

This is the peer reviewed version of the following article:

Combining Identity Features and Artifact Analysis for Differential Morphing Attack Detection / Di Domenico, Nicolò; Borghi, Guido; Franco, Annalisa; Maltoni, Davide. - 14233:(2023), pp. 100-111. (Proceedings of the 22nd International Conference on Image Analysis and Processing, ICIAP 2023 Udine, Italy September 11-15, 2023) [10.1007/978-3-031-43148-7_9].

Springer Science and Business Media Deutschland GmbH
Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

01/05/2026 01:01

(Article begins on next page)



Combining Identity Features and Artifact Analysis for Differential Morphing Attack Detection

Nicolò Di Domenico, Guido Borghi^(✉), Annalisa Franco, and Davide Maltoni

Dipartimento di Informatica - Scienza e Ingegneria (DISI), University of Bologna,
47521 Cesena, Italy

{nicolo.didomenico, guido.borghi, annalisa.franco, davide.maltoni}@unibo.it

Abstract. Due to the importance of the Morphing Attack, the development of new and accurate Morphing Attack Detection (MAD) systems is urgently needed by private and public institutions. In this context, D-MAD methods, *i.e.* detectors fed with a trusted live image and a probe tend to show better performance with respect to S-MAD approaches, that are based on a single input image. However, D-MAD methods usually leverage the identity of the two input face images only, and then present two main drawbacks: they lose performance when the two subjects look alike, and they do not consider potential artifacts left by the morphing procedure (which are instead typically exploited by S-MAD approaches). Therefore, in this paper, we investigate the combined use of D-MAD and S-MAD to improve detection performance through the fusion of the features produced by these two MAD approaches.

Keywords: Morphing Attack · Morphing Attack Detection · Differential MAD (D-MAD) · Single image MAD (S-MAD) · Feature Fusion

1 Introduction

Through an image morphing algorithm, it is possible to merge two images into one. In particular, this process can be applied to face images to create an intermediate one which includes facial characteristics of the two contributing subjects. A *Morphing Attack* [10] employs the aforementioned process to break the unique link between an official document and its owner: specifically, a subject with no criminal records (*accomplice*) can apply for a passport using a morphed mugshot picture to conceal the identity of a *criminal*. Indeed, several studies [29, 32] have shown the effectiveness of this attack, capable of fooling both the human control (*e.g.* a police officer) and the current commercial Facial Recognition Systems.

In particular, the morphing attack poses a significant security threat to Automated Border Control (ABC) gates located at international airports. These systems are designed to automatically verify the facial image stored in the electronic Machine Readable Travel Document (eMRTD) against a live image captured at the gate. Indeed, the presence of a morphed face can effectively bypass these

security checks, allowing both the criminal and the accomplice to pass through the gate. Therefore, it is essential to develop robust and efficient *Morphing Attack Detection* (MAD) algorithms [26] capable of detecting the presence of a morphed face automatically, not only when the document is used by the criminal – *i.e.* the primary task – but also by the accomplice.

Recently, several MAD methods have been proposed in the literature. Generally, these algorithms are classified into two families of approaches [3]: *Single image MAD* (S-MAD) and *Differential MAD* (D-MAD). S-MAD methods receive as input a single image, they examine only the potentially morphed mugshot picture and mainly rely on the potential traces (*e.g.* artifacts) left by the morphing process [2]. Differently, D-MAD methods compare the potentially morphed image against a trusted one and then their main hypothetical usage is at the airport gates, in which the document image is compared with the live-captured one. Two examples of input couples of a D-MAD system are reported in Fig. 1.




Fig. 1. Input example of a D-MAD system, consisting of a morphed image (center) and a live acquisition depicting the criminal (Subject 1) or the accomplice (Subject 2).

From a general point of view, D-MAD methods exhibit greater performance than S-MAD methods in detecting morphed mugshot photos [32]. Unfortunately, since D-MAD systems are mainly based on the comparison of the two face identities provided in input, their efficacy is worsened with input images are similar (*i.e.* the morphed image is created from look-alike subjects or the morphing factor privileges one of the contributing subjects). Therefore, we focus this work on the development of D-MAD methods that exploit also artifact-related information, in order to improve their performance with similar identities or even extend the use of D-MAD systems to the document enrollment procedure (in which the ID image is very similar to the applicant to fool the human examiner).

Starting from the observation that D-MAD methods usually tend not to consider the presence of artifacts left by the morphing procedure, we explore different strategies for combining S-MAD and D-MAD features: in particular, we investigate the performance of the SoA D-MAD approach [31] by introducing

an S-MAD module that operates only on the suspected morphed image. The underlying idea is that the S-MAD module can improve the final accuracy since it can detect visible or invisible artifacts produced by the morphing process that are normally overlooked by a more traditional D-MAD algorithm. Experimental results reveal that the proposed method improves the accuracy, especially in detecting morphed images in couples in which the accomplice is present (*i.e.* when the identity features extracted from the two images are very similar).

2 Proposed Method

The proposed method, depicted in Fig. 2, mainly consists of two different modules, *i.e.* S-MAD and D-MAD: the first module is responsible for the extraction of feature from the potentially morphed image, while the second one extracts features from the same image and the live probe. These features are then merged – through a feature fusion procedure investigated in the following sections – to create the input for the final classifier that produces in the output the final morphing detection score. As a classifier, we adopt a Multi-Layer Perceptron (MLP), with an architecture of 3 hidden layers of size 250, 125, and 64 with the sigmoid activation function in the final neuron. The MLP is trained using the BCE loss function, Adam [14] as optimizer with an initial learning rate of $5 \cdot 10^{-4}$ and an early stopping procedure in order to prevent overfitting, stopping the training after 5 epochs without a minimum improvement of 10^{-3} in the validation loss.




Fig. 2. Overview of the proposed method. As shown, S-MAD and D-MAD modules extract different features that are fused together and used by the final MLP classifier.

2.1 S-MAD Module

The S-MAD module consists of a backbone, specifically we adopt an Inception-Resnet V1 [33] architecture, pre-trained on the VGG-Face2 [5] dataset. In particular, the model is fine-tuned on several morphing datasets obtained with various morphing algorithms (described in Sect. 3.1), producing images ranging from low to medium quality. Moreover, we run a supplemental fine-tuning process to improve the algorithm’s performance on heavily compressed, ICAO-compliant [35], JPEG images. A 512 dimensional feature vector is finally obtained removing the last fully connected layer of the architecture exploited

for the classification task. For the training procedure, we adopt the Stochastic Gradient Descent (SGD) with a learning rate of 10^{-3} and the early-stopping procedure exploited for the training of the whole method. No momentum decay is exploited. This module is developed leveraging the Revelio framework¹.

2.2 D-MAD Module

As the D-MAD module, we take inspiration from the solution proposed in [31] that can be regarded as the current state of the art, as also shown in the results published on the FVC-onGoing platform [1].

Specifically, we use a ResNet-50 [12] network trained for the face recognition task [23] through the ArcFace loss [7], to extract the embeddings of the facial input images. Since the input is represented by two images, this module outputs two different embeddings of size 512 that are combined through a subtraction, and then the final feature is represented by a single embedding with the same size of 512. Authors show that the ArcFace loss function produces robust embeddings since it tends to maximize the geodesic distance between different identities and that the produced embeddings contain therefore information exclusively related to the input face identity. Our implementation is based on the publicly available Deepface² framework.

3 Experimental Validation

3.1 Datasets

For our experimental evaluation, we employ several publicly available datasets, with varying quality levels, briefly described and discussed in the following.

- *Progressive Morphing Database* (PMDb) [11]: it is a collection of 1108 morphed images generated by applying a public morphing algorithm to AR [18], FRGC [22], and Color Feret [21] datasets. The dataset contains 280 subjects, divided into 134 males and 146 females. No manual retouching procedures have been applied to enhance the visual quality of the images. As a result, the images may contain artifacts such as blurred areas or ghosts.
- *Idiap Morph* [27,28]: it is a collection of five datasets created using different morphing algorithms (OpenCV [17], FaceMorpher [24], StyleGAN [13], WebMorph [6], and AMSL [20]), created starting from face images from the Feret [21], FRGC [22], and Face Research Lab London Set (FRL) [6] datasets. The visual quality of the morphed images generated with the OpenCV and FaceMorpher morphing algorithms is negatively affected by artifacts in both the background and foreground. Morphed faces generated with the StyleGAN algorithm exhibit typical GAN-related textures [37]. The AMSL morphing algorithm is used to generate 2175 morphed images from 102 adult faces: these images are compressed to a maximum size of 15 kB, simulating the process to embed an image into the chip of the eMRTD.

¹ <https://miatbiolab.csr.unibo.it/revelio-framework>.

² <https://github.com/serengil/deepface>.

- *MorphDB* [11]: it is a dataset of 100 morphed images generated using the Squirz Morph 2.1 [36] algorithm, applied to images from the Color Feret [21] and FRGC [22] datasets. This dataset is composed of 50 male and 50 female subjects. As all images are manually retouched, their visual quality is excellent. This element makes this dataset particularly challenging, although the limited number of images may make it unsuitable for conducting an extensive performance review of a MAD algorithm.
- *FEI* [34]: it is a dataset generated using the images contained in the *FEI Face Database*, which includes 200 subjects, equally split between male and female. All faces are mainly represented by subjects between 19 and 40 years old with distinct appearances, hairstyles, and accessories. This dataset contains 6000 morphed images obtained with three different morphing algorithms, namely FaceFusion [9], UTW [25], and NTNU [25], employing two different morphing factors (0.3 and 0.5).

3.2 Experimental Protocol

Experimental results are split into three distinct scenarios, according to the identity of the trusted live image: i) *Criminal*: contains bona fide attempts (*i.e.* the document image is not morphed) and morphed attempts where the document image is morphed and the live image belongs to the criminal subject; ii) *Accomplice*: contains bona fide and morphed attempts in which the live image comes from the accomplice. iii) *Both*: this scenario contains all the couples belonging to the criminal and accomplice ones.

A sample for each kind of couple can be found in Fig. 1.

In all the following results, we focus on the performance obtained in the *Accomplice* scenario that represents the practical case in which the accomplice presents the morphed image for the enrollment procedure. As mentioned, due to the greater similarity between the subjects present in both pictures, these experiments are generally considered more challenging than the previously mentioned one (*Criminal*) for D-MAD methods.

3.3 Metrics

To evaluate and compare MAD systems, there are several metrics commonly used for assessing their performance [30]: *Bona Fide Presentation Classification Error Rate* (BPCER), which represents the proportion of bona fide images incorrectly classified as morphed, and *Attack Presentation Classification Error Rate* (APCER), which represents the proportion of morphed images incorrectly labeled as bona fide. They are formulated as follows:

$$\text{BPCER}(\tau) = \frac{1}{N} \sum_{i=1}^N H(b_i - \tau), \quad \text{APCER}(\tau) = 1 - \left[\frac{1}{M} \sum_{i=1}^M H(m_i - \tau) \right] \quad (1)$$

In both definitions, τ is the score threshold on which b_i, m_i , the detection scores, are compared; $H(x) = \{1 \text{ if } x > 0, 0 \text{ otherwise}\}$ is defined as a step function. Typically, the BPCER is measured with respect to a given APCER value, *i.e.* $B_{0.1}$, $B_{0.05}$ and $B_{0.01}$, representing the lowest BPCER with APCER $\leq 10\%$, $\leq 5\%$, $\leq 1\%$, respectively. Ideally, a MAD algorithm employed in a real-world setting would need to operate at a low APCER (*i.e.* letting almost no criminals through) of around 0.1%, while maintaining an acceptable corresponding BPCER (*i.e.* generating few false positives) of around 1%.

The *Equal Error Rate* (EER), *i.e.* the error rate for which both BPCER and APCER are equal, is usually reported as a single value.

4 Experimental Results

4.1 S-MAD and D-MAD Module Assessment

Firstly, we assess the performance of S-MAD and D-MAD modules separately. For the S-MAD module, we test the same network described in Sect. 2.1, while in the D-MAD module we add the MLP architecture described in Sect. 2.2 that acts as a classifier. In this manner, we aim to understand the detection capabilities of each module, and these results offer a useful baseline to better analyze the performance of the proposed method in the following experiments.

Results are reported in Table 1. As expected, the metrics obtained through the S-MAD module are identical regardless of the type of couple, as in both cases only the same suspected morphed images are used. Besides, the D-MAD module provides considerably better performance than the S-MAD one when the live image contains the criminal, indicating that a trusted, live-capture image proves to be effective in tackling the task by comparing the two input identities. Moreover, we observe that not only the performance gap between S-MAD and D-MAD is nearly canceled when the accomplice is present in the live-capture image, but also the D-MAD performance is significantly worsened (about +10% in EER): we prove that the greater similarity between the identities makes the classification task more challenging.

Table 1. Morphing detection scores obtained on the FEI test set. Results are reported in terms of Equal Error Rate (EER), the lowest BPCER related to APCER $\leq 10\%$, $\leq 5\%$, and $\leq 1\%$, respectively.

Module	Accomplice			Criminal			Both		
	EER	$B_{0.05}$	$B_{0.01}$	EER	$B_{0.05}$	$B_{0.01}$	EER	$B_{0.05}$	$B_{0.01}$
S-MAD	.186	.360	.515	.186	.360	.515	.186	.360	.515
D-MAD	.180	.470	.827	.085	.147	.447	.141	.343	.767

4.2 Investigation on Feature Fusion

Previous results suggest the opportunity of exploring the combination of S-MAD and D-MAD classifiers so that the overall performance of the system does not only rely on the identity present in the live image. Therefore, we test the proposed method (see Sect. 2), merging the input features through different approaches described as follows. In addition, we also test our method by replacing the MLP with an SVM classifier, trained using an RBF kernel with a $C = 3$ regularization factor and a γ kernel coefficient which is inversely proportional to the variance of the training data received in input. This choice has been driven by the use of SVM in many morphing-related works in the literature that suggest the importance of this type of classifier in the MAD field.

Firstly, we start our investigation with a simple concatenation for the feature fusion produced by the S-MAD and D-MAD modules. Results are reported in the first line of Table 2 and show that the concatenation (indicated with the letter C) provides results that are similar to those obtained with the S-MAD module in the accomplice scenario. This behavior suggests that S-MAD features have a strong impact on the classification, outweighing the features provided by D-MAD and thus negating the benefits they bring in comparing the identities.

Table 2. Morphing detection scores obtained on the FEI test set across different classifiers and feature fusion techniques. C stands for concatenation, MM for Min-Max, and MV for Mean-Variance (see Sect. 4.2).

Fusion	Class.	Accomplice			Criminal			Overall		
		EER	B _{0.05}	B _{0.01}	EER	B _{0.05}	B _{0.01}	EER	B _{0.05}	B _{0.01}
C	MLP	.168	.317	.510	.168	.345	.520	.168	.330	.515
	SVM	.175	.248	.458	.175	.275	.460	.175	.265	.458
MM	MLP	.132	.245	.463	.138	.265	.463	.135	.260	.463
	SVM	.140	.278	.475	.160	.320	.475	.147	.295	.475
MV	MLP	.175	.338	.543	.185	.398	.563	.181	.365	.555
	SVM	.195	.317	.495	.257	.423	.530	.225	.387	.522

To further investigate our hypothesis, we run a *t-distributed Stochastic Neighbor Embedding* (t-SNE) [16] dimensionality reduction on the input features, divided both by source (*i.e.* D-MAD or S-MAD) and by class (*i.e.* bona fide or morphed). The resulting plot, shown in Fig. 3, highlights how the features can easily be separated by their respective source, suggesting that they may occupy different portions of the feature space. However, there is no clear separation between bona fide and morphed feature vectors; this reinforces the hypothesis that the classifier could be prioritizing the S-MAD features while disregarding those generated by the D-MAD module. Moreover, to test how the high dimensionality of the two concatenated vectors might affect the system's

performance, we employ the PCA algorithm [19], which indicates that the optimal intrinsic feature dimensionality is only one less than the original. Thus, we infer that all features of both vectors may be required to achieve the best results and that other fusion methods must be investigated.

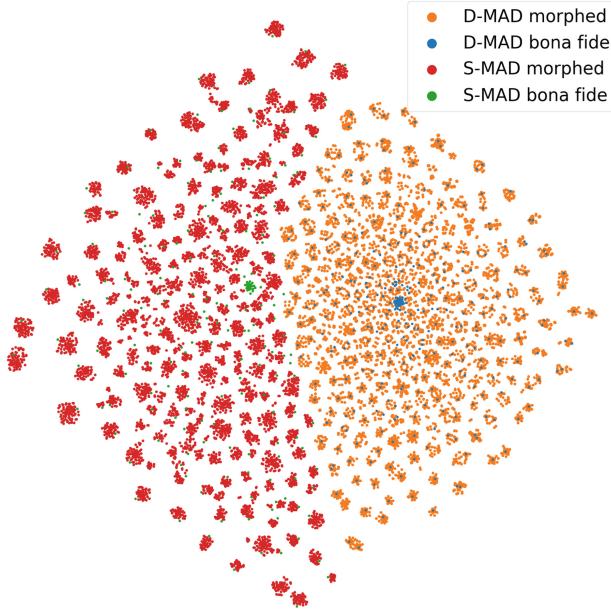


Fig. 3. Visualization of the *t*-distributed Stochastic Neighbor Embedding (t-SNE) [16] of D-MAD and S-MAD feature vectors divided by ground truth (best on screen).

Therefore, we investigate two additional fusion strategies: i) *Min-Max* (MM): before concatenating the two feature vectors, they are separately rescaled to have each component in the $[0, 1]$ range; ii) *Mean-Variance* (MV): before concatenating the two feature vectors, they are separately rescaled to have each component with mean value $\mu = 0$ and variance $\sigma = 1$.

Results of the above-mentioned experiments are reported in Table 2. We observe the performance gap that was previously found between MLPs and SVMs is not present when the two features are merged together; indeed, the former almost always outperforms the latter. Besides, the Mean-Variance strategy provides unsatisfactory results, which are worse than the simple concatenation strategy. An ex-post numerical analysis on the normalized feature vectors used for training shows that, even when each component is rescaled to have $\mu = 0$ and $\sigma = 1$, the D-MAD and S-MAD features still show significant differences in range. This could be a possible explanation for the great performance of the Min-Max fusion strategy, thus proving that translating the two feature vectors to the same numeric range helps improve the model’s performance.

In addition to this investigation, we also test the performance impact of including the cosine similarity (C_S) [15] between the two D-MAD feature vectors. The underlying idea comes from the fact that, as explained in [7], the embeddings produced by the model are optimized so that the geodesic angle between each identity is maximized. Therefore, the cosine similarity between the embeddings obtained from both the suspected morphed and live images should be approximately 1 when no morphing algorithm is applied; on the contrary, if the similarity is closer to -1, then we can assume that the two presented identities are too far apart and therefore some morphing process has taken place.

To determine if there is a tangible performance improvement, we train an MLP whose input is composed of both the D-MAD and S-MAD features with the best fusion strategy found, *i.e.* the *Min-Max*, as well as the cosine similarity between the two original embeddings produced by the ArcFace [7] loss. Moreover, inspired by the chosen fusion strategy, we investigate whether to translate the cosine similarity from its $[-1, 1]$ range to $[0, 1]$. Experimental results are reported in Table 3: they show that adding the cosine similarity provides a tangible performance improvement only when left in its original range. On the contrary, if the cosine similarity is translated into the $[0, 1]$ range the model’s performance is considerably worsened.

Table 3. Morphing detection scores obtained on the FEI test set with and without employing the cosine distance C_S . “-” symbol denotes that the range of the distance has kept unchanged.

C_S range	Accomplice			Criminal			Overall		
	EER	B _{0.05}	B _{0.01}	EER	B _{0.05}	B _{0.01}	EER	B _{0.05}	B _{0.01}
-	.132	.245	.463	.138	.265	.463	.135	.260	.463
$[-1, 1]$.125	.237	.468	.125	.235	.440	.125	.235	.445
$[0, 1]$.132	.237	.465	.141	.290	.470	.136	.270	.468

4.3 Comparison with the State of the Art

Finally, we test the performance of our proposed method against the current D-MAD literature methods. We compare our proposed algorithm (in particular, we select the best configuration obtained, *i.e.* the Min-Max feature fusion and the cosine similarity in the $[-1, 1]$ range) against the methods proposed in [4, 31]. Experimental results are reported in Table 4.

It is worth noting that the proposed method overcomes both competitors when the identity in the live image belongs to the accomplice, consistently in all metrics reported, suggesting that S-MAD features can effectively improve the performance of D-MAD methods. However, when the criminal is present, the proposed algorithm has still room for improvement indicating the need to investigate further feature fusion methods and to develop specific MAD techniques to differently address image pairs with criminals and accomplices.

Table 4. Morphing detection scores obtained on the FEI test set through the proposed methods with respect to the current literature solutions.

Method	Accomplice			Criminal			Overall		
	EER	B _{0.05}	B _{0.01}	EER	B _{0.05}	B _{0.01}	EER	B _{0.05}	B _{0.01}
[31]	.175	.475	.780	.066	.085	.310	.129	.343	.690
[4]	.153	.345	.563	.060	.095	.370	.115	.257	.515
Ours	.125	.237	.468	.125	.235	.440	.125	.235	.445

Lastly, we also test the proposed method through the FVC-onGoing platform [8] on the sequestered DMAD-SOTAMD_D-1.0 benchmark, even though in this dataset the available morphed images are only compared with the criminal or bonafide subjects. We obtain an EER of about 10%, a worse result with respect to the method [31] that is SotA in couples with criminal (with an EER = 4.5%), but also a better result with respect to the algorithm proposed in [11] (EER = 14%), showing that machine learning-based techniques yield overall better results and represent promising solutions. These results suggest the need for the development of a strategy able to select the best MAD algorithm: specifically, this system should be able to detect whether the criminal or the accomplice is present in the live image, and then either use, for instance, the standalone state-of-the-art D-MAD algorithm (*e.g.* [31]) or the method proposed in this paper. In this manner, we may be able to overcome the limitations of both systems, which respectively underperform when the live image contains the accomplice or the criminal.

5 Conclusions and Future Works

In this paper, we have investigated the fusion of S-MAD and D-MAD features, to create a system that is capable of detecting morphed images in the challenging scenario in which the accomplice is used for comparison. Experimental results reveal that effectively combining the two kinds of embeddings is not a trivial task. Specifically, we demonstrate that the features produced by S-MAD and D-MAD methods occupy different regions of the feature space, and their normalization in a predefined range improves the model’s overall effectiveness. Moreover, we show that including a further feature represented by the cosine distance between the two embeddings produced by the D-MAD feature extractor improves the algorithm’s performance. As a future work, we plan to improve the overall performance by developing a model able to preliminarily discriminate between couples with the accomplice or the criminal, thus enabling the selective use of a specific method to address the D-MAD task.

Acknowledgment. This work is part of the iMARS project. The project received funding from the European Union’s Horizon 2020 research and innovation program under Grant Agreement No. 883356. Disclaimer: this text reflects only the author’s views, and the Commission is not liable for any use that may be made of the information contained therein.

References

1. Biolab: FVC-onGoing. <https://biolab.csr.unibo.it/fvcongoing/>
2. Borghi, G., Franco, A., Graffieti, G., Maltoni, D.: Automated artifact retouching in morphed images with attention maps. *IEEE Access* **9**, 136561–136579 (2021)
3. Borghi, G., Graffieti, G., Franco, A., Maltoni, D.: Incremental training of face morphing detectors. In: *2022 26th International Conference on Pattern Recognition (ICPR)*, pp. 914–921. IEEE (2022)
4. Borghi, G., Pancisi, E., Ferrara, M., Maltoni, D.: A double siamese framework for differential morphing attack detection. *Sensors* **21**(10), 3466 (2021)
5. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: *13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi'an, China, May 15–19 (2018)*
6. DeBruine, L., Jones, B.: Face research lab London set. *Psychol. Methodol. Des, Anal* (2017)
7. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019 (2019)*
8. Dorizzi, B., et al.: Fingerprint and on-line signature verification competitions at ICB 2009. In: *Advances in Biometrics, Third International Conference, ICB 2009, Alghero, Italy, June 2–5, 2009. Proceedings. Lecture Notes in Computer Science*, vol. 5558 (2009)
9. FaceFusion: Facefusion. <https://www.wearemoment.com/FaceFusion/>
10. Ferrara, M., Franco, A., Maltoni, D.: The magic passport. In: *IEEE International Joint Conference on Biometrics, Clearwater, IJCB 2014, FL, USA, September 29 - October 2, 2014 (2014)*
11. Ferrara, M., Franco, A., Maltoni, D.: Face demorphing. *IEEE Trans. Inf. Forensics Secur.* **13**(4) (2018)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778 (2016)
13. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of StyleGAN. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8110–8119 (2020)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings (2015)*. <https://arxiv.org/abs/1412.6980>
15. Li, B., Han, L.: Distance weighted cosine similarity measure for text classification. In: Yin, H., et al. (eds.) *IDEAL 2013. LNCS*, vol. 8206, pp. 611–618. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-41278-3_74
16. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(11) (2008)
17. Mallick, S.: Face morph using OpenCV - C++ / Python. <https://learnopencv.com/face-morph-using-opencv-cpp-python/>
18. Martinez, A., Benavente, R.: The AR face database: Cvc technical report, 24 (1998)
19. Minka, T.: Automatic choice of dimensionality for PCA. In: *Advances in Neural Information Processing Systems*, vol. 13 (2000)
20. Neubert, T., Makrushin, A., Hildebrandt, M., Kraetzer, C., Dittmann, J.: Extended stirtrace benchmarking of biometric and forensic qualities of morphed face images. *IET Biometrics* **7**(4), 325–332 (2018)

21. Phillips, P.J., Wechsler, H., Huang, J., Rauss, P.J.: The FERET database and evaluation procedure for face-recognition algorithms. *Image Vision Comput.* **16**(5) (1998)
22. Phillips, P.J., et al.: Overview of the face recognition grand challenge. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 1. IEEE (2005)
23. Pini, S., Borghi, G., Vezzani, R., Maltoni, D., Cucchiara, R.: A systematic comparison of depth map representations for face recognition. *Sensors* **21**(3), 944 (2021)
24. Quek, A.: FaceMorpher morphing algorithm. <https://github.com/alyssaq/face-morpher>
25. Raja, K., et al.: Morphing attack detection-database, evaluation platform, and benchmarking. *IEEE Trans. Inf. Forensics Secur.* **16**, 4336–4351 (2020)
26. Raja, K.B., et al.: Morphing attack detection-database, evaluation platform, and benchmarking. *IEEE Trans. Inf. Forensics Secur.* **16**, 4336–4351 (2021)
27. Sarkar, E., Korshunov, P., Colbois, L., Marcel, S.: Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks. arXiv preprint [arXiv:2012.05344](https://arxiv.org/abs/2012.05344) (2020)
28. Sarkar, E., Korshunov, P., Colbois, L., Marcel, S.: Are gan-based morphs threatening face recognition? In: ICASSP 2022–2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2959–2963. IEEE (2022)
29. Scherhag, U., Debiassi, L., Rathgeb, C., Busch, C., Uhl, A.: Detection of face morphing attacks based on PRNU analysis. *IEEE Trans. Biom. Behav. Identity Sci.* **1**(4) (2019)
30. Scherhag, U., Rathgeb, C., Merkle, J., Breithaupt, R., Busch, C.: Face recognition systems under morphing attacks: a survey. *IEEE Access* **7**, 23012–23026 (2019)
31. Scherhag, U., Rathgeb, C., Merkle, J., Busch, C.: Deep face representations for differential morphing attack detection. *IEEE Trans. Inf. Forensics Secur.* **15** (2020)
32. Scherhag, U., 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, Darmstadt, Germany, September 20–22, 2017. LNI, vol. P-270 (2017)
33. Szegedy, C., et al.: Going deeper with convolutions. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015 (2015)
34. Thomaz, C., Giraldi, G.: A new ranking method for principal components analysis and its application to face image analysis. *Image and Vision Comput.* **28**(6), 902–913 (2010)
35. Wolf, A.: ICAO: Portrait quality (reference facial images for MRTD), version 1.0. standard. International Civil Aviation Organization (2018)
36. xiberpix: Sqirlz morphing algorithm. <https://sqirlz-morph.it.uptodown.com/windows>
37. Zhang, X., Karaman, S., Chang, S.F.: Detecting and simulating artifacts in GAN fake images. In: 2019 IEEE International Workshop on Information Forensics and Security (WIFS), pp. 1–6. IEEE (2019)