Morph-PIPE: Plugging in Identity Prior to Enhance Face Morphing Attack Based on Diffusion Model

Haoyu Zhang[1], Raghavendra Ramachandra[1], Kiran Raja[1], and Christoph Busch[1,2]

[1] Norwegian University of Science and Technology, Department of Information Security and Communication Technology, NTNU Gjøvik Teknologiveien 22, 2815 Gjøvik, Norway. * haoyu.zhang@ntnu.no

[2] Hochschule Darmstadt, Schöfferstraße 3, 64295 Darmstadt, Germany

Abstract. Face-morphing attacks (MA) aim to deceive Face Recognition Systems (FRS) by combining the face images of two or more subjects into a single face image. To evaluate the vulnerability of existing FRS and further develop countermeasures against potential attacks, it is necessary to create diverse morphing algorithms that produce high visual quality and have strong attack potential on FRS. In this work, we propose a novel morphing algorithm using a diffusion model and adding identity prior to strengthening attack potential on the FRS. Compared to existing works using diffusion models, our method can add explicit control of the morph generation process through identity manipulation. We benchmark our proposed approach on an ICAO-compliant face morphing dataset against state-of-the-art (SOTA) morphing algorithms, including one baseline using the diffusion model and two representative morphing algorithms. The results indicate an improvement in the performance of the morphing attack potential compared to the baseline algorithm using diffusion while it achieves comparable attack strength to other SOTA morphing generation algorithms which rely on tedious manual intervention in the creation of morphed images.

Keywords: Face Morphing Attack · Diffusion Models · Morph Generation · Vulnerability Analysis · Morphing Attack Potential

1 Introduction

Face recognition systems have been widely deployed in various security applications, such as automatic border control (ABC) [17]. However, it has been shown that the improvement in the generalizability and recognition performance of the FRS also makes them more vulnerable to different types of attacks [22] [28]. A face morphing attack deceives a face recognition system by generating a

* This work was supported by the European Union's Horizon 2020 Research and Innovation Program under Grant 883356.

single photo that is combined with face images from two or more subjects. Face-morphing attacks from a generation and detection perspective have become a central topic in research on FRS [10, 26] to secure the identity control process. A face image submitted by an applicant to obtain a valid passport or visa leaves the opportunity for a malicious actor to enrol manipulated/morphed images by finding an accomplice whose identity resembles very closesly[28].

To evaluate the vulnerability of existing FRS against morphing attacks and develop robust Morphing Attack Detection (MAD) algorithms, it is necessary to have diverse morph generation methods which can generate high-quality morphs. The morphed images of high quality can be used for investigating the vulnerability of FRS and further the development of detection solutions. The strength of face morphing attacks on FRS is mainly based on two aspects: the visual/perceptual image quality (e.g., image sharpness and no ghost artifacts) and the morphing attack potential (MAP) on FRS (the success rate when using morphs to attack FRS). Traditional morph generation methods are based on the alignment of facial landmark points [2] [25] [21] [3] [4]; however, they usually generate many ghost artifacts. Improved methods [11] have been developed using post-processing techniques, such as color equalization and swapping of the face region. With the development of deep learning, researchers have proposed the use of Generative Adversarial Networks (GAN) [13] to generate morphs [6] [29] [30] [8] [20]. Meanwhile, Denoising Diffusion Probabilistic Models (DDPM) [14] have been proposed and have achieved considerable diversity and generalizability in image generation compared to GANs, which can be suitable for morphing biometric samples. Recent works have used Diffusion Models to generate morphs [5] [7] by manipulating the latent code.

However, existing approaches for generating morphs using diffusion are based only on the averaged interpolation of the stochastic code and high-level semantic code without prioritizing the identity information. As a result, this does not guarantee linear mapping between the change in identity information and the interpolation of the latent code. In other words, the mid-point of the interpolation may not result in a morphed face that is the most similar morph to both the contributing subjects. Explicitly adding control of the identity prior can increase the morphing attack potential of generated morphs [30]. However, the diffusion-based models for generating morphs have not considered the identity prior to increasing the attack strength. Meanwhile, generating an image using GANs only needs to forward the latent vector through the pre-trained generator, but diffusion models require several steps of inference to gradually denoise into an image. It is computationally costly to add control of identity prior by end-to-end optimization on the morphed latent vector as MIPGAN [30].

Hence, we propose an approach to plug in the identity prior to enhancing the morphing attack potential of a diffusion-based morphing algorithm (MorphPIPE) by optimizing the interpolation factor when fusing the latent representation of face images. An overview of the approach is shown in Figure 1 and will be detailly described in Section 3. To evaluate the performance of our proposed method, we create a new face morphing dataset using face images from FRGC

[23] and generate morphs based on our proposed method and other baseline methods. The attack strength of our method against other SOTA approaches is benchmarked by vulnerability analysis on FRS. We summarize our contributions to this work as follows:

- A novel morphing method termed 'Morph-PIPE' using a diffusion model with efficient control of identity prior is proposed. The identity control is added to optimize and search for the interpolation factor when morphing latent vectors, targeting to increase the possibility of successful attack.
- We conduct an extensive vulnerability analysis on two different FRS using standardized metric - morphing attack potential [12] and benchmark the results with 3 other morphing algorithms on the FRGC-based dataset [23]. The results indicate a considerable attack potential among all morphing algorithms. Compared with the baseline SOTA method using the Diffusion Model, the proposed method showed an improvement in the effectiveness of attacking the FRS. Compared to other SOTA methods, the proposed method also achieved comparable performance in terms of attack potential. The implementation will be released only with agreements for ethical usage [3].

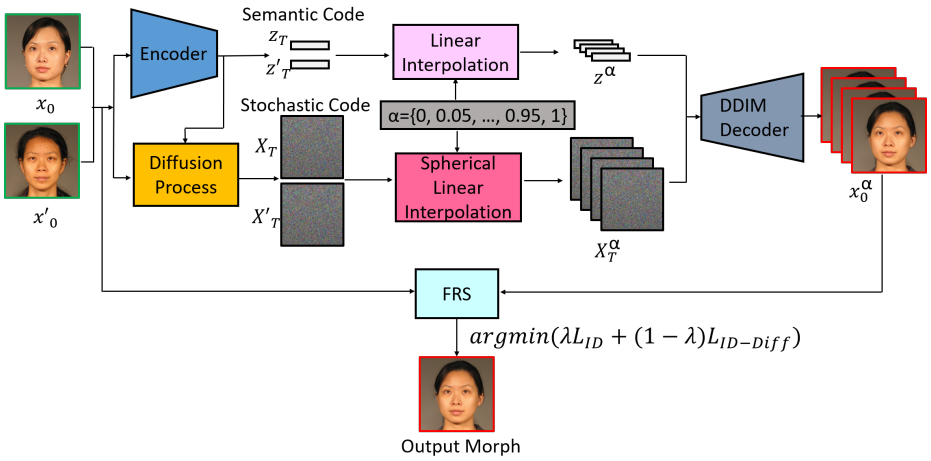

Fig. 1. Overview of our proposed Morph-PIPE method using identity control.

[3] https://share.nbl.nislab.no/HaoyuZhang/face-morphing-attack-ntnu-morph-pipe-public

2 Background

The naive method of morphing is simply averaging the RGB values of each pixel in the image, which would result in severe ghost artifacts. The traditional method of automatically generating morphing attacks is based on aligning the face content by the help of facial landmark points (e.g., eyes, nose, and mouth). Subsequently, based on the correspondence of the same landmark point from different subjects, the pixels can be warped and averaged locally. There are various warping methods, including free-form deformation, deformation by moving least squares, and deformation based on mass–spring [28]. The most common method for generating morphing attacks is Delaunay triangulation [2], where the same pair of triangles is affine transformed and pixel values are blended with a factor. In morphing attacks, morphing multiple identities makes it more challenging to achieve an effective attack; hence, it is usually assumed to generate a morph between two subjects with a contributing factor (i.e., blending factor) of 0.5. Due to differing geometry of facial structures, landmark-based morphing algorithms result in morphed images with ghost artifacts, especially around eye regions, nostrils and hair regions. Hence, many works subsequently have focused on automated post-processing (e.g., crop-pasting of the face region and color equivalence) [11].

Also, morphs can be generated by deep learning techniques. GANs are image generation models that learn the mapping from a known distribution to the data distribution in an adversarial training fashion, where the known distribution can be considered as the latent space of the model. By morphing in the latent space instead of in the pixel space, ghost artifacts can be avoided. Hence, the GAN-based morphing methods can be divided into three parts: projecting the image sample into the latent space, interpolation in the latent space, and reconstruction. Damer et al. [6] first proposed to use Bi-directional GAN where an encoder is included to generate morphs in which the visual quality and generalizability are highly limited due to the performance of the backbone model. Subsequently, Venkatesh et al. [29] applied StyleGAN [18] with a more disentangled latent space and generated morphs with high image resolution and visual quality. However, interpolation in the latent space cannot be fully mapped using the interpolation of identity information. To create a strong morphing attack on the FRS, Zhang et al. [30] proposed an end-to-end optimization framework with identity loss referring to both subjects simultaneously. Damer et al. [8] have also explored using a hybrid method where the landmark-based morph is first generated and then reconstructed by GAN to eliminate the ghost artifacts. However, the reconstruction process also introduces additional noise to the identity information. Kelly et al. [20] proposed further disentangling the identity information within the latent space of StyleGAN and studied the case of attacking a known FRS. In general, the advantage of a GAN-based morph is the possibility of including an identity prior as a control, and there are fewer ghost artifacts than landmark-based morphing algorithms. However, the GAN-based morphs contain specific GAN traces. Overall, the attacking potential is lower than that of landmark-based methods [30].

Recently, diffusion models have achieved considerable performance in the field of image generation, especially in conditioned image generation after the Denoising Diffusion Probabilistic Models (DDPM) paper [14] for its high generalizability. The Diffusion Model constructs a parameterized Markov chain and learns to forward diffuse from the data sample to the noise and inverse process denoising back to the image. Considering the generation of morphs as an image generation task conditioned on face images from contributing subjects, it is worth studying the diffusion model for generating morphs. The DDIM [24] includes a semantic encoder to encode high-level semantics from an image and control the diffusion process. Damer et al [7] and Blasingame et al. [5] proposed the use of the DDIM to generate morphs by linear and spherical interpolation of the latent code with a factor $\alpha = 0.5$. The latter work also studied different interpolation configurations for fusing the stochastic latent code. Compared to GANs, Diffusion models are more computationally costly when inferencing because it takes several steps of image generation for a single output image when reversing the image from noise. To enhance the attack potential of these diffusion-based morphing algorithms, we propose our PIPE method to plug in identity prior and control the morphing process.

3 Proposed Method

The overview of our proposed method is illustrated in Figure 1. The existing SOTA Diffusion-based methods are based on simple interpolation in the latent space with a fixed interpolation factor $\alpha = 0.5$. There's no explicit identity control during the morphing process. Meanwhile, in GAN-based morphing algorithms, identity control is typically applied to the latent code using gradient information traced back from its reconstructed image. However, this is challenging owing to the high computational cost of sequential inferences in Diffusion Models. Hence, we propose an add-on method, where we generate a set of morphs in parallel, given a set of blending factors $\alpha = \frac{n}{N-1}, n = 0, 1, ..., N-1$, where N is the total number of interpolation steps. Finally, we used a pre-trained ArcFace [9] face recognition system as the identity prior to searching for the best *alpha* factor based on the defined metrics.

Overall, the detailed algorithm can be described as follows. First, a pair of contributing images x_0, x'_0 are encoded into semantic codes using pre-trained encoder from [24]:

$$z = Enc(x_0), z' = Enc(x'_0), \tag{1}$$

and then together diffused into stochastic code X_T, X'_T. Linear interpolation is applied to blend the semantic condition:

$$z^\alpha = lerp(z, z', \alpha), \tag{2}$$

and spherical interpolation is applied to blend latent code as suggested by Song et al. [27]:

$$X_T^\alpha = slerp(X_T, X'_T, \alpha). \tag{3}$$

Further, condition z^{α} is used to control at each time T during the denoising diffusion process from noises X_T^{α} to morphs x_0^{α}. Finally, alpha is chosen such that the loss function based on extracted identity prior can be minimized: $argmin(\lambda \cdot L_{ID}(x_0^{\alpha}, x_0, x_0') + (1-\lambda) \cdot L_{ID-Diff}(x_0^{\alpha}, x_0, x_0'))$, where weight factor λ

$$L_{ID}(x_0^{\alpha}, x_0, x_0') = 1 - \frac{FRS(x_0^{\alpha}) \cdot FRS(x_0)}{\|FRS(x_0^{\alpha})\| \|FRS(x_0)\|} + 1 - \frac{FRS(x_0^{\alpha}) \cdot FRS(x_0')}{\|FRS(x_0^{\alpha})\| \|FRS(x_0')\|}, \quad (4)$$

$$L_{ID-Diff}(x_0^{\alpha}, x_0, x_0') = \left| \frac{FRS(x_0^{\alpha}) \cdot FRS(x_0')}{\|FRS(x_0^{\alpha})\| \|FRS(x_0')\|} - \frac{FRS(x_0^{\alpha}) \cdot FRS(x_0)}{\|FRS(x_0^{\alpha})\| \|FRS(x_0)\|} \right|. \quad (5)$$

These two loss functions are inspired by [30], where L_{ID} measures how similar the generated morphs are to each contributing subject, and $L_{ID-Diff}$ measures how balanced the similarity of generated morphs is to each contributing subject. This overall process indicates that we optimize the morphing factor α by searching in its sampling space to minimize the defined loss function based on identity prior. For other hyperparameters, in this work we empirically select the interpolation step $N = 21$ and weight factor $w = 0.5$.

4 Experiments and Results

In this section, we describe the experimental settings and benchmark results. Figure 2 shows an example of the dataset and morphs generated using different algorithms. For the morphing algorithms,

- LMA-UBO [11]: an advanced landmark-based morphing algorithm with post-processing including copy-pasting of face region.
- MIPGAN-II [30]: end-to-end optimization framework based on StyleGAN2 [19] as the backbone.
- MorDiff [7]: state-of-the-art morphing algorithm using Diffusion Models, the baseline of our method.
- Morph-PIPE: proposed morphing algorithm, based on MorDiff and plugged in identity prior to enhance the attack potential.

We construct a new dataset based on FRGCv2 [23] dataset, including in total 140 subjects, over 1000 mated samples, and 2500 morph pairs per morphing algorithm. The morphing pairs were selected without crossing the gender and subject pairing based on the similarity score within the dataset as detailed in an earlier work[30].

To address sustainability and reduce computational overhead, our algorithm is based on the existing DDIM model weights [24] trained on the FFHQ [18] dataset with 256×256 image resolution. ArcFace was used to perform the final filtering step. It should be noted that the ArcFace FRS we used for optimizing the interpolation factor and the ArcFace FRS in vulnerability analysis are not the same model.

4.1 Evaluation Protocol and Metrics

We use the Morphing Attack Potential (MAP) [12], a standardized metric in ISO/IEC CD 20059 [16] to measure the possibility of a successful attack on multiple FRSs with multiple attempts. Following the configuration proposed by the original authors [12], the DeepFace library [1] is used to implement the evaluation protocol. Four FRSs were selected: ArcFace, Dlib, Facenet, and VGGface. As our dataset is selected to be compliant with ICAO 9303 [15] and hence to be representative of passport quality, the number of mated samples varies. However, the MAP metric requires the same number of mated samples for each morphing subject pair, so the minimum number of mated samples has to be selected. Hence in this experiment, each morph is tested in a single attempt for each contributing subject and in a total of two attempts.

4.2 Results

The quantified benchmarking results are shown in the four tables in Figure 3. For each entry $MAP[r, c]$ in the tables, this indicates the possibility that one morphed image successfully attacked r times at all c FRSs. A larger MAP value indicates a higher morphing attack potential and a stronger morphing attack generation algorithm.

For visual quality demonstrated in Figure 2, the proposed method can generate morphs with high visual fidelity. As Morph-PIPE only optimizes the interpolation factor in the MorDiff framework, it achieves a similar performance of visual quality with minor differences. Compared to MIPGAN-II, diffusion-based methods have shown to be more natural in jawline, hairline, and skin texture. However, due to the complexity of the backbone model, the image resolution of morphs generated by these two diffusion-based algorithms is currently limited to 256×256 pixels while MIPGAN-II supports 1024×1024. For morphs generated by the landmark-based LMA-UBO algorithm, their skin texture is the most natural than other selected deep-learning-based morphs, however, distortion during warping and swapping can still be noticed when the difference of landmark points and skin tone within contributing faces is obvious.

Figure 3 shows the tables of MAP evaluation for different morphing algorithms. The table follows the visualization example in [12]. In each entry of a table, $MAP[r, c]$ reports the proportion of morphs that can successfully achieve a match decision against at least r attempts for each contributing subject by at least c FRSs. It is shown that by using the proposed Morph-PIPE method, the generated attacks were stronger or the same for all cases compared to the baseline MorDiff model. This demonstrates the effectiveness of explicitly adding identity control using the proposed methods to make the baseline algorithm more general and robust.

Among all the morphing algorithms benchmarked in this experiment, attacks generated by MIPGAN-II were the strongest for single attempts. However, the LMA-UBO method showed better generality at four FRS and was stronger for all cases given two attempts. Our proposed method is weaker than MIPGAN-II and LMA-UBO but remains comparable.

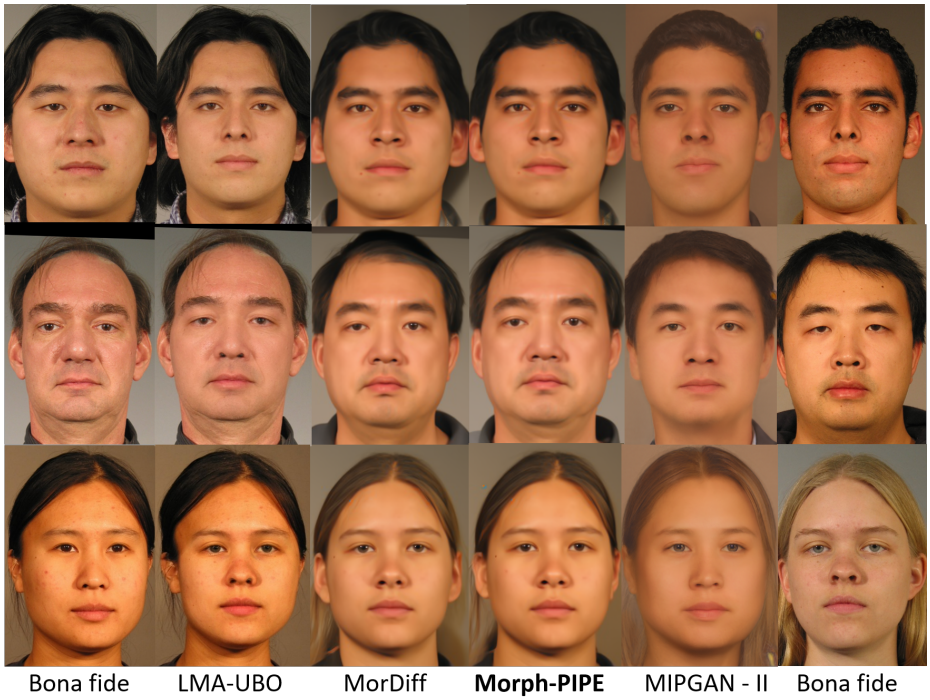

Fig. 2. Examples of our morphing dataset: Each row demonstrates a pair of subjects with their contributing bona fide data and morphs generated by different morphing algorithms.

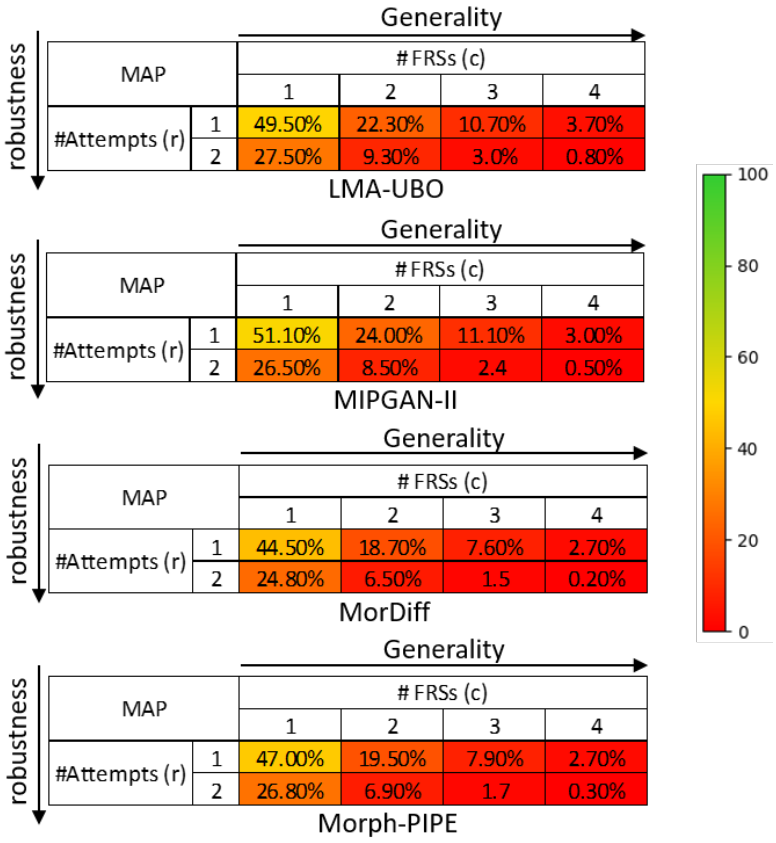

Fig. 3. MAP benchmarking results of LMA-UBO, MIPGAN-II, MorDiff, and Morph-PIPE.

4.3 Limitations

Owing to computational limitations, we only targeted the method without the inference of the diffusion model during optimization. As it is reasonable to assume that an attacker will attempt to generate a single morphing attack, it would also be helpful to investigate identity control by latent mapping or end-to-end optimization, or fine-tune the backbone diffusion model on an existing dataset. It would also be interesting to generalize the proposed Morph-PIPE method to optimize the morphing factor in other types of morphing algorithms, instead of diffusion-based algorithms.

5 Conclusion

Noticing that the existing SOTA lacks explicit identity control, we propose an efficient Morph-PIPE method to use the FRS model as an identity prior and optimize the morphing factor by searching and minimizing the defined loss term for each morph pair. Finally, we benchmarked the proposed method with the baseline and other selected representative morphing algorithms. The results indicate an improved attack potential in all cases compared with the baseline model without using Morph-PIPE. The benchmarking results of the proposed method are also comparable to those of other SOTA morphing algorithms. Overall, diffusion models are suitable for generating high-quality morphs with minor artifacts and considerable attack potential. The proposed Morph-PIPE method can be efficient and effective in enhancing the existing diffusion-based morphing algorithm, which also demonstrates the benefit of adding explicit identity control for the generation of morphs. Future work could involve spending more computational resources and exploring identity control based on the inference of the diffusion model during optimization.

6 Ethical Considerations

This work has proposed a morphing attack generation algorithm, while the intention is to illustrate the potential risk by analysing the vulnerability of existing open-source face recognition systems and contributing to the development of robust morphing attack detection algorithms in future works. We condemn any behaviour of misusing the proposed technique to create misleading content or actual manipulation attacks. The implementation of this work will be released with agreements only for ethical scientific usage.

References

1. Deepface website, https://github.com/serengil/deepface/
2. Face morph using opencv (2017), http://www.learnopencv.com/face-morph-using-opencv-cpp-python/
3. Facemorpher (2018), https://github.com/yaopang/FaceMorpher/

4. Webmorph (2018), https://webmorph.org/
5. Blasingame, Z., Liu, C.: Diffusion models for stronger face morphing attacks. arXiv preprint arXiv:2301.04218 (2023)
6. Damer, N., Saladié, A.M., Braun, A., Kuijper, A.: Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network. In: 2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS). pp. 1–10 (Oct 2018). https://doi.org/10.1109/BTAS.2018.8698563
7. Damer, N., Fang, M., Siebke, P., Kolf, J.N., Huber, M., Boutros, F.: Mordiff: Recognition vulnerability and attack detectability of face morphing attacks created by diffusion autoencoders. arXiv preprint arXiv:2302.01843 (2023)
8. Damer, N., Raja, K., Süßmilch, M., Venkatesh, S., Boutros, F., Fang, M., Kirchbuchner, F., Ramachandra, R., Kuijper, A.: Regenmorph: Visibly realistic gan generated face morphing attacks by attack re-generation. In: Advances in Visual Computing: 16th International Symposium, ISVC 2021, Virtual Event, October 4-6, 2021, Proceedings, Part I. pp. 251–264. Springer (2021)
9. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4690–4699 (2019)
10. Ferrara, M., Franco, A., Maltoni, D.: The magic passport. In: IEEE International Joint Conference on Biometrics. pp. 1–7. IEEE (2014)
11. Ferrara, M., Franco, A., Maltoni, D.: Decoupling texture blending and shape warping in face morphing. In: 2019 International Conference of the Biometrics Special Interest Group (BIOSIG). pp. 1–5. IEEE (2019)
12. Ferrara, M., Franco, A., Maltoni, D., Busch, C.: Morphing attack potential. In: 2022 International Workshop on Biometrics and Forensics (IWBF). pp. 1–6. IEEE (2022)
13. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM **63**(11), 139–144 (2020)
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems **33**, 6840–6851 (2020)
15. INTERNATIONAL CIVIL AVIATION ORGANIZATION: PORTRAIT QUALITY (REFERENCE FACIAL IMAGES FOR MRTD) (2017), date published: 2017-04-01
16. ISO/IEC JTC1 SC37 Biometrics: ISO/IEC CD 20059. Vulnerability of Biometric Recognition Systems with Respect to Morphing Attacks. International Organization for Standardization (2023)
17. Jain, A.K., Li, S.Z.: Handbook of face recognition, vol. 1. Springer (2011)
18. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
19. 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)
20. Kelly, U.M., Spreeuwers, L., et al.: Worst-case morphs: a theoretical and a practical approach. In: 2022 International Conference of the Biometrics Special Interest Group (BIOSIG). pp. 1–5. IEEE (2022)
21. 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)

22. Ngan, M., Ngan, M., Grother, P., Hanaoka, K., Kuo, J.: Face recognition vendor test (frvt) part 4: Morph-performance of automated face morph detection. US Department of Commerce, National Institute of Standards and Technology (2023), https://pages.nist.gov/frvt/reports/morph/frvt_morph_report.pdf
23. Phillips, P.J., Flynn, P.J., Scruggs, T., Bowyer, K.W., Jin Chang, Hoffman, K., Marques, J., Jaesik Min, Worek, W.: Overview of the face recognition grand challenge. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). pp. 947–954 vol. 1 (June 2005). https://doi.org/10.1109/CVPR.2005.268
24. Preechakul, K., Chatthee, N., Wizadwongsa, S., Suwajanakorn, S.: Diffusion autoencoders: Toward a meaningful and decodable representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10619–10629 (2022)
25. Raghavendra, R., Raja, K.B., Venkatesh, S., Busch, C.: Face morphing versus face averaging: Vulnerability and detection. In: IEEE International Joint Conference on Biometrics (IJCB). pp. 555–563 (2017)
26. Raja, K., Ferrara, M., Franco, A., Spreeuwers, L., Batskos, I., de Wit, F., Gomez-Barrero, M., Scherhag, U., Fischer, D., Venkatesh, S.K., Singh, J.M., Li, G., Bergeron, L., Isadskiy, S., Ramachandra, R., Rathgeb, C., Frings, D., Seidel, U., Knopjes, F., Veldhuis, R., Maltoni, D., Busch, C.: Morphing attack detection-database, evaluation platform, and benchmarking. IEEE Transactions on Information Forensics and Security **16**, 4336–4351 (2021). https://doi.org/10.1109/TIFS.2020.3035252
27. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
28. Venkatesh, S., Ramachandra, R., Raja, K., Busch, C.: Face morphing attack generation and detection: A comprehensive survey. IEEE Transactions on Technology and Society **2**(3), 128–145 (2021). https://doi.org/10.1109/TTS.2021.3066254
29. Venkatesh, S., Zhang, H., Raghavendra, R., Raja, K., Damer, N., Busch, C.: Can gan generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection. In: 2020 International Workshop on Biometrics and Forensics (IWBF). pp. 1–6. IEEE (2020)
30. Zhang, H., Venkatesh, S., Ramachandra, R., Raja, K., Damer, N., Busch, C.: Mipgan—generating strong and high quality morphing attacks using identity prior driven gan. IEEE Transactions on Biometrics, Behavior, and Identity Science **3**(3), 365–383 (2021)