

# MorDIFF: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Diffusion Autoencoders

Naser Damer<sup>1,2</sup>, Meiling Fang<sup>1,2</sup>, Patrick Siebke<sup>1</sup>, Jan Niklas Kolf<sup>1,2</sup>  
Marco Huber<sup>1,2</sup>, Fadi Boutros<sup>1</sup>

<sup>1</sup>Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany

<sup>2</sup>Department of Computer Science, TU Darmstadt, Darmstadt, Germany

**Abstract**—Investigating new methods of creating face morphing attacks is essential to foresee novel attacks and help mitigate them. Creating morphing attacks is commonly either performed on the image-level or on the representation-level. The representation-level morphing has been performed so far based on generative adversarial networks (GAN) where the encoded images are interpolated in the latent space to produce a morphed image based on the interpolated vector. Such a process was constrained by the limited reconstruction fidelity of GAN architectures. Recent advances in the diffusion autoencoder models have overcome the GAN limitations, leading to high reconstruction fidelity. This theoretically makes them a perfect candidate to perform representation-level face morphing. This work investigates using diffusion autoencoders to create face morphing attacks by comparing them to a wide range of image-level and representation-level morphs. Our vulnerability analyses on four state-of-the-art face recognition models have shown that such models are highly vulnerable to the created attacks, the MorDIFF, especially when compared to existing representation-level morphs. Detailed detectability analyses are also performed on the MorDIFF, showing that they are as challenging to detect as other morphing attacks created on the image- or representation-level. Data and morphing script are made public <sup>1</sup>.

**Index Terms**—Face Recognition, Morphing Attack, Diffusion

## I. INTRODUCTION

Face recognition (FR) systems, despite their high accuracy, are vulnerable to many attacks, one of which is face morphing attack. Face morphing attacks aim at creating face images that are verifiable to be the face of multiple identities, which can lead to building faulty identity links in operations like border checks. This was shown by Ferrara et al. [15] when they proved that one attack image can successfully match more than one person, automatically and by human experts. When such morphing attacks are used in association with travel or identity documents, it can allow multiple subjects to verify their identity to the alphanumeric information of the document. This possible false link to the identity information might enable a number of illegal actions related to financial transactions, human trafficking, illegal immigration, among others.

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science, and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

<sup>1</sup><https://github.com/naserdamer/MorDIFF>

Face morphing can be performed on the image-level, commonly by interpolating facial landmarks in the morphed images and blending the texture information [16], [32]. Morphing can also be performed on the representation-level by interpolating face image representations and decoding this interpolated representation into a face morphing attack. This representation-level morphing has been so far performed using GAN architectures, with latent (representation) size depending on the number of GAN training samples, leading to low reconstruction fidelity. Typically, despite being prone to blending artifacts, the image-level morphing led to attacks that much more strongly preserve the identities of the morphed images, in comparison to GAN-based representation-level morphs. With recent advancements in the diffusion autoencoders [20], [30], [37] have aimed at avoiding the pitfalls of GAN architectures to significantly enhance the reconstruction fidelity. This makes diffusion autoencoders a perfect candidate to perform representation-level morphing with exceptional identity-preservation and realistic image appearance.

This work investigates the vulnerability of state-of-the-art FR models to morphing attacks created on the representation-level by diffusion autoencoders. To perform this, we created our diffusion-based morphing attacks, the MorDIFF, by interpolating the semantics and stochastic latent representation of the two face images to be morphed. This interpolated latent is later decoded into the morphing attack image. We investigated the vulnerability of a set of FR models to MorDIFF attacks, proving their high ability to attack FR systems in comparison to other morphing techniques. We additionally investigated the detectability of MorDIFF attacks by a set of morphing attack detection (MAD) solutions based on different backbones and training data, pointing out the challenging nature of detecting MorDIFF attacks, among other attacks.

## II. RELATED WORKS

Face morphing is either performed on the image-level or on the representation level. First face morphing attacks were created on the image-level by detecting facial landmarks in the source images to be morphed. The facial landmarks are interpolated and the texture is blended, resulting in what is commonly referred to as landmark-based morphs (LMA). Variations of this process were used in the literature, such as the work of Ferrara et al. [16] and Ramachandra et al.

[32]. An early comparison [34] of such methods has shown that the approach used in [6], [32] achieved the relatively strong face morphing attacks, i.e. high identity preservation of the morphed identities. Further variations of this process performed the interpolation individually on partial facial parts [31], producing attacks that are harder to detect by the MAD. The listed image-level morphs, and others that are used in the experiment of this work, have various degrees of image artifacts introduced by the fact that the identity interpolation is performed on the image-level [41].

Given the advances in GAN architectures and their performance in producing synthetic images, and to avoid the disadvantages of image-level interpolation, MorGAN GAN-based morphing approach was proposed in [11]. The first representation-level face morphing attacks were MorGAN, where the solution encoded the images to be morphed into the GAN latent space and interpolated these latent representations. The interpolated latent vector is then decoded by the GAN generator to produce the morphed face image. These early representation-level morphs preserved the identities moderately and were of low resolution, but were proven to be hard to detect if they were not used to train the MAD [8], [13]. A post-generation cascaded enhancement step was added in a follow-up work on the MorGAN generator to increase the image perceptual quality, nonetheless, with the same identity preservation level [7]. Based on the concept introduced by MorGAN [11], Venkatesh et al. created morphed images with better identity preservation qualities and a much more realistic appearance [40]. This was mainly driven by the use of the StyleGAN architecture by Karras et al. [26]. Later on, also based on the StyleGAN architecture [26], the MIPGAN I and II were introduced to generate images with higher identity preservation [41]. This was enforced by a training loss function that optimizes the identity interpolation on the latent-level. A hybrid family of morphing approaches did use the GAN architecture, but as a post-processing step after performing the morphing on the image-level. ReGenMorphs [10] introduced this concept to take advantage of the identity-preservation qualities of the image-level morphing while encoding and decoding the morphed image in a GAN architecture to produce a new face (without blending artifacts) that posses the morphed identity information.

The existing image-level morphing techniques produce the desired identity-preservation properties of a strong attack, however, are linked to visible blending artifacts. The existing representation-level morphing techniques are all GAN-based. Such a method preserves the desired identity to a much lower degree than image-based morphs and is linked to synthetic-like generation artifacts when the manipulations are performed on the latent space [11], [40], [41]. This is caused by the relatively low reconstruction fidelity of GAN architectures, which are themselves dependent on the size of training data [30]. This is what diffusion models targeted by achieving higher fidelity and thus act as a perfect candidate to facilitate more realistic and more identity-preserving representation-level morphing attacks, as investigated in this work. A recent not-peer-reviewed

pre-print was posted a few days before the submission of this work where such a concept was investigated [1], however with vulnerabilities being measured with extremely outdated and outperformed FR solutions, making the presented results practically irrelevant. Regardless of the morphing techniques, all possible morphing attacks should be considered and novel attacks should be foreseen to help build defensive mechanisms before such attacks are used by actual attackers, especially for highly realistic and strong attacks.

### III. DIFFUSION-BASED FACE MORPHING

This section describes the morphing technique used to create the MorDIFF morphs. The approach is based on the work of Preechakul et al. [30] and is described in the following with the necessary background information.

Previous morphing generation process used the GAN architecture [7], [11], [40] and generated good morph attacks. Despite the fact that GAN can learn meaningful representations in the latent spaces, the reconstruction fidelity is limited due to the latent size depending on the number of training samples. In contrast to GAN, diffusion probabilistic model (DPM) [20], [37] has achieved higher visual fidelity in image generation. However, the latent code of DPM lacks semantic meaning, which makes it harder for diverse representation learning in comparison with GAN-based models. To address such issues, diffusion autoencoder [30] was proposed to separately infer both semantic and stochastic information from the inputs. Therefore, we use the diffusion autoencoder [30] to generate high fidelity and smooth face morphing between two input face images  $x_I^1$  and  $x_I^2$ .

The diffusion autoencoder contains two encoders. One is a semantic encoder mapping  $x_I^1$  and  $x_I^2$  to semantic latent representations  $\mathbf{z}_s^1$  and  $\mathbf{z}_s^2$ . The other one is a stochastic encoder to obtain the latent code  $\mathbf{x}_t^1$  and  $\mathbf{x}_t^2$  from inputs. Therefore, two subcodes are encoded from  $x_I^1$  and  $x_I^2$  to  $(\mathbf{z}_s^1, \mathbf{x}_t^1)$  and  $(\mathbf{z}_s^2, \mathbf{x}_t^2)$ . Then, a linear interpolation *Lerp* and a spherical linear interpolation *SLerp* (as recommended in [36] and [38]) are used to obtain a representative latent code  $\mathbf{z}(\lambda)$  with  $\lambda \in [0, 1]$  from the obtained subcodes. The interpolation can be formulated as following:

$$Lerp(\mathbf{z}_s^1, \mathbf{z}_s^2; \lambda) = \lambda \mathbf{z}_s^1 + (1 - \lambda) \mathbf{z}_s^2 \quad (1)$$

$$SLerp(\mathbf{x}_t^1, \mathbf{x}_t^2; \lambda) = \frac{\sin \theta \lambda \mathbf{x}_t^1}{\sin \theta} + \frac{\sin \theta (1 - \lambda) \mathbf{x}_t^2}{\sin \theta} + \quad (2)$$

where

$$\theta = \frac{\arccos(\mathbf{x}_t^1 \cdot \mathbf{x}_t^2)}{\|\mathbf{x}_t^1\| \|\mathbf{x}_t^2\|} \quad (3)$$

Therefore, the latent code  $\mathbf{z}(\lambda)$  can be formulated as  $\mathbf{z}(\lambda) = (Lerp(\mathbf{z}_s^1, \mathbf{z}_s^2; \lambda), SLerp(\mathbf{x}_t^1, \mathbf{x}_t^2; \lambda))$ . Finally, a conditional DPM is used as decoder to produce the morphing attack, denoted as MorDIFF in our case. The used diffusion autoencoder training details are provided in [30] and we used the pre-trained autoencoder available publicly <sup>2</sup>.

<sup>2</sup><https://github.com/phizaz/diffae>



Fig. 1: Samples of the MorDIFF (right most column) attacks, the baseline attacks (created by FaceMorpher, OpenCV, WebMorph, MIPGAN I and II), and the bona fide images that were morphed to create the attack (two left most columns). The image-level morphs (FaceMorpher, OpenCV, WebMorph) show the traditional blending artifacts, while the representation-level morphs (MIPGAN-I and II) show typical streaking GAN artifacts. The MorDIFF attacks in these samples show much lower level of generative artifacts, as expected when comparing diffusion-based generation to GAN-based generation [30].

In the whole morphing process, semantic encoder is able to manipulate the facial features for partial identity information preserving and stochastic encoder provides the detailed information to produce high visual fidelity morphing attack samples, as shown in Figure 1.

#### IV. THE MORDIFF DATASET

The MorDIFF dataset extends over the SYN-MAD 2022 competition dataset [22] and uses the same morphing pairs to enable a comparable dataset set. The MorDIFF and the SYN-MAD 2022 datasets are both based on the Face Research Lab London (FRL) dataset [14]. The FRL dataset contains images of 102 different individuals and provides high-quality frontal images created in a controlled scenario with a wide range of different ethnicities. All individuals present in the dataset signed consent for their images to be “used in lab-based and web-based studies in their original or altered forms and to illustrate research.” For the morph generation, we limited the data to the frontal images of the dataset, similar to SYN-MAD 2022.

The pairs are selected by splitting the frontal images of the FRL dataset depending on the provided gender and ex-

pression (neutral or smiling). ElasticFace-Arc [3] is then used to generate embeddings of the images and these embeddings are then compared with cosine similarity within the split to find the most similar faces. The 250 most similar face pairs are selected, resulting in 250 female neutral pairs, 250 female smiling pairs, 250 male neutral pairs, and 250 male smiling pairs, in total 1000 MorDIFF attack images and the 204 bona fide images of the SYN-MAD 2022 dataset.

The SYN-MAD 2022 dataset, which we use as a benchmark, contains morphed images from 5 different approaches, three image-level (FaceMorpher (commercial-of-the-shelf), OpenCV [28] and Webmorph (online tool <sup>3</sup>)) and two representation-level GAN-based (MIPGAN I [41] and MIPGAN II [41]). We follow the same morph pair selection protocol (described above) and now add the MorDIFF attacks to the benchmark, which will be publicly released. Other GAN-based representation-level morphing techniques, such as the MorGAN [11], EMorGAN [7], and StyleGAN morphs [40] were not considered, as their quality and identity preservation

<sup>3</sup><https://webmorph.org/>

ability is far inferior MIPGAN I and II as shown in [41]).

Samples of the MorDIFF attacks are shown in Figure 1 along with the morphed bonafide images and the morphing attacks created by the five baseline methods. Looking carefully at the images, one notices that the image-level morphs (OpenCV, FaceMorpher, and WebMorph) do, as expected have higher shadowing and blending artifacts in comparison to the representation-level morphs (MIPGAN I and II, and the MorDIFF). The MIPGAN I and II contain streaking artifacts common to the GAN-generated faces. This kind of artifact is almost non-existent in the samples of the MorDIFF attacks.

We considered measuring the statistical and perceptual image quality of the different morphs and found that previous works have shown no clear correlation between the image quality and the realistic appearance when dealing with Morphing attacks (typically of ICAO standard [23] and of very high perceived quality) [41]. Such investigation have proven to lead to faulty conclusions, and thus we opted out of such study. In [41], the PSNR and SSIM values showed an insignificant difference between the different attacks, and a minimum difference was inconsistent with the perceived visual quality. Other work showed that the visibly unrealistic images of MorGAN have higher statistical quality metrics (6 different metrics [7]) than other attacks. In a recent study, researchers have shown that operations that apparent morphing artifacts do not consistently affect the estimated quality across a large number of quality estimation strategies [17], [18], however, especially utility metrics, can be used to differentiate between attacks and bona fide to a small degree [17]. However, studies on the human observer’s ability to detect the MorDIFF attack are necessary and are planned to evaluate the relative perceived quality of these attacks, as conducted on previous attack methods in [19].

## V. VULNERABILITY OF FACE RECOGNITION TO MORDIFF ATTACKS

We evaluated the vulnerability of four FR systems to the MorDIFF attacks in comparison to the five baseline attacks. The chosen FR systems are two of the top-performing large-scale FR solution, i.e. ElasticFace (ElasticFace-Arc) [3] and CurricularFace [21], and two top performing light-weight FR solutions that can be implemented in embedded and portable devices, i.e. the MixFaceNet [2] and PocketNet [4]. CurricularFace and ElasticFace are based on ResNet-100 architectures and have 55.52M parameters each with 24192.51 MFLOPs, MixFaceNet consists of 3.07M parameters with 451.7 MFLOPs and PocketNet has 1.75M parameters and 1099.15 MFLOPs. On the frequently used benchmark Labeled Faces in the Wild [42], CurricularFace and ElasticFace achieve both an accuracy of 99.80%, MixFaceNet scores 99.60% and PocketNet 99.58%. On the AgeDB-30 benchmark [29], the accuracy is respectively 98.32% for CurricularFace, 98.35% for ElasticFace, 96.63% for MixFaceNet and 97.17% for PocketNet. All the FR models used are the official releases by the respective authors.

We present the vulnerability results in a quantifiable manner by listing the Mated Morphed Presentation Match Rate

(MMPMR) [35] and the Fully Mated Morphed Presentation Match Rate (FMMPMR) [35] based on the decision threshold at the false match rate of (FMR) 1% and 0.1%, MMPMR100 and MMPMR1000, respectively.

Table I presents the results of our vulnerability analyses. On all four considered FR solutions, and for both operating points (FMR of 1% and 0.1%), the FR is more vulnerable to MorDIFF attacks in comparison to other representation-level attacks. This is evident by the large difference between the MMPMR values (MMPMR100 and MMPMR1000) scored by the MorDIFF attacks in comparison to those scored by the MIPGAN I and II attacks. As expected, FR solutions are more vulnerable to the blending artifact-prone image-level morphing techniques in comparison to the MIPGAN I and II attacks, especially the image-level morphs created by OpenCV and WebMorph. This gap between the FR vulnerability to image-level and representation-levels morphs became very tight when comparing the image-level morphs to MorDIFF attacks. MorDIFF attacks score MMPMR values that are in most cases extremely close to the best-performing image-level morphs. One can also notice that for all attack types, and especially for lower FMR (0.1% ,i.e. MMPMR1000) the vulnerability of the slightly higher performing large-scale FR models (ElasticFace and CurricularFace) is higher than the vulnerability of the compact light-weight models (MixFaceNet and PocketNet).

## VI. DETECTABILITY OF MORDIFF ATTACKS

To measure the detectability of the MorDIFF attacks, along with the baseline attacks, we deploy two well-performing MADs: MixFaceNet-MAD and Inception-MAD.

MixFaceNet-MAD employs the backbone of MixFaceNet [2] for face verification and identification. Considering its lower computation complexity (FLOPs) and high accuracy, MixFaceNet was used successfully for face MAD [9] by adding a classification layer after embedding layer to make the MAD decision. We followed the same modification and settings in [9] to conduct experiments. Inception-MAD uses the Inception-v3 [39] network architecture as a backbone. This architecture has been used successfully for MAD [12], [33]. Both models are trained from scratch by using the cross-entropy loss function and Adam optimizer with a learning rate of  $10^{-4}$  and a weight decay of  $10^{-5}$ , where the input size for Inception-MAD is  $229 \times 229 \times 3$  and for MixFaceNet-MAD is  $224 \times 224 \times 3$ . To avoid overfitting and for a fair comparison, an early stopping technique with the patience of 20 epochs.

Three instances of each of the MADs are evaluated, and each of these instances is trained on a different database. This results in a realistic evaluation protocol where all the MADs are trained on datasets of different origins than the evaluation ones (different source datasets and different identities). The three training datasets used in experiments are SMDD, LMA-DRD, and MorGAN-LMA.

SMDD [9] dataset is a synthetic-based MAD dataset that 25,000 bona fide images are created by StyleGAN2-ADA [25], [40], and 15,000 morph attacks are generated based on such bona fide by using OpenCV morphing technique [28].

FR model		ElasticFace [3]		CurricularFace [21]		MixFaceNet [2]		PocketNet [4]	
Morphing technique		MMPMR100	MMPMR1000	MMPMR100	MMPMR1000	MMPMR100	MMPMR1000	MMPMR100	MMPMR1000
Image level	OpenCV	0.997	0.980	0.996	0.986	0.996	0.963	0.996	0.970
	FaceMorpher	0.962	0.913	0.970	0.935	0.972	0.931	0.979	0.941
	WebMorph	0.990	0.986	0.988	0.988	0.988	0.984	0.988	0.988
Representation level	MIPGAN-I	0.980	0.845	0.962	0.890	0.981	0.887	0.991	0.900
	MIPGAN-II	0.953	0.778	0.953	0.832	0.973	0.836	0.977	0.857
	MorDIFF (our)	<b>0.990</b>	<b>0.948</b>	<b>0.995</b>	<b>0.968</b>	<b>0.992</b>	<b>0.958</b>	<b>0.996</b>	<b>0.949</b>

TABLE I: Vulnerability of four face recognition solutions to the MorDIFF and the baseline attacks. Under all evaluation settings, the MorDIFF attacks were stronger (FR more vulnerable) than the other representation-level morphing attacks (indicated by the highest MMPMR values in this category in bold). The vulnerability to the MorDIFF attacks is, in many evaluation settings, very close to the strongest image-level morphing attack (highest MMPMR values indicated by italic font in that category).

Considering the privacy-friendly characteristic and applicability of SMDD for generalized MADs, we use SMDD as one of the training sets. In addition to synthetic MAD data, we utilized two real MAD datasets, LMA-DRD and MorGAN-LMA. LMA-DRD includes 123 bona fide samples selected from VGGFace2 dataset [5] and 88 morphed attacks generated by OpenCV morphing [28]. These attacks were then created by the re-digitized (print and scan) method. MorGAN-LMA contains 750 bona fides selected from CelebA [27] and 500 morphs created by OpenCV landmark-based morphing [28].

The MAD performance (detectability) is presented by the Attack Presentation Classification Error Rate (APCER), i.e. the proportion of attack images incorrectly classified as bona fide samples, at a fixed (1%, 10%, and 20%) Bona fide Presentation Classification Error Rate (BPCER), i.e. the proportion of bona fide images incorrectly classified as attack samples, as defined in the ISO/IEC 30107-3 [24]. Additionally, the Detection Equal Error Rate (EER), i.e. the value of APCER or BPCER at the decision threshold where they are equal, is reported.

The detectability results are presented in Table II. Keeping in mind that all the MAD evaluation scenarios presented here are realistic cross-dataset evaluations, all the morphing attacks were not perfectly detected. MADs trained on the large-scale synthetic-based SMDD tend to perform better than other MADs, which might be due to the higher diversity in the training data. For most MADs, the MorDIFF was not the most, nor the least accurately detected morph type. Thus they can be considered to be challenging, similarly to other attack types, to be detected by the MADs. In general, the MixFaceNet-MAD performed slightly better than the Inception-MAD in most cases, whether to detect MorDIFF or other attacks, keeping in mind that this MAD (trained on SMDD) was the baseline for the recent SYN-MAD 2022 competition [22]. MorDiff attacks seem to be also detected relatively better when the MADs are trained on the SMDD dataset.

## VII. CONCLUSION

This paper investigated using diffusion autoencoders to generate morphing attacks by interpolating the semantics and stochastic latent representation of the two face images to be morphed, which are later decoded to create a morphed image, i.e. the MorDIFF. We showed that state-of-the-art FR solutions are extremely vulnerable to this representation-level morphing technique, especially when compared to the recent GAN-based representation-level morphing solutions such as MIPGAN I and II. The MorDIFF attacks are also challenging to detect,

MAD	Train data	Test data	EER (%)	APCER (%) @ BPCER (%)		
				1.00	10.00	20.00
MixFaceNet-MAD	SMDD	FaceMorph	4.60	5.50	3.60	2.90
		MIPGAN_I	16.70	75.80	22.20	14.50
		MIPGAN_II	20.62	81.58	32.03	20.62
		OpenCV	8.33	36.38	6.50	3.76
		WebMorph	18.20	74.00	24.00	17.60
		MorphDiffusion	8.50	33.40	7.40	4.10
	LMAD-DRD	FaceMorph	5.60	11.20	3.30	1.20
		MIPGAN_I	14.40	64.90	19.80	7.70
		MIPGAN_II	11.51	48.65	13.51	4.30
		OpenCV	16.37	72.46	25.61	12.09
		WebMorph	21.60	82.80	46.60	23.80
		MorphDiffusion	21.40	81.60	43.90	22.30
	MorGAN-LMA	FaceMorph	8.00	13.90	7.30	0.00
		MIPGAN_I	14.60	93.00	30.40	0.00
		MIPGAN_II	18.92	93.19	34.33	0.00
OpenCV		9.86	78.76	13.41	0.00	
WebMorph		15.80	92.40	34.00	0.00	
MorphDiffusion		13.50	89.30	27.20	0.00	
Inception-MAD	SMDD	FaceMorph	0.00	1.70	0.00	0.00
		MIPGAN_I	10.90	50.90	13.70	5.70
		MIPGAN_II	16.22	82.48	25.83	11.41
		OpenCV	7.52	28.66	5.49	3.05
		WebMorph	18.00	85.20	27.40	13.40
		MorphDiffusion	5.30	17.20	3.50	2.50
	LMAD-DRD	FaceMorph	61.00	99.90	99.60	97.40
		MIPGAN_I	41.30	99.60	89.50	70.40
		MIPGAN_II	39.74	99.60	88.19	63.37
		OpenCV	8.84	40.96	8.23	2.13
		WebMorph	20.20	85.60	41.00	17.00
		MorphDiffusion	95.40	100.00	100.00	100.00
	MorGAN-LMA	FaceMorph	0.80	0.70	0.00	0.00
		MIPGAN_I	46.10	98.70	89.20	78.00
		MIPGAN_II	35.84	96.80	74.28	56.26
OpenCV		9.34	45.43	9.25	2.95	
WebMorph		19.00	70.00	31.80	18.40	
MorphDiffusion		26.50	87.00	53.10	32.50	

TABLE II: The detectability of the MorDIFF attacks and the baseline attacks using six different MADs (two architectures and 3 training datasets). The MorDIFF attacks are challenging to detect, as are other attack types, especially in the realistic scenario of cross-dataset evaluation we present here. Please note that the rows with EER values above 50% (which is not realistic) indicate that the detection scores of the MAD, in such cases, indicated that the bona fide are closer to being predicted as attacks than the attacks themselves.

similarly to other morphing attacks. The MorDIFF data will be publicly released and it follows the same generation protocol of the recent MAD evaluation benchmark, SYN-MAD 2022, to maintain research reproducibility and comparison. Future research will look into the generalization benefits of including the MorDIFF morphs in the training of MAD solutions.

## REFERENCES

- [1] Zander Blasingame and Chen Liu. Diffusion models for stronger face morphing attacks. *CoRR*, abs/2301.04218, 2023. 2
- [2] Fadi Boutros, Naser Damer, Meiling Fang, Florian Kirchbuchner, and Arjan Kuijper. Mixfacenets: Extremely efficient face recognition networks. In *IJCB*, pages 1–8. IEEE, 2021. 4, 5

- [3] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *CVPR Workshops*, pages 1577–1586. IEEE, 2022. 3, 4, 5
- [4] Fadi Boutros, Patrick Siebke, Marcel Klemt, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Pocketnet: Extreme lightweight face recognition network using neural architecture search and multistep knowledge distillation. *IEEE Access*, 10:46823–46833, 2022. 4, 5
- [5] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *FG*, pages 67–74. IEEE Computer Society, 2018. 5
- [6] Naser Damer, Viola Boller, Yaza Wainakh, Fadi Boutros, Philipp Terhörst, Andreas Braun, and Arjan Kuijper. Detecting face morphing attacks by analyzing the directed distances of facial landmarks shifts. In *GCPR*, volume 11269, pages 518–534. Springer, 2018. 2
- [7] Naser Damer, Fadi Boutros, Alexandra Mosegui Saladie, Florian Kirchbuchner, and Arjan Kuijper. Realistic dreams: Cascaded enhancement of gan-generated images with an example in face morphing attacks. In *BTAS*, pages 1–10. IEEE, 2019. 2, 3, 4
- [8] Naser Damer, Jonas Henry Grebe, Steffen Zienert, Florian Kirchbuchner, and Arjan Kuijper. On the generalization of detecting face morphing attacks as anomalies: Novelty vs. outlier detection. In *BTAS*, pages 1–5. IEEE, 2019. 2
- [9] Naser Damer, César Augusto Fontanillo López, Meiling Fang, Noémie Spiller, Minh Vu Pham, and Fadi Boutros. Privacy-friendly synthetic data for the development of face morphing attack detectors. In *CVPR Workshops*, pages 1605–1616. IEEE, 2022. 4
- [10] Naser Damer, Kiran B. Raja, Marius Süßmilch, Sushma Venkatesh, Fadi Boutros, Meiling Fang, Florian Kirchbuchner, Raghavendra Ramachandra, and Arjan Kuijper. Regenmorph: Visibly realistic GAN generated face morphing attacks by attack re-generation. In *ISVC (I)*, volume 13017, pages 251–264. Springer, 2021. 2
- [11] Naser Damer, Alexandra Mosegui Saladie, Andreas Braun, and Arjan Kuijper. Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network. In *BTAS*, pages 1–10. IEEE, 2018. 2, 3
- [12] Naser Damer, Noémie Spiller, Meiling Fang, Fadi Boutros, Florian Kirchbuchner, and Arjan Kuijper. PW-MAD: pixel-wise supervision for generalized face morphing attack detection. In *ISVC (I)*, volume 13017, pages 291–304. Springer, 2021. 4
- [13] Naser Damer, Steffen Zienert, Yaza Wainakh, Alexandra Mosegui Saladie, Florian Kirchbuchner, and Arjan Kuijper. A multi-detector solution towards an accurate and generalized detection of face morphing attacks. In *FUSION*, pages 1–8. IEEE, 2019. 2
- [14] Lisa DeBruine and Benedict Jones. Face Research Lab London Set. 4 2021. 3
- [15] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. The magic passport. In *IJCB*, pages 1–7. IEEE, 2014. 1
- [16] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. Face demorphing. *IEEE Trans. Inf. Forensics Secur.*, 13(4):1008–1017, 2018. 1
- [17] Biying Fu and Naser Damer. Face morphing attacks and face image quality: The effect of morphing and the unsupervised attack detection by quality. *IET Biometrics*, 11(5):359–382, 2022. 4
- [18] Biying Fu, Noémie Spiller, Cong Chen, and Naser Damer. The effect of face morphing on face image quality. In *BIOSIG*, LNI. Gesellschaft für Informatik e.V., 2021. 4
- [19] Sankini Rancha Godage, Froy Lovåsaly, Sushma Venkatesh, Kiran Raja, Raghavendra Ramachandra, and Christoph Busch. Analyzing human observer ability in morphing attack detection -where do we stand? *IEEE Trans. on Technology and Society*, pages 1–1, 2022. 4
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1, 2
- [21] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: Adaptive curriculum learning loss for deep face recognition. In *CVPR*, pages 5900–5909. Computer Vision Foundation / IEEE, 2020. 4, 5
- [22] Marco Huber, Fadi Boutros, Anh Thi Luu, Kiran B. Raja, Raghavendra Ramachandra, Naser Damer, Pedro C. Neto, Tiago Gonçalves, Ana F. Sequeira, Jaime S. Cardoso, João Tremeço, Miguel Lourenço, Sergio Serra, Eduardo Cermeño, Marija Ivanovska, Borut Batagelj, Andrej Kronovsek, Peter Peer, and Vitomir Struc. SYN-MAD 2022: Competition on face morphing attack detection based on privacy-aware synthetic training data. In *IJCB*, pages 1–10. IEEE, 2022. 3, 5
- [23] International Civil Aviation Organization, ICAO. Machine readable passports – part 9 – deployment of biometric identification and electronic storage of data in eMRTDs. *Civil Aviation Organization (ICAO)*. 4
- [24] International Organization for Standardization. ISO/IEC DIS 30107-3:2016: Information Technology – Biometric presentation attack detection – P. 3: Testing and reporting, 2017. 5
- [25] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, 2020. 4
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410. Computer Vision Foundation / IEEE, 2019. 2
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738. IEEE Computer Society, 2015. 5
- [28] Satya Mallick. Face morph using opencv — c++ / python. *LearnOpenCV*, 1(1), 2016. 3, 4, 5
- [29] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: The first manually collected, in-the-wild age database. In *CVPR Workshops*, pages 1997–2005. IEEE Computer Society, 2017. 4
- [30] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *CVPR*, pages 10609–10619. IEEE, 2022. 1, 2, 3
- [31] Le Qin, Fei Peng, Sushma Venkatesh, Raghavendra Ramachandra, Min Long, and Christoph Busch. Low visual distortion and robust morphing attacks based on partial face image manipulation. *IEEE Trans. Biom. Behav. Identity Sci.*, 3(1):72–88, 2021. 2
- [32] Ramachandra Raghavendra, Kiran B. Raja, Sushma Venkatesh, and Christoph Busch. Face morphing versus face averaging: Vulnerability and detection. In *IJCB*, pages 555–563. IEEE, 2017. 1, 2
- [33] Raghavendra Ramachandra, Sushma Venkatesh, Kiran B. Raja, and Christoph Busch. Detecting face morphing attacks with collaborative representation of steerable features. In *CVIP (I)*, volume 1022 of *AISC*, pages 255–265. Springer, 2018. 4
- [34] Ulrich Scherhag, Jonas Kunze, Christian Rathgeb, and Christoph Busch. Face morph detection for unknown morphing algorithms and image sources: a multi-scale block local binary pattern fusion approach. *IET Biom.*, 9(6):278–289, 2020. 2
- [35] Ulrich Scherhag, Andreas Nautsch, Christian Rathgeb, Marta Gomez-Barrero, Raymond N. J. Veldhuis, Luuk J. Spreuwers, Maikel Schils, Davide Maltoni, Patrick Grother, Sébastien Marcel, Ralph Breithaupt, Ramachandra Raghavendra, and Christoph Busch. Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting. In *BIOSIG*, volume P-270 of *LNI*, pages 149–159. GI / IEEE, 2017. 4
- [36] Ken Shoemake. Animating rotation with quaternion curves. In *SIGGRAPH*, pages 245–254. ACM, 1985. 2
- [37] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2256–2265. JMLR.org, 2015. 1, 2
- [38] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 2
- [39] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826. IEEE Computer Society, 2016. 4
- [40] Sushma Venkatesh, Haoyu Zhang, Raghavendra Ramachandra, Kiran B. Raja, Naser Damer, and Christoph Busch. Can GAN generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection. In *IWBF*, pages 1–6. IEEE, 2020. 2, 3, 4
- [41] Haoyu Zhang, Sushma Venkatesh, Raghavendra Ramachandra, Kiran Bylappa Raja, Naser Damer, and Christoph Busch. MIPGAN - generating strong and high quality morphing attacks using identity prior driven GAN. *IEEE Trans. Biom. Behav. Identity Sci.*, 3(3):365–383, 2021. 2, 3, 4
- [42] T. Zheng and W. Deng. Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments. Technical Report 18-01, Beijing University of Posts and Telecommunications, February 2018. 4