific entity, leading to submission restrictions and limited accessibility. In this work we will use the public results of NIST FRVT MORPH to compare with our approach.

6. Experiments

6.1. Differential Benchmarking

We performed experiments of the fused classification strategy with the binary classification baseline and tested

¹<https://github.com/iurii-m/MorDeeply>

Table 2. Comparison with differential image morphing detection methods by $APCER@BPCER = (0.1, 0.01)$ from the NIST FRVT MORPH benchmark.

Method	$APCER@BPCER = \delta$							
	Visa-Border		Manual		MIPGAN-II		Print+Scan	
	$\delta = 0.1$	$\delta = 0.01$	$\delta = 0.1$	$\delta = 0.01$	$\delta = 0.1$	$\delta = 0.01$	$\delta = 0.1$	$\delta = 0.01$
Scherhag et al. [36]	0.013	0.212	0.055	0.357	0.004	0.134	0.012	0.176
Kashiani et al. [18]	0.447	0.901	0.873	0.989	0.182	0.481	0.842	0.996
Lorenz et al. [20]	0.432	1.000	0.634	1.000	0.168	1.000	0.732	1.000
Ferrara et al. [14]	0.966	0.999	0.689	0.969	0.004	0.751	0.070	0.280
Ours	0.232	0.555	0.531	0.872	0.359	0.859	0.680	0.926
Ours(FR)	0.087	0.453	-	-	-	-	0.125	0.568

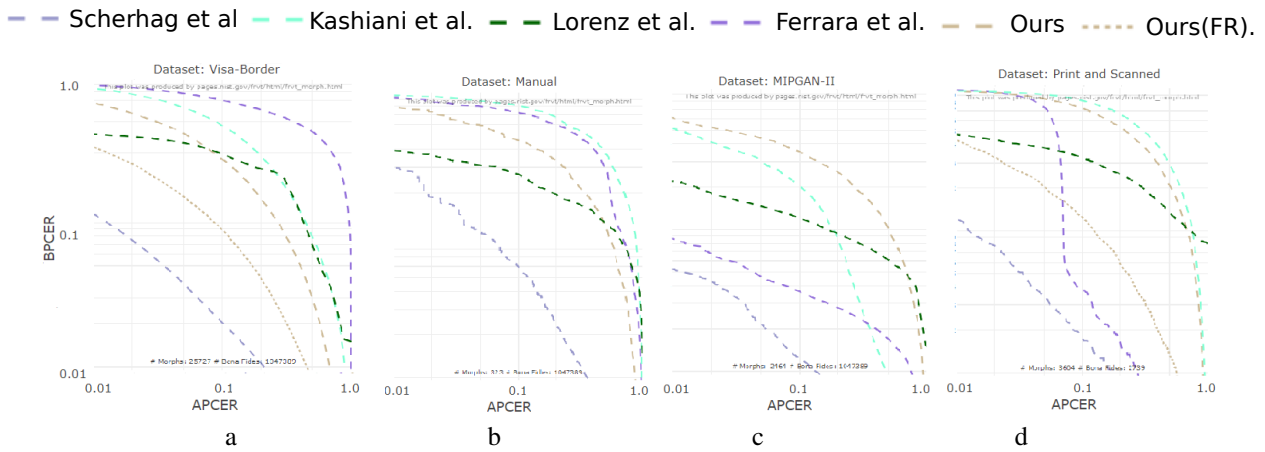


Figure 3. DET curves for the D-MAD algorithms by NIST FRVT MORPH benchmark. a) protocol Visa-Border; b) protocol Manual; c) protocol MIPGAN-II; d) protocol Print + Scanned.

those cases in our custom benchmarks. The baseline is implemented on the same setup (see Figure 1) where the identity classification components are disabled and the training is driven by a single loss component in Eq. 3.

In all the cases we use the EfficientNetB3 [42] backbone network with input image size 300×300 . It is trained with SGD optimizer for 5 epochs with momentum 0.9 and linearly decreasing learning rate from 0.01 to 0.0001. The batch size is 28.

Our results (see Figure 2, and Table 1) demonstrate the superiority of our approach over the baseline. Fused classification allow to generalize the detection performance to the unseen data and scenarios. Also we conclude that the V2 strategy (where morphs are disentangled from their original classes) is superior then the V1 and allow to achieve better MAD performance.

6.2. NIST FRVT MORPH

We evaluate the performance of our top-performing model (Fused Classification V3) by comparing it with several state-of-the-art (SOTA) D-MAD approaches, which have public results on the FRVT NIST MORPH Benchmark [26]. We perform comparison in several protocols: *Visa-Border* (25727 Morphs); *Manual* (323 Morphs); *MIPGAN-II* (2464 Morphs); *Print + Scanned* (3604 Morphs). All protocols in the comparison utilize a substantial collection of $\sim 1M$ bona fide images. The performance evaluation is conducted using the metrics $APCER@BPCER = (0.1, 0.01)$.

Our performance results (see Table 2, Figure 3) are comparable to the leaders in several benchmarks. Also, our method does not demonstrate bias to a particular morphing generative strategy and has the most stable performance across all protocols in comparison to other approaches.

We also present the algorithm (Ours(FR))(see Table 2), where the morphing detection signal of our fused classifica-

tion detector is multiplied with the similarity, which is given by a face recognition model. This algorithm indeed demonstrates a superior result in all the benchmarks, where it was tested. This indicates that currently differential face morphing benchmarks share many similarities with 1-1 face verification protocols and can be approached with only a strong face recognition model in hands. Such an approach is reasonable from the practical security perspective of detecting impostors, but not very correct from the academic perspective, since the face morphing detector in differential cases is prompted to detect identity non-matched pairs that may not have face morphs at all.

Despite the fact that this feature does not pose significant risks, it should be taken into account when developing algorithms for differential face morphing detection or benchmarks for their evaluation.

7. Conclusion

In this paper, we focus on differential face morphing detection and propose a novel deep learning method that incorporates sophisticated face recognition tasks and employs a fused classification scheme for morphs classification. We propose public benchmark utilities for differential face morphing detection. Additionally, we raise several questions on the differences in the vision of differential face morphing detection from academic and security application perspectives.

References

- [1] P. Aghdaie, B. Chaudhary, S. Soleymani, J. Dawson, and N. Nasrabadi. Attention Aware Wavelet-based Detection of Morphed Face Images. *2021 IEEE IJCB*, pages 1–8, 2021. [2](#)
- [2] Artificial Intelligence Laboratory of FEI in São Bernardo do Campo, São Paulo, Brazil. FEI face database, 2006. <https://fei.edu.br/cef/facedatabase.html>. [5](#)
- [3] Biometric System Laboratory. UBO-Morpher, 2018. <http://biolab.csr.unibo.it/Research.asp>. (accessed: September 1, 2022). [2](#)
- [4] G. Borghi, E. Pancisi, M. Ferrara, and D. Maltoni. A double siamese framework for differential morphing attack detection. *Sensors*, 21:3466, 05 2021. [3](#)
- [5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *International Conference on FG*, 2018. [4](#)
- [6] N. Damer, M. Fang, P. Siebke, J. N. Kolf, M. Huber, and F. Boutros. Mordiff: Recognition vulnerability and attack detectability of face morphing attacks created by diffusion autoencoders, 2023. [2](#)
- [7] N. Damer, A. M. Saladié, A. Braun, and A. Kuijper. Morgan: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Generative Adversarial Network. In *2018 IEEE 9th International Conference on BTAS*, pages 1–10, 2018. [2](#)
- [8] N. Damer, N. Spiller, M. Fang, F. Boutros, F. Kirchbuchner, and A. Kuijper. Pw-mad: Pixel-wise supervision for generalized face morphing attack detection. In *Advances in Visual Computing*, pages 291–304, 2021. [2](#)
- [9] L. Debiasi, U. Scherhag, C. Rathgeb, A. Uhl, and C. Busch. PRNU-based detection of morphed face images. *2018 IWBF*, pages 1–7, 2018. [2](#)
- [10] L. DeBruine and B. Jones. Face research lab london set, May 2017. https://figshare.com/articles/dataset/Face_Research_Lab_London_Set/5047666/3. [5](#)
- [11] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *2019 Conference on CVPR*, pages 4685–4694, 2019. [2](#)
- [12] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. *IJCB 2014 - 2014 IEEE/IAPR*, 12 2014. [2](#)
- [13] M. Ferrara, A. Franco, and D. Maltoni. Face demorphing. *IEEE Transactions on Information Forensics and Security*, 13(4):1008–1017, 2018. [3](#)
- [14] M. Ferrara, A. Franco, and D. Maltoni. Face morphing detection in the presence of printing/scanning and heterogeneous image sources. *IET Biometrics*, 10, 02 2021. [6](#)
- [15] Y. Huang, Y. Wang, Y. Tai, X. Liu, P. Shen, S. Li, J. Li, and F. Huang. CurricularFace: Adaptive Curriculum Learning Loss for Deep Face Recognition. In *The IEEE/CVF Conference on CVPR*, June 2020. [2](#)
- [16] International Organization for Standardization. ISO/IEC 30107–3:2017. Information Technology—Biometric Presentation Attack Detection — Part 3: Testing and Reporting. ISO/IEC JTC 1/SC 37 Biometrics, 09 2017. [5](#)
- [17] T. Karras, S. Laine, and T. Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *2019 IEEE/CVF Conference on CVPR*, pages 4396–4405, 2019. [2, 4](#)
- [18] H. Kashiani, S. M. Sami, S. Soleymani, and N. M. Nasrabadi. Robust ensemble morph detection with domain generalization. In *IEEE International IJCB 2022*, pages 1–10. IEEE, 2022. [6](#)
- [19] M. Kim, A. K. Jain, and X. Liu. Adaface: Quality adaptive margin for face recognition. *2022 IEEE/CVF Conference on CVPR*, pages 18729–18738, 2022. [2](#)
- [20] S. Lorenz, U. Scherhag, C. Rathgeb, and C. Busch. Morphing attack detection: A fusion approach. In *IEEE Fusion*, 2021. [6](#)
- [21] I. Medvedev, F. Shadmand, and N. Gonçalves. Mordeephy: Face morphing detection via fused classification. In *Proceedings of ICPRAM*, pages 193–204. SciTePress, 2023. [1, 2, 3, 4, 5](#)
- [22] I. Medvedev, J. Tremoço, B. Mano, L. E. Santo, and N. Gonçalves. Towards understanding the character of quality sampling in deep learning face recognition. *IET Biometrics*, 11(5):498–511, 2022. [2](#)
- [23] Q. Meng, S. Zhao, Z. Huang, and F. Zhou. Magface: A universal representation for face recognition and quality assessment. In *2021 IEEE/CVF Conference on CVPR*, pages 14220–14229, 2021. [2](#)

- [24] T. Neubert, C. Kraetzer, and J. Dittmann. A Face Morphing Detection Concept with a Frequency and a Spatial Domain Feature Space for Images on eMRTD. In *Proceedings of the ACM Workshop*, pages 95–100, 07 2019. 2
- [25] NIST. International Face Performance Conference, 2022. <https://www.nist.gov/news-events/events/2022/11/international-face-performance-conference-ifpc-2022>. (accessed: May 1, 2023). 1
- [26] NIST. NIST FRVT MORPH, 2022. https://pages.nist.gov/frvt/html/frvt_morph.html. 5, 6
- [27] Le Qin, Fei Peng, and Min Long. Face morphing attack detection and localization based on feature-wise supervision. *IEEE TIFS*, 17:3649–3662, 2022. 3
- [28] Thelen R. and Horchert J. Aktivisten schmuggeln Fotomontage in Reisepass, 2018. <https://www.spiegel.de/netzwelt/netzpolitik/biometrie-im-reisepass-peng-kollektiv-schmuggelt-fotomontage-in-ausweis-a-1229418.html>. (accessed: May 1, 2023). 1
- [29] R. Raghavendra, K. B. Raja, and C. Busch. Detecting morphed face images. In *2016 IEEE 8th International Conference on BTAS*, pages 1–7, 2016. 2
- [30] R. Raghavendra, K. B. Raja, S. Venkatesh, and C. Busch. Transferable Deep-CNN Features for Detecting Digital and Print-Scanned Morphed Face Images. *2017 IEEE Conference on CVPRW*, pages 1822–1830, 2017. 2
- [31] K. Raja, M. Ferrara, A. Franco, L. Spreeuwers, I. Batskos, F. Wit, M. Gomez-Barrero, U. Scherhag, D. Fischer, S. Venkatesh, J. M. Singh, G. Li, L. Bergeron, S. Isadskiy, R. Raghavendra, C. Rathgeb, D. Frings, U. Seidel, F. Knopjes, and C. Busch. Morphing Attack Detection - Database, Evaluation Platform and Benchmarking. *IEEE TIFS*, PP:1–1, 11 2020. 5
- [32] R. Ramachandra, S. Venkatesh, K. Raja, and C. Busch. Towards making morphing attack detection robust using hybrid scale-space colour texture features. In *2019 IEEE 5th International Conference on ISBA*, pages 1–8, 2019. 2
- [33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. B., A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. 2
- [34] E. Sarkar, P. Korshunov, L. Colbois, and Sébastien Marcel. Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks. *arXiv preprint*, Oct. 2020. 5
- [35] 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 *2017 5th IWBF*, pages 1–6, 2017. 2
- [36] U. Scherhag, C. Rathgeb, J. Merkle, and C. Busch. Deep Face Representations for Differential Morphing Attack Detection. *IEEE TIFS*, 15:3625–3639, 2020. 3, 6
- [37] School of Natural Sciences University of Stirling. Psychological Image Collection of Stirling, 1998. <http://pics.stir.ac.uk>. (accessed: September 1, 2022). 5
- [38] F. Schroff, D. Kalenichenko, and J. Philbin. FaceNet: A unified embedding for face recognition and clustering. In *2015 IEEE Conference CVPR*, pages 815–823, 2015. 2
- [39] Y. Shi, X. Yu, K. Sohn, M. Chandraker, and A. Jain. Towards Universal Representation Learning for Deep Face Recognition. In *Proceedings of 2020 IEEE/CVF CVPR*, pages 6816–6825, 06 2020. 2
- [40] J. Sun, W. Yang, J. Xue, and Q. Liao. An Equalized Margin Loss for Face Recognition. *IEEE Transactions on Multimedia*, pages 1–1, 2020.
- [41] Y. Sun, X. Wang, and X. Tang. Deeply learned face representations are sparse, selective, and robust. *2015 IEEE Conference on CVPR*, pages 2892–2900, 2015. 2
- [42] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 09–15 Jun 2019. 6
- [43] The Community Research and Development Information Service. Image Manipulation Attack Resolving Solutions, 2020. <https://cordis.europa.eu/project/id/883356>. (accessed: September 1, 2022). 1
- [44] Matjaž Torkar. Morphing cases in slovenia. NIST IFPS, 2022. Ministry of the Interior Police, Slovenia. 1
- [45] L. Wandzik, G. Kaeding, and R. V. Garcia. Morphing Detection Using a General-Purpose Face Recognition System. In *2018 26th EUSIPCO*, pages 1012–1016, 2018. 2
- [46] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A Discriminative Feature Learning Approach for Deep Face Recognition. In *Computer Vision – ECCV 2016*, pages 499–515, 2016. 2
- [47] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer, and C. Busch. MIPGAN - Generating Robust and High Quality Morph Attacks Using Identity Prior Driven GAN. *ArXiv*, abs/2009.01729, 2020. 2