

Can GAN Generated Morphs Threaten Face Recognition Systems Equally as Landmark Based Morphs?

- Vulnerability and Detection

Sushma Venkatesh[†], Haoyu Zhang[†], Raghavendra Ramachandra[†], Kiran Raja[†], Naser Damer[‡], Christoph Busch[†]

[†]Norwegian Biometrics Laboratory, Norwegian University of Science and Technology (NTNU), Norway

[‡]Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany.

{sushma.venkatesh; haoyu.zhang; raghavendra.ramachandra;kiran.raja;christoph.busch}@ntnu.no
{naser.damer}@igd.fraunhofer.de

Abstract—The primary objective of face morphing is to combine face images of different data subjects (e.g. an malicious actor and an accomplice) to generate a face image that can be equally verified for both contributing data subjects. In this paper, we propose a new framework for generating face morphs using a newer Generative Adversarial Network (GAN) - StyleGAN. In contrast to earlier works, we generate realistic morphs of both high-quality and high resolution of 1024×1024 pixels. With the newly created morphing dataset of 2500 morphed face images, we pose a critical question in this work. (i) *Can GAN generated morphs threaten Face Recognition Systems (FRS) equally as Landmark based morphs?* Seeking an answer, we benchmark the vulnerability of a Commercial-Off-The-Shelf FRS (COTS) and a deep learning-based FRS (ArcFace). This work also benchmarks the detection approaches for both GAN generated morphs against the landmark based morphs using established Morphing Attack Detection (MAD) schemes.

Index Terms—Biometrics, Face Recognition, Morphing, Attack, Vulnerability

I. INTRODUCTION

Due to the widespread deployment of biometric-based identification and verification of individuals, it is essential to observe a biometric characteristic that is reliable, user-friendly and easy to capture. Face biometrics is well suited for this purpose due to its popularity and widespread use for biometric authentication. Moreover we consider the ease of capturing from a distance in a non-intrusive manner and also the recently achieved high recognition accuracy. These properties further enable that face recognition is to be used in various applications that are attributed with high security requirements like border control. However, biometric face recognition systems (FRS) are known to be highly vulnerable to presentation attacks (aka., spoofing attacks) against the capture device [14]. In addition face recognition system can be deceived during the enrolment process by providing manipulated images [4].

Among the different types of attacks against FRS, face morphing has gained momentum because of the high impact it poses on border control security. The morphing process




Fig. 1: Comparison of morphed images generated using Landmark (LM) , MorGAN and StyleGAN

enables a malicious actor to generate a morphed image by using an accomplice's face image in a seamless manner [4]. The process introduces a significant threat to the border control scenario as it is easy to obtain a passport document with a morphed image. This fact is also due to the limitations of current passport issuance protocols in which digital images are submitted in a self-supervised manner by an applicant for passport renewal through web services in countries like New Zealand, Estonia and Ireland. In other countries there exists no live enrolment in the passport renewal process, on the contrary the facial image is provided by an applicant in printed form and is subsequently scanned and re-digitized. This leaves an opportunity for the applicant to morph the face image prior to submitting it in the passport application.

A. Morph generation process and limitations

Early work on face morphing attacks [4] demonstrated the vulnerability of FRS with respect to morphed facial images, while to the same extend human experts could be fooled [15] [5]. Following the recent works towards detecting morph attacks on both digital image and re-digitized (print-scanned) image [11]–[13] we must state that this area of research is still

in a premature state. The crucial part of a morphing attack is the generation of high quality morphed facial image, which is ICAO compliant and can attack a deployed FRS with high probability. In the literature, there exist two different ways of generating morphed face images namely (a) Landmark based morphed face generation and (b) Deep learning-based morphed face generation. In landmark-based morph generation, given two images, the landmarks of both facial images are obtained and the Delaunay triangulation is generated for both images. Subsequently alpha blending is performed to obtain a single morphed image based on averaged Delaunay triangles. The majority of the recently published literature is based on open-source morphing tools [12] which are based on landmark constrained Delaunay triangles.

A deep learning-based approach in contrast involves synthesizing a morphed face image by using a Generative Adversarial Network (GAN). Limited works are reported in the literature on using GAN for morph generation [3]. The first reported work in this direction is based on the MorGAN [3] in which morphed images are generated corresponding to a image resolution of 64×64 pixels. Recently, the morphed images generated using MorGAN were super-resolved to have a incrementally larger dimension of 120×120 pixels [2]. It is important to note that the images generated using both approaches incorporating GAN [3] [2] are not ICAO compliant and hence have very limited use in real-life attacks. Irrespective of the morph image generation approach, it is essential that one needs to generate a high-quality image that can pose a high threat potential, when presented to a human expert in the control procedure while the passport issuance is carried out or to surpass a FRS during Automatic Border Control (ABC) crossing scenario.




Fig. 2: Illustration of minimal artifacts in morphed images generated using StyleGAN versus landmark based face morphing.

Motivated to address the limitation of low quality images generated by the previous GAN architectures, in this work, we present a new approach to generate high quality morph images. The recent improvement made in GAN architectures has enabled us to generate a high quality facial images with a resolution of 1024×1024 pixels using StyleGAN [8]. This is achieved by embedding the images into latent space which

is further optimized to synthesize the high quality and high resolution image [1]. As illustrated in Figure 1 the morphed images generated using StyleGAN can be observed to be superior in terms of quality, resolution and visual depiction.

Further, as noted from Figure 2, a number of artifacts can be easily handled in an automatic manner with the newly proposed approach, which is capable of suppressing visual artifacts. The clear superiority of the newly proposed approach can be noted around the iris regions, where double edges are inherently dealt with. While it well known fact that landmark based morphs threaten FRS to a high degree [4], [12], one can easily conclude the amount of extra time and resources that is anticipated to make the morphs visually appealing by removing the artifacts.

While the superior quality of face images can be achieved through the newly proposed approach and eventually reaching compliance to ICAO standards, we raise some fundamental questions.

- Despite the high quality of morphed images do they scale up to threaten a FRS in the same manner as the landmark based morphs, which typically exhibit large artifacts?
- To what degree can current MAD mechanisms detect such GAN based attacks on FRS, when the processing is limited to the digital domain?

In the course of answering the above questions, we can summarize the contributions of this work as follows:

- A new approach to generate morphed face images using the StyleGAN is presented.
- A new face morphing dataset comprising of $2500 \times 3 = 7500$ morphed images is generated using the StyleGAN and MorGAN approach. In order to compare the new approach using the GAN methodology, this work also constructs a corresponding landmark based morph dataset.
- To quantify the threats from GAN based morphed face images, a comprehensive vulnerability analysis is conducted using both, a commercial FRS (COTS) and an open-source FRS (ArcFace).
- In order to give an insight into the detection challenges of such attacks, this work also reports a detailed evaluation of MAD mechanisms on both GAN based and landmark based morphed face images.

In the rest of the paper, Section II describes the morph generation process proposed in this work using StyleGAN. Section III provides the details regarding the quantitative experiments indicating the vulnerability of FRS and the detection challenge. With remarks on future works in this direction, we draw the conclusion in Section IV.

II. MORPHED FACE GENERATION USING STYLEGAN

In this section, we present the StyleGAN based face morph generation to achieve high quality face morphs. Figure 3 depicts the block diagram of the proposed framework for the morphed face generation using a StyleGAN architecture [8]. Given the latent code L_1 of the faces, the StyleGAN

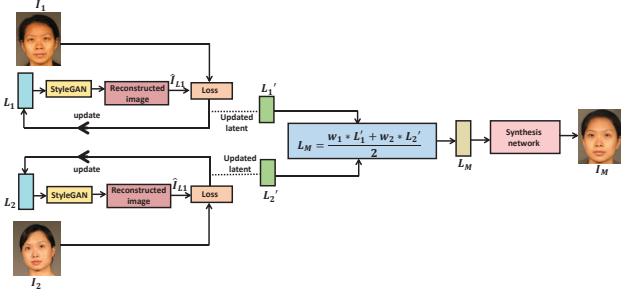


Fig. 3: Block diagram of the morphed face image using StyleGAN

[1] maps the inputs to an intermediate latent space (W) through the mapping network. The mapping layer consists of 8 fully connected layers that are serially connected. In this work, we force a strategy to embed the face image into the latent space (W_f), which is inspired by earlier work [1]. This process enables us to synthesize the data-subject-specific morphed face. The embedded latent space for a particular face is then passed through the synthesis network consisting of 18 layers, in order to control the adaptive instance normalization (AdaIN). As a direct result, we obtain the representation in 18 multiple latent spaces, each with a dimension of 512, which is further concatenated. For a given face image I_1 , the embedding is carried out by optimizing a loss function that measures the similarity between I_1 and the reconstructed image \hat{I}_{L_1} using the corresponding latent code L_1 . To maintain the perceptual fidelity a loss is computed as the weighted combination of VGG-16 perceptual loss [1] as given below:

$$PL = \min \sum_{i=1}^4 \frac{\lambda_i}{N_i} \|F_j(\hat{I}_{L_1}) - F_j(I_1)\|_2^2 \quad (1)$$

Where, F_j is the feature output of VGG-16 layer $conv1_1$, $conv1_2$, $conv3_2$ and $conv4_2$ respectively, $\lambda_i = 1$ and N_j is the number of scalars in the j^{th} layer. The optimization is carried out using Adam optimizer with a $\beta_1 = 0.5$.

We have selected the perceptual loss based on the visual quality of the morphed image that can reflect the suitability for border control applications. Let the final reconstructed image correspond to I_1 and L_1 be \hat{I}_{L_1} and the corresponding updated latent code be L_1' . We follow the same procedure mentioned above for the second image I_2 to get again a reconstructed image \hat{I}_{L_2} and the corresponding updated latent code denoted L_2' . The morphing operation is carried out by averaging the latent code as follows:

$$L_M = \frac{w_1 * L_1' + w_2 * L_2'}{2} \quad (2)$$

Finally, L_M is passed through the synthesis network to generate the morphed image that has a resolution of 1024×1024 pixels, where w_1 and w_2 indicate the weights, which we have chosen to be $w_1 = w_2 = 0.5$.

A. Differences of proposed approach with earlier works

In contrast to earlier works [3], to avoid the bias of morph generation with known set (closed-set), the StyleGAN is trained using the disjoint face dataset from FFHQ dataset [8] consisting of high quality face images. As it can be observed from Figure 1, the morphed face images generated using StyleGAN have higher perceptual fidelity as compared to MorGAN based morphed images and are equally comparable to landmark based morphed generation. It can be noticed that, the MorGAN [3] based morph generation indicates low-quality images that are not ICAO compliant rendering them not suitable for passport applications. As a secondary note, the MorGAN based images also indicate a poor visual similarity to the contributing subjects, while landmark based morphs exhibit stronger artifacts that are clearly visible in Figure 1.

Intrigued by the high fidelity of morphed face images, we take a detailed analysis guided by a sample image to compare it against the landmark based morph generation. As observed in Figure 2, the ghosting artifacts in landmark based morphing can be prominently seen due to the misalignment of landmarks leading to several artifacts, especially in the ocular, mouth and nose region. It is interesting to observe that the proposed StyleGAN based morph generation did not create any perceptual noise. The example demonstrates the high quality of the generated image, when compared to the MorGAN based approach.

While contrary to a landmark based morphed faces, the proposed StyleGAN based morph generation does not indicate a strong geometrical resemblance as it is the case for a landmark based approach. Motivated by such visual observations and superior quality of morph images achieved with the proposed approach and accounting for the lower geometrical resemblance of contributing subjects, we conduct a detailed analysis of threats to FRS as detailed in the next section.

III. EXPERIMENTS AND RESULTS

In order to measure the impact of the proposed approach of morph generation, we first create a new dataset of morphed images created from 140 unique data subjects. With the newly generated morph dataset, we first investigate and report the vulnerability of FRS and compare it with the vulnerability reported in similar earlier work using MorGAN [3] and traditional landmark based morphing. Further, we also analyze the detection potential of morphed faces generated using the proposed framework with StyleGAN.

A. Database Generation

We introduce a new morphed face database created from 140 individuals that include 47 female and 93 male data subjects. The facial images are derived from the FRGC-V2 face database [10]. The newly generated database is subdivided into two sets for training and testing that consists of independent data subjects with no overlap between the splits. The training set consists of 690 bona fide images and 1190 morphed images. The testing set consists of 580 bona fide and 1310 morphed images. To effectively analyze the vulnerability

and provide a comparison to earlier works, we have generated morph images using three different techniques, which include (i) Landmark-Based (ii) MorGAN and (iii) proposed StyleGAN approach. Care is exercised to generate morphed images with similar facial appearance within same gender category. Additionally, to guarantee high quality of the newly generated dataset constraints of high quality illumination and no pose is imposed before creating the morphs. The guidelines laid out in earlier works [12] [16] are followed to obtain a database of high relevance for morphing attack detection.

B. Evaluation Metrics for Vulnerability Analysis

We measure the vulnerability of FRS following the guidelines of Frontex and setting the operating threshold to FAR = 0.1 % (for both FRS). We further follow the realistic evaluation protocol where the morph image is created by using two face images corresponding to a malicious actor and an accomplice. We compute the vulnerability by enrolling a given morphed face image $M_{I_1,2}$ and probing the corresponding contributing subjects I_1 and I_2 with an image from a different FRGC-session. We further obtain the comparison scores S_1 and S_2 for both images I_1 and I_2 against the morphed image. The morphed image $M_{I_1,2}$ is only considered a threat if and only if the comparison scores S_1 and S_2 succeed to cross the preset threshold at FAR = 0.1%. If the condition is not met, we simply consider that the morphed image is not a real threat as the comparison scores are not able to successfully verify the morphed image against both contributing subjects making the morphing attack not realistic. We term this new metric as *Fully Mated Morphed Presentation Match Rate (FMMPMR)* and compute it in general form as:

$$FMMPMR = \frac{1}{P} \sum_{M,P} (S1_M^P > \tau) AND (S2_M^P > \tau) \dots AND (Sk_M^P > \tau) \quad (3)$$

Where $P = 1, 2, \dots, p$ represent the number of attempts made by presenting all the probe images from the contributing subject against M^{th} morphed image, $K = 1, 2, \dots, k$ represents the number of contributing data subjects to the constitution of the generated morphed image (in our case $K = 2$), Sk_M^P represents the comparison score of the K^{th} contributing subject obtained with P^{th} attempt (in our case the P^{th} probe image from the dataset) corresponding to M^{th} morph image and τ represents the threshold value corresponding to FAR = 0.1%.

When compared to the existing metric MMPMR [16], the FMMPMR considers the number of attempts (that are assessed jointly with contributing subjects) with regards to each face morphed images and thus reflect the realistic vulnerability of a FRS. The MMPMR [16] is designed to measure vulnerability only on the morphed image in a joint set rather than a number of attempts on each morphed image. Hence, the MMPMR fails to reflect the number of attempts (by contributing subjects) made against the corresponding morphing image to determine the vulnerability of FRS.

In this work, the COTS threshold is set at $\tau = 0.5$ based on the NIST FRVT test reports as recommended by the COTS provider while ArcFace FRS threshold is set at $\tau = 0.36$ base on the face recognition trials on FRGC-v2 dataset. The higher the value of the FMMPMR the higher the threat from morphed images and correspondingly a higher vulnerability of FRS towards morphed images must be stated.

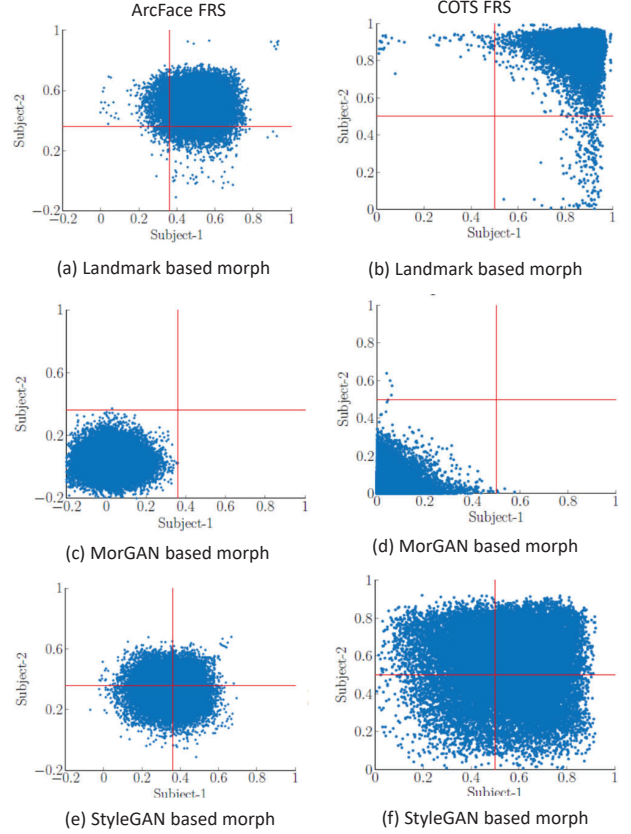


Fig. 4: Vulnerability analysis using COTS and ArcFace. The scatter plots represents the comparison scores of morphed face image against two contributing subjects. The red lines indicate the threshold corresponding to FAR = 0.1%

C. Results from Vulnerability analysis

In this section, we present the vulnerability analysis using two different Face Recognition Systems (FRS) (i) a Commercial off the Shelf face recognition system (COTS), Cognitec FaceVACS-SDK Version 9.4.2 ¹ and (ii) an Open-source deep learning based FRS (ArcFace). To effectively benchmark the results we also compare two different State-Of-The-Art (SOTA) morph generation techniques such as landmark based morph generation [6] and MorGAN based morph generation [3].

Figure 4 shows the scatter plot of the comparison scores obtained from two different FRS on images obtained using

¹outcome not necessarily constitutes the best the algorithm can do

three different types of morphed face generation approaches. Table I indicates the quantitative values of both FMMPMR and MMPMR computed from two different FRS for all three cases of face morph generation techniques. Based on the obtained results the key observations made are listed below:

TABLE I: Vulnerability analysis FMMPMR(%) and MMPMR(%)

Morphing Type	FMMPMR (%)		MMPMR (%)	
	ArcFace	COTS	ArcFace	COTS
Landmark based Morph [6]	86.04	98.91	95.76	100
MorGAN [3]	0	0	0	0
StyleGAN	21.1	41.49	39	64.68

The following are the main observations from our experiments.

- Landmark based face morph generation indicates a high threat to FRS (analogously high vulnerability of FRS to such images) compared to that of two other morph generation methods. This can be attributed to the fact that the landmark based morph generation preserves both texture and geometrical structure of the morphed image corresponding to its contributing subjects.
- The analysis of the experimental results also show that MorGAN based morph generated images do not pose a severe threat to FRS. The potential reason for this can be due to low quality generated morph image (64×64 pixels). A careful observation of the images also revealed the degradation of texture and geometry in the generated morphed images. As a caveat, we note that the MorGAN network is not re-trained (or fine-tuned) on closed dataset of contributory subjects. The conscious choice was made to investigate the generalisation of GANs for morph face image generation and study the threats.
- StyleGAN based morph generation method shows relatively higher degree of threats when compared with MorGAN. Despite higher threats, the images from proposed approach of morph generation did not compete against the landmark based methods. An introspection into this indicates the quality difference of FFHQ dataset versus the employed FRGC-V2 dataset. Specifically, the pre-trained morph generator is trained on FFHQ dataset which has very different characteristics than FRGC-V2 dataset leading the network to mimic the characteristics of the FFHQ dataset. Another aspect for the lower degree of threat is due to lack of geometric correspondence of facial structure in morphed faces when compared to that of the landmark based face morphing. The lower geometrical correspondence despite the high visual quality fails in verification stage from FRS.
- When compared to ArcFace, the COTS indicates a higher vulnerability for both landmarks and StyleGAN based morph attack detection due to the high accuracy of ver-

ifying the subjects in COTS under different data capture conditions as expected in operational scenario. Thus, COTS while making itself robust about certain degree of degraded data, also accepts the morphs to a higher degree.

- Table I also indicates the distinction between FMMPMR and MMPMR metric used to quantify the vulnerability. The MMPMR reports high values in comparison to FMMPMR as it does not account for the number of attempts per morphing image. Further, we have also measured the Relative Morph Match Rate (RMMR) [16] that can account for the True Acceptance Rate of the FRS. Since both FRS employed in this paper have reached TAR = 100%, the RMMR is the same as the FMMPMR/MMPMR.

TABLE II: Quantitative performance of state-of-the-art MAD techniques on StyleGAN dataset

Morphing Type	Algorithms	D-EER(%)	BPCER(%) @ APCER	
			=5(%)	=10(%)
Landmark based Morph [6]	HoG-SVM	10.29	17.66	10
	LBP-SVM	15.42	29.15	22.98
	Color Textures	1.57	0.51	0.17
	CAN	4.8	4.63	2.4
MorGAN [3]	HoG-SVM	0	0	0
	LBP-SVM	0	0	0
	Color Textures	0	0	0
	CAN	0	0	0
StyleGAN	HoG-SVM	0.04	0	0
	LBP-SVM	0.68	0	0
	Color Textures	0	0	0
	CAN	0.36	0	0

D. Performance Metrics for MAD

The performance of Morphing Attack Detection (MAD) techniques are presented using the ISO/IEC 30107-3 metrics [7] such as Attack Presentation Classification Error Rate (APCER (%)) which defines the proportion of attack images incorrectly classified as bona fide images and Bona fide Presentation Classification Error Rate (BPCER (%)) in which bona fide images incorrectly classified as attack images [7] along with the Detection Equal Error Rate (D-EER (%)).

E. MAD Detection Performance

In this section, we report the detection performance of MAD techniques to understand the impact of different types of morphing techniques. We have therefore selected four different MAD techniques - LBP-SVM [11], HoG-SVM [18], color denoising [17], Context aggregation Network (CAN) [18] based on the recent benchmarks. Table II indicates the MAD performance on all three different morph generation techniques.

Compared to three different morph generation methods, landmark based technique indicates a relatively high challenge

for the detection techniques, when compared to that of GAN based techniques. However on the same kind of morph generation approach, the recent technique based on color texture indicates the lowest error rates with D-EER(%) of 1.57(%). While it is noted that the GAN generated morphs are easier to detect, a possible reason can be attributed to the residual noise [9] that is associated with GAN in generating these morphed images. Even though StyleGAN can generate a high quality images with a resolution of 1024×1024 pixels, the inherent noise in the generated morph images make enables to detect them. This is not the case for landmark based morph images, which do not contain such characteristic noise.

F. Limitations and Future Directions

Observing the results from the empirical evaluation of different approaches of morph generation both for threats to FRS and ability to detect the morphs, we note certain limitations in the current work as listed below.

- The GAN based morph generation does not impose the landmark correspondence leading to high quality images but not with high facial similarity in geometrical appearance to contributing subjects. This has lead to lower threat to FRS in comparison to landmark based morphs. Future works in this direction can focus on imposing such a constraint in the latent space, in order to increase the threat to FRS.
- Despite the accuracy of MAD being very high, it can be primarily attributed to digital pixel level information helping to detect the attacks. A print and scan of the the same morphed images can further reveal the real challenge in detecting the morphing attacks as the print-scan cycle loses the pixel level soft-information in the image.

The future works in this direction will lead to establishing the real threat landscape on FRS from the GAN generated morphed face images.

IV. CONCLUSION

This work investigated the feasibility of generating high quality morph generation and proposed a new approach using StyleGAN. The proposed approach resulted in morphed face images with a dimension of 1024×1024 pixels and no visual artifacts. To indicate the real threat potential to FRS, the morphed face images generated from proposed StyleGAN were analyzed using a commercial FRS and an open-source FRS. Further, to provide a fair comparison to earlier works, MorGAN and Landmark based approaches were benchmarked on the same set of data by creating a new morphed face database. The set of experiments clearly indicate the that StyleGAN based morphed face images do show threats to FRS but to a much lower degree as compared to traditional landmark based morph generation techniques. While detecting the attacks stemming from GAN approaches is relatively easy in the digital domain, the real challenge of detecting them after the print-scan process is still not explored. In summary, we answer the question - *Can GAN Generated Morphs Threaten*

Face Recognition Equally as Landmark Based Morphs?, our experimental results indicates with a clear no in digital domain alone.

REFERENCES

- [1] R. Abdal, Y. Qin, and P. Wonka. Image2stylegan: How to embed images into the stylegan latent space? *CoRR*, abs/1904.03189, 2019.
- [2] N. Damer, F. Boutros, A. Moseguí Saladié, F. Kirchbuchner, and A. Kuijper. Realistic dreams: Cascaded enhancement of gan-generated images with an example in face morphing attacks. 10 2019.
- [3] 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 Biometrics Theory, Applications and Systems (BTAS)*, pages 1–10, Oct 2018.
- [4] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. In *IEEE International Joint Conference on Biometrics*, pages 1–7. IEEE, 2014.
- [5] M. Ferrara, A. Franco, and D. Maltoni. *Face Recognition Across the Imaging Spectrum*, chapter On the Effects of Image Alterations on Face Recognition Accuracy, pages 195–222. Springer International Publishing, 2016.
- [6] M. Ferrara, A. Franco, and D. Maltoni. Decoupling texture blending and shape warping in face morphing. In *2019 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5. IEEE, 2019.
- [7] ISO/IEC JTC1 SC37 Biometrics. *ISO/IEC 30107-3. Information Technology - Biometric presentation attack detection - Part 3: Testing and Reporting*. International Organization for Standardization, 2017.
- [8] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019.
- [9] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and improving the image quality of stylegan. *arXiv preprint arXiv:1912.04958*, 2019.
- [10] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, Jin Chang, K. Hoffman, J. Marques, Jaesik Min, and W. Worek. Overview of the face recognition grand challenge. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, pages 947–954 vol. 1, June 2005.
- [11] R. Raghavendra, K. B. Raja, and C. Busch. Detecting Morphed Face Images. In *8th IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, pages 1–8, 2016.
- [12] R. Raghavendra, K. B. Raja, S. Venkatesh, and C. Busch. Face morphing versus face averaging: Vulnerability and detection. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 555–563, 2017.
- [13] R. Raghavendra, K. B. Raja, S. Venkatesh, and C. Busch. Transferable deep-cnn features for detecting digital and print-scanned morphed face images. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1822–1830. IEEE, 2017.
- [14] R. Ramachandra and C. Busch. Presentation attack detection methods for face recognition systems: A comprehensive survey. *ACM Computing Surveys (CSUR)*, 50(1):1–37, 2017.
- [15] D. Robertson, R. S. Kramer, and A. M. Burton. Fraudulent id using face morphs: Experiments on human and automatic recognition. *PLoS ONE*, 12(3):1–12, 2017.
- [16] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. Veldhuis, L. Spreeuwens, M. Schils, D. Maltoni, P. Grother, S. Marcel, B. Ralph, R. Raghavendra, and C. Busch. Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting. In *2017 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–7. IEEE, 2017.
- [17] S. Venkatesh, R. Ramachandra, K. Raja, L. Spreeuwens, R. Veldhuis, and C. Busch. Morphed face detection based on deep color residual noise. In *ninth International Conference on Image Processing Theory, Tools and Applications (IPTA 2019)*, pages 1–5. IEEE, 2019.
- [18] S. Venkatesh, R. Ramachandra, K. Raja, L. Spreeuwens, R. Veldhuis, and C. Busch. Detecting morphed face attacks using residual noise from deep multi-scale context aggregation network. In *2020 Winter Conference on Applications of Computer Vision (WACV 2020)*. IEEE, 2020.