

# Leveraging Diffusion For Strong and High Quality Face Morphing Attacks

Zander Blasingame<sup>✉</sup> and Chen Liu<sup>✉</sup>, *Senior Member, IEEE*

**Abstract**—Face morphing attacks seek to deceive a Face Recognition (FR) system by presenting a morphed image consisting of the biometric qualities from two different identities with the aim of triggering a false acceptance with one of the two identities, thereby presenting a significant threat to biometric systems. The success of a morphing attack is dependent on the ability of the morphed image to represent the biometric characteristics of both identities that were used to create the image. We present a novel morphing attack that uses a Diffusion-based architecture to improve the visual fidelity of the image and the ability of the morphing attack to represent characteristics from both identities. We demonstrate the effectiveness of the proposed attack by evaluating its visual fidelity via Fréchet Inception Distance (FID). Also, extensive experiments are conducted to measure the vulnerability of FR systems to the proposed attack. The ability of a morphing attack detector to detect the proposed attack is measured and compared against two state-of-the-art GAN-based morphing attacks along with two Landmark-based attacks. Additionally, a novel metric to measure the relative strength between different morphing attacks is introduced and evaluated.

**Index Terms**—Morphing Attack, GAN, Vulnerability Analysis, Face Recognition, Diffusion Models

## 1 INTRODUCTION

FACE recognition (FR) systems have become one of the most common biometric modalities used for identity verification across a wide range of modern-day applications, from trivial tasks such as unlocking a smart phone to official businesses such as banking, e-commerce, and law enforcement. Unfortunately, while FR systems can often reach low false rejection and acceptance rates [1], they are especially vulnerable to a new class of emerging attacks, known as the face morphing attack [2], [3], [4], [5]. The face morphing attack aims to compromise a fundamental property of biometric security, *i.e.*, the one-to-one mapping from biometric data to the associated identity. This compromise is achieved by creating a morphed face which contains biometric data of both identities in such a manner that presenting *one* morphed image triggers a match with *two* disjoint identities, violating the fundamental principle.

This poses a significant threat towards FR systems, especially in application cases such as e-passports and border access. Notably the e-passport scenario, wherein the applicant submits a passport photo either in digital or printed form, is especially vulnerable to face morphing attack. This is particularly relevant for countries such as New Zealand, Estonia, and Ireland, where e-passports are used for both issuance and renewal of documents [6]. In 2018 a German activist was reported to have received a German passport with a photo of his face morphed with an Italian politician [7]. Moreover, digital morphs can be easily generated hence offer a low-cost attack in the digital domain. It is not uncommon for a digital image submitted to a document submission portal to not have its authenticity verified by a human agent [8]. Critically, an

adversary who is blacklisted from accessing a certain system can create a morph with a non-blacklisted individual to gain access.

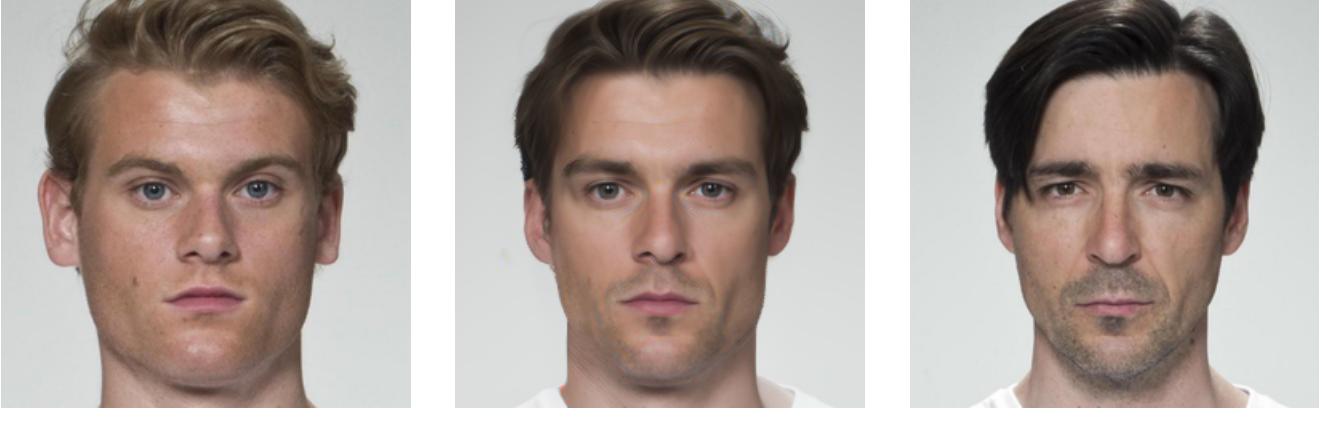
Due to the severity of face morphing attacks, an abundance of algorithms have been developed to identify these attacks [2], [9], [10], [11]. Methods for Morphing Attack Detection (MAD) can be broadly characterized into two classes based on the manner in which they obtain the features used for detection, *i.e.*, handcrafted features or deep features. Handcrafted features are used in the so-called *classical* algorithms which seek to find evidence of the morphing attack in the pixel domain, whether that is evidence of degradation in image quality [10], residual noise left from the morphing attack [12], or local geometric features such as Local Binary Patterns (LBP) [13], Binarized Statistical Image Feature (BSIF) [14], and Local Phase Quantization (LPQ) [4]. Conversely, deep features are used with *deep learning*-based algorithms. These features are extracted by a deep Convolutional Neural Network (CNN), often a pre-trained network such as ResNet150 [8], [15]. Generally, the best success has been found using deep CNN-based features in comparison to handcrafted features on both digital and print-scan face data [16], [17].

Comparatively, there has been less research on face morphing attack algorithms. Similar to the two classes of MAD algorithms, there exist two broad classes of face morphing attacks: Landmark-based attacks and deep learning-based attacks. Landmark-based morphing attacks use local features to create the morphed image by warping and aligning the landmarks within each face then create a morphed face by pixel-wise compositing. Landmark-based attacks have been shown to be effective against FR systems [18]. Recent work has enhanced the effectiveness of Landmark-based attacks by using adversarial perturbation [19]. In contrast, deep learning-based morphing attacks use a machine learning model to embed the original bona fide faces into a semantic representation which are then combined to produce a new representation that contains information from both identities. This new representation is then used to generate a morphed face.

However, nearly all state-of-the-art deep learning based morph

• Z.W. Blasingame is with the Department of Electrical and Computer Engineering, Clarkson University, Potsdam, NY, 13699.  
E-mail: blasinzw@clarkson.edu

• C. Liu is with the Department of Electrical and Computer Engineering, Clarkson University, Potsdam, NY, 13699.  
E-mail: cliu@clarkson.edu

(a) Identity *a*

(b) Morphed image

(c) Identity *b*

Fig. 1: Example of the proposed Diffusion-based morphing attack. Samples are from FRLL dataset.

methods were based on the Generative Adversarial Network (GAN) framework [20], [21], with the primary difference being architectural improvements such as using the StyleGAN2 architecture [22] over the vanilla GAN architecture, or changes to the morph generation pipeline as seen in the Morphing through Identity Driven Prior GAN (MIPGAN) [23]. At the same time, there exists a handful of alternative state-of-the-art deep generative models outside the GAN framework which offer their own advantages in terms of visual fidelity, semantic representation capabilities, and inference speed.

In particular, a class of generative models collectively known as “Diffusion models” have been shown to possess high visual fidelity, even beating state-of-the-art GANs in visual fidelity [24], at the cost of increased inference time. As visual fidelity and semantic representation abilities are far more important for the potency of a morphing attack than inference speed, we present a novel methodology for generating strong face morphing attacks by leveraging Diffusion-based methods. Figure 1 shows an exemplary morphed image generated by the proposed attack constructed from two identities from the FRLL dataset [25]. We summarize the contributions of the proposed work as shown below:

- We propose a novel method for generating morphed faces by using a Diffusion-based model which calculates twin embeddings to generate images of high visual fidelity.
- We evaluate our proposed attack against four other morphing attacks with extensive experiments assessing the vulnerability of three FR systems across three different datasets.
- The proposed morphing attack is evaluated via the Fréchet Inception Distance, a quantitative measure of visual fidelity.
- The proposed attack is further evaluated on its ability to evade detection from MAD algorithms trained against other morphing attacks.
- We introduce a novel metric to measure the strength of one morphing attack relative to another.
- We present a small-scale study on the impact of pre-processing in the FR pipeline on the vulnerability of FR systems to morphing attacks.
- An exploration of different interpolation techniques on the proposed morphing attack is presented and evaluated.

## 2 PRIORITY WORK

Several face morphing attacks have been developed by researchers with the morph generating process generally using face landmarks or deep learning. In particular, we compare our proposed Diffusion-based attack against four state-of-the-art morphing attacks, two deep learning-based and two Landmark-based; namely, the OpenCV, FaceMorpher, StyleGAN2, and MIPGAN-II face morphing attacks. These schemes represent different types of morphing attacks and provide a substantial baseline to measure the performance of the proposed Diffusion-based attack. To the best of our knowledge all previously proposed deep learning-based attacks generate the morphed images via a type of GAN architecture [18], [20], [23], [26].

### 2.1 Landmark-Based Morphing Attacks

The FaceMorpher and OpenCV attacks were chosen as they are commonly used to represent Landmark-based attacks [11], [18], [27].

**FaceMorpher** is an open-source algorithm that uses the STASM landmark detector [28], [29]. From the landmarks on the images Delaunay triangles are formed, which are then warped and blended together. The areas outside the landmarks are averaged, typically introducing strong artifacts in the neck and hair regions of the image [30].

**OpenCV** morphing attack uses the open-source OpenCV library with a 68-point annotator from the Dlib library [31]. The images and associated landmarks are used to form Delaunay triangles. Then, in a similar manner to FaceMorpher, the landmarks are warped and blended. In contrast to the approach of FaceMorpher, the areas outside the landmarks do not consist of an averaged image, but rather additional Delaunay triangles. However, these morphs also exhibit strong artifacts outside the facial area due to the missing landmarks [30].

### 2.2 GAN-Based Morphing Attacks

As mentioned earlier, prior deep learning-based attacks have used a GAN architecture for the morph generation process [27], [32]. GANs are a type of deep generative model which seeks to learn the sampling process for some data distribution  $\mathbb{P}_{data}$  on  $\mathcal{X}$ , *i.e.*, given some simple distribution  $Z \sim p(z)$  on  $\mathcal{Z}$ , the generator  $G : \mathcal{Z} \rightarrow \mathcal{X}$  is to learn  $G(Z) \sim \mathbb{P}_{data}$ . A discriminator, sometimes

called the critic,  $D : \mathcal{X} \rightarrow [0, 1]$  is trained adversarially against the generator in a minimax game described by

$$\min_G \max_D \mathbb{E}_{x \sim \mathbb{P}_{data}} \log D(x) + \mathbb{E}_{z \sim p(z)} \log(1 - (D \circ G)(z)) \quad (1)$$

where the discriminator attempts to get better at distinguishing synthetic samples from genuine samples, while the generator tries to get better at deceiving the discriminator. For a GAN-based morphing attack, it becomes necessary for an encoding algorithm  $E : \mathcal{X} \rightarrow \mathcal{Z}$  to exist, which can embed images in the latent space such that the inversion has low distortion  $(G \circ E)(x) \approx x$ . The latent codes for two identities are then averaged to produce a new latent code representing the morphed face which is passed to the generator. Notably, there exists a trade-off between the inversion distortion and editability of the latent embeddings [33]. Damer *et al.* proposed to use the GAN architecture for generating morphs by combining two latent codes encoded from two real identities to create a morphed code [20]. This proposed attack, known as MorGAN, was based on a modification to the vanilla GAN architecture with the addition of an explicitly defined encoder architecture that was trained jointly with the generator via a modified adversarial loss formulation. Since then the StyleGAN2 [18] and MIPGAN-II [23] attacks have improved upon the MorGAN formulation by improving the GAN architecture, loss formulation, or encoding algorithm.

**StyleGAN2** offers a host of improvements over the standard GAN implementation that enables the architecture to achieve state-of-the-art image quality when generating high resolution images. The StyleGAN2 model was pre-trained on the Flickr-Faces-HQ (FFHQ) dataset [34]. The faces were then cropped to possess the same landmark alignment as in the FFHQ dataset. Following the approach in [22], the images,  $x_a, x_b$ , are embedded by optimizing an initial latent code through stochastic gradient descent, minimizing the perceptual loss between the generated image and target image. After each embedded latent code,  $z_a, z_b$ , is found, a morphed latent code is created by linearly interpolating between the two,  $z_{ab} = \text{lerp}(z_a, z_b; 0.5)$ . Lastly, the interpolated latent code is passed to the StyleGAN2 synthesis network to get the morphed image  $x_{ab}$ . The StyleGAN2 morphs are strong when used with images containing a uniform background, which makes them especially powerful when used in conjunction with the Face Research Lab London (FRLL) dataset [25].

**MIPGAN-II** proposes an extension on StyleGAN2 by adding an optimization procedure for the latent vector used in creating the morphed image [23]. The StyleGAN2 portion of MIPGAN-II was pre-trained on the FFHQ dataset. The two bona fide images are embedded into the latent space using the StyleGAN2 optimization procedure. The latent code is initially constructed as  $z_0 = \text{lerp}(z_a, z_b; 0.5)$ . For  $n$  epochs the latent code is optimized to minimize a combination of perceptual loss, identity loss, identity difference loss, and Multi-Scale Structural Similarity loss, finding a fully optimized latent  $z_n$ . The latent code  $z_n$  is then passed to the StyleGAN2 synthesis network to create the morphed image. As MIPGAN-II presents a refinement on StyleGAN2 for the application of morphing attack, it possesses similar advantages and disadvantages that StyleGAN2 morphs offer.

### 3 DIFFUSION-BASED MORPHING ATTACK

Unlike GANs which learn the sampling process for the data distribution through adversarial training between the generator and the critic, diffusion-based and score-based generative models

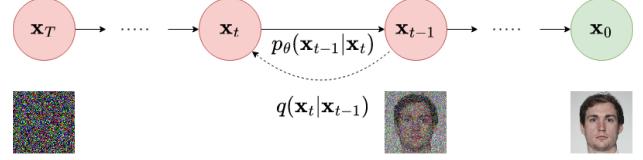


Fig. 2: The forward and reverse Diffusion processes.

learn the data distribution via a denoising process through varying noise levels. Diffusion-based models can achieve image fidelity superior to state-of-the-art generative models, matching even the acclaimed BigGAN-deep [35] model, while maintaining better coverage of the data distribution [24]. For these reasons we propose a morphing attack that uses Diffusion-based methods as the generative process.

#### 3.1 Diffusion Models

Given data distribution  $q(\mathbf{x}_0)$  on data space  $\mathcal{X}$ , the goal is to learn a model  $p_\theta(\mathbf{x}_0)$  approximating  $q(\mathbf{x}_0)$  which can be easily sampled. Denoising Diffusion Probabilistic Models (DDPMs) [36] are latent variable models of the form

$$p_\theta(\mathbf{x}_0) = \int p_\theta(\mathbf{x}_{0:T}) \, d\mathbf{x}_{1:T} \quad (6)$$

where  $\{\mathbf{x}_t\}_{t=1}^T \in \mathcal{X}$  are latent variables and  $T \in \mathbb{N}$ . The reverse process is a Markov chain starting at  $p_\theta(\mathbf{x}_T) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ , *i.e.*, a normal distribution with mean vector  $\mathbf{0}$  and variance  $\mathbf{I}$  (the identity matrix), with Gaussian transitions

$$p_\theta(\mathbf{x}_{0:T}) = p_\theta(\mathbf{x}_0) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t) \quad (7)$$

The diffusion (forward) process is fixed to a Markov chain that gradually adds Gaussian noise to the original sample  $\mathbf{x}_0$  according to variance schedule  $\{\beta_t\}_{t=1}^T$ , such that

$$q(\mathbf{x}_{1:T}, \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t \mid \mathbf{x}_{t-1}) \quad (8)$$

and  $q(\mathbf{x}_t \mid \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$ . See Figure 2 for an illustration of this process. The transition probability  $p(\mathbf{x}_{t-1} \mid \mathbf{x}_t)$  is likely to be very complex, unless the gap between  $t$  and  $t - 1$  is very small, *i.e.*,  $T \rightarrow \infty$ . In this case,  $p(\mathbf{x}_{t-1} \mid \mathbf{x}_t)$  can be modelled as  $\mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \sigma_t)$ , where  $\boldsymbol{\mu}_\theta : \mathcal{X} \times \mathbb{N} \rightarrow \mathcal{X}$  is an estimator at step  $t$  parameterized by  $\theta$ . Ho *et al.* [36] proposes to use the following form:

$$\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) \right) \quad (9)$$

where  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ , and  $\boldsymbol{\epsilon}_\theta : \mathcal{X} \times \mathbb{N} \rightarrow \mathcal{X}$  is a function approximator parameterized by  $\theta$ , which learns to predict the noise added to  $\mathbf{x}_0$  to get  $\mathbf{x}_t$ . This is achieved by using a U-Net, a type of CNN consisting of several skip connections formed at each resolution size in the architecture to model  $\boldsymbol{\epsilon}_\theta$  [37]. Like variational autoencoder (VAE), the model is trained by optimizing the variational bound on the negative log likelihood.

Relaxing the constraint that the inference process has to be Markovian leads to another kind of diffusion model known as the Denoising Diffusion Implicit Model (DDIM) as proposed by Song

**Algorithm 1** Diffusion Morphing Algorithm with DDIM scheduler where  $\sigma_t = 0$ .**Require:** The following components:

- 1) Bona fide images,  $\mathbf{x}_0^{(a)}, \mathbf{x}_0^{(b)} \in \mathcal{X}$
- 2) Noise prediction model  $\epsilon_\theta : \mathcal{X} \times \mathcal{Z} \times \mathbb{N} \rightarrow \mathcal{X}$
- 3) Image space preprocessing function,  $\xi : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$
- 4) Image space interpolation function,  $\ell_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \times [0, 1] \rightarrow \mathcal{X}$
- 5) Latent space interpolation function,  $\ell_{\mathcal{Z}} : \mathcal{Z} \times \mathcal{Z} \times [0, 1] \rightarrow \mathcal{Z}$
- 6) Timing sub-schedule  $\{\tau_i\}_{i=1}^N$

**procedure** DIFFUSIONMORPH( $\mathbf{x}_0^{(a)}, \mathbf{x}_0^{(b)}, \{\tau_i\}_{i=1}^N$ ) $\mathbf{z}_a \leftarrow E(\mathbf{x}_0^{(a)})$ 

▷ Calculate semantic latent codes

 $\mathbf{z}_b \leftarrow E(\mathbf{x}_0^{(b)})$  $\mathbf{x}_0^{(a)} \leftarrow \xi(\mathbf{x}_0^{(a)}, \mathbf{x}_0^{(b)})$ 

▷ Preprocess images passed to stochastic encoder

 $\mathbf{x}_0^{(b)} \leftarrow \xi(\mathbf{x}_0^{(b)}, \mathbf{x}_0^{(a)})$ **for**  $i \leftarrow 1, 2, \dots, N$  **do** $\mathbf{x}_{\tau_i}^{(a)} \leftarrow \sqrt{\alpha_{\tau_i}} \left( \frac{\mathbf{x}_{\tau_{i-1}}^{(a)} - \sqrt{1 - \alpha_{\tau_{i-1}}} \cdot \epsilon_\theta(\mathbf{x}_{\tau_{i-1}}^{(a)}, \mathbf{z}_a, \tau_{i-1})}{\sqrt{\alpha_{\tau_{i-1}}}} \right) + \sqrt{1 - \alpha_{\tau_i}} \epsilon_\theta(\mathbf{x}_{\tau_{i-1}}^{(a)}, \mathbf{z}_a, \tau_{i-1})$  ▷ Forward pass of diffusion algorithm $\mathbf{x}_{\tau_i}^{(b)} \leftarrow \sqrt{\alpha_{\tau_i}} \left( \frac{\mathbf{x}_{\tau_{i-1}}^{(b)} - \sqrt{1 - \alpha_{\tau_{i-1}}} \cdot \epsilon_\theta(\mathbf{x}_{\tau_{i-1}}^{(b)}, \mathbf{z}_b, \tau_{i-1})}{\sqrt{\alpha_{\tau_{i-1}}}} \right) + \sqrt{1 - \alpha_{\tau_i}} \epsilon_\theta(\mathbf{x}_{\tau_{i-1}}^{(b)}, \mathbf{z}_b, \tau_{i-1})$ **end for** $\mathbf{x}_T^{(ab)} \leftarrow \ell_{\mathcal{X}}(\mathbf{x}_T^{(a)}, \mathbf{x}_T^{(b)}; 0.5)$ 

▷ Stochastic code interpolation

 $\mathbf{z}_{ab} \leftarrow \ell_{\mathcal{Z}}(\mathbf{z}_a, \mathbf{z}_b; 0.5)$ 

▷ Semantic code interpolation

**for**  $i \leftarrow N, N-1, \dots, 1$  **do** $\mathbf{x}_{\tau_{i-1}}^{(ab)} \leftarrow \sqrt{\alpha_{\tau_{i-1}}} \left( \frac{\mathbf{x}_{\tau_i}^{(ab)} - \sqrt{1 - \alpha_{\tau_i}} \epsilon_\theta(\mathbf{x}_{\tau_i}^{(ab)}, \mathbf{z}_{ab}, \tau_i)}{\sqrt{\alpha_{\tau_i}}} \right) + \sqrt{1 - \alpha_{\tau_{i-1}}} \cdot \epsilon_\theta(\mathbf{x}_{\tau_i}^{(ab)}, \mathbf{z}_{ab}, \tau_i)$  ▷ Diffusion generative process**end for****return**  $\mathbf{x}_0^{(ab)}$ **end procedure**

$$q_\sigma(\mathbf{x}_{t-1} \mid \mathbf{x}_t, \mathbf{x}_0) = \mathcal{N} \left( \sqrt{\alpha_{t-1}} \mathbf{x}_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \frac{\mathbf{x}_t - \sqrt{\alpha_t} \mathbf{x}_0}{\sqrt{1 - \alpha_t}}, \sigma_t^2 \mathbf{I} \right) \quad (2)$$

$$\mathbf{x}_{t-1} = \underbrace{\sqrt{\alpha_{t-1}} \left( \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\alpha_t}} \right)}_{\text{Predicted } \mathbf{x}_0} + \underbrace{\sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \epsilon_\theta(\mathbf{x}_t, t)}_{\text{Direction pointing to } \mathbf{x}_t} + \underbrace{\sigma_t \epsilon_t}_{\text{Random Noise}} \quad (3)$$

$$\mathbf{x}_{\tau_{i-1}} = \sqrt{\alpha_{\tau_{i-1}}} \left( \frac{\mathbf{x}_{\tau_i} - \sqrt{1 - \alpha_{\tau_i}} \epsilon_\theta(\mathbf{x}_{\tau_i}, \tau_i)}{\sqrt{\alpha_{\tau_i}}} \right) + \sqrt{1 - \alpha_{\tau_{i-1}}} \cdot \epsilon_\theta(\mathbf{x}_{\tau_i}, \tau_i) \quad (4)$$

$$\mathbf{x}_{\tau_i} = \sqrt{\alpha_{\tau_i}} \left( \frac{\mathbf{x}_{\tau_{i-1}} - \sqrt{1 - \alpha_{\tau_{i-1}}} \cdot \epsilon_\theta(\mathbf{x}_{\tau_{i-1}}, \mathbf{z}, \tau_{i-1})}{\sqrt{\alpha_{\tau_{i-1}}}} \right) + \sqrt{1 - \alpha_{\tau_i}} \epsilon_\theta(\mathbf{x}_{\tau_{i-1}}, \mathbf{z}, \tau_{i-1}) \quad (5)$$

et al. [38]. Consider a family  $\mathcal{Q}$  of inference distributions indexed by a real vector  $\sigma \in \mathbb{R}_{\geq 0}^T$

$$q_\sigma(\mathbf{x}_{1:T} \mid \mathbf{x}_0) = q_\sigma(\mathbf{x}_T \mid \mathbf{x}_0) \prod_{t=2}^T q_\sigma(\mathbf{x}_{t-1} \mid \mathbf{x}_t, \mathbf{x}_0) \quad (10)$$

where  $q_\sigma(\mathbf{x}_T \mid \mathbf{x}_0) = \mathcal{N}(\sqrt{\alpha_T} \mathbf{x}_0, (1 - \alpha_T) \mathbf{I})$  and for all  $t > 1$  Equation (2) holds. This leads to the forward equation, Equation (3). Such a formulation allows for a deterministic generation of an image when  $\sigma_t = 0$ .

Another advantage of the DDIM framework is the ability to sample the forward trajectory more quickly by skipping timesteps. Let  $\{\tau_i\}_{i=1}^N$  be a monotonically increasing sequence of samples with  $N < T$ . By convention  $\tau_N = T$  and an initial timestep  $\tau_0 = 0$  is inserted. The inference equation can be reformulated into Equation (4) where the parameter  $N$  controls a trade-off between sample speed and visual fidelity.

Preechakul et al. [39] proposed a Diffusion Autoencoder by employing a conditional DDIM. The DDIM is conditioned on a semantic representation  $\mathbf{z} \in \mathcal{Z}$  by modifying  $\epsilon_\theta^{(t)}$  to take  $\mathbf{z}$  as an additional input. An encoding network  $E : \mathcal{X} \rightarrow \mathcal{Z}$  is introduced to learn the semantic latent code  $\mathbf{z} = E(\mathbf{x})$ . In this framework a twin pair of latent codes are created for a given image,  $(\mathbf{z}, \mathbf{x}_T)$ , the semantic and stochastic latent codes. In theory, the semantic code controls semantic information like hair color, face shape, and other high level concepts, whereas the stochastic code controls stochastic variations in the image. Instead of sampling  $\mathbf{x}_T \sim q_\sigma(\mathbf{x}_T \mid \mathbf{x}_0)$ , Preechakul et al. proposed Equation (5) as a deterministic forward process. This implementation necessitates that the semantic code  $\mathbf{z}$  be generated first before creating the stochastic code.

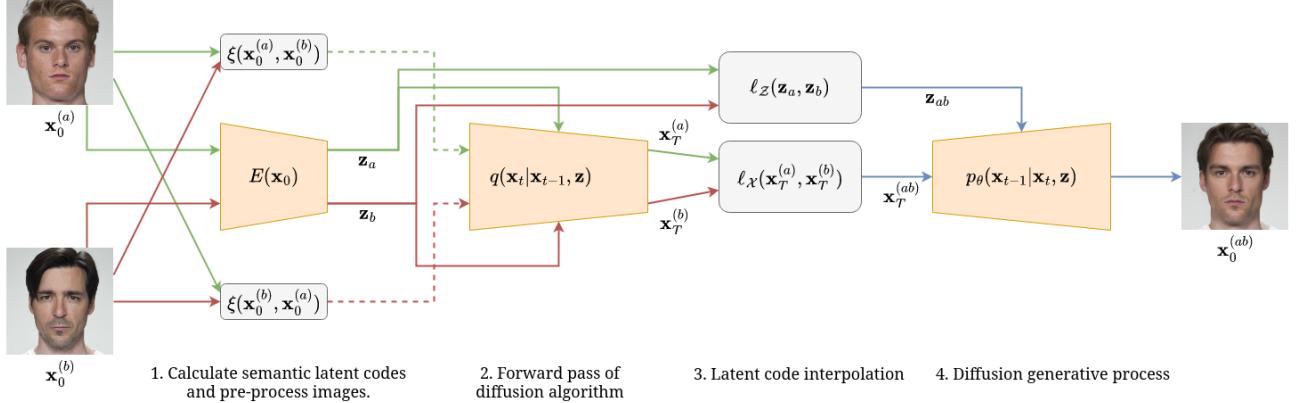


Fig. 3: Proposed architecture for Diffusion-based morphs, where the green traces indicate variables associated with identity  $a$ , likewise red traces denote identity  $b$ , and blue traces for the morphed identity  $ab$ .

### 3.2 Proposed Morphing Algorithm

We propose a novel process for the creation of morphed images by employing both the stochastic and semantic encoders. In particular, let  $x_a, x_b \in \mathcal{X}$  be two bona fide images of identities  $a, b$ , and let  $x_0^{(a)} = x_a$  and  $x_0^{(b)} = x_b$ . Algorithm 1 outlines the structure of the proposed Diffusion-based morphing attack, hereafter called the Diffusion attack for simplicity, with additional illustration provided in Figure 3.

Beyond the core components of the DDIM and semantic encoder, three additional functions are added to the architecture, namely, the image space preprocessing function,  $\xi : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$ , image space interpolation function,  $\ell_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \times [0, 1] \rightarrow \mathcal{X}$ , and latent space interpolation function,  $\ell_{\mathcal{Z}} : \mathcal{Z} \times \mathcal{Z} \times [0, 1] \rightarrow \mathcal{Z}$ .

The interpolation functions are used to interpolate between the semantic and stochastic values by some factor  $\gamma \in [0, 1]$ . The image space preprocessing function is used to prepare the image passed to the semantic encoder. The simplest form of the interpolation function is the linear interpolation function,  $\text{lerp}(a, b; \gamma) = \gamma a + (1 - \gamma)b$ , where linear interpolation was found to be the best choice for  $\ell_{\mathcal{Z}}$ . However, Song *et al.* [38] suggests the usage of spherical linear interpolation for  $\ell_{\mathcal{X}}$ . For a vector space  $V$  and two vectors  $u, v \in V$ , the spherical interpolation by a factor of  $\gamma$  is given as

$$\text{slerp}(u, v; \gamma) = \frac{\sin((1 - \gamma)\theta)}{\sin \theta} u + \frac{\sin(\gamma\theta)}{\sin \theta} v \quad (11)$$

where  $\theta = \frac{\arccos(u \cdot v)}{\|u\| \|v\|}$ .

While the semantic code provides most of the fundamental information, such as positioning of facial features, the stochastic code is used to provide information on the details not explicitly associated with the identity, but necessary for the realism of the generated image. By altering the stochastic code, details such as direction of strands of hair, clothing, *etc.*, are altered whilst the identity of the image is preserved [39]. Unlike the rather straightforward nature of linearly interpolating between the semantic codes to produce an image with key identifying characteristics of both identities, the nature of the stochastic code can lead to images of low visual fidelity if the interpolation is not done carefully. In particular, linear interpolation between two stochastic codes does guarantee a smooth interpolation between the stochastic details in the images. For this reason the preprocessing function,  $\xi$ , is used to prepare the image passed to the stochastic encoder. One strategy

is to “pre-morph” the image when extracting the stochastic details, *i.e.*,  $\xi$  performs an image space morph of the image with the goal of reducing the artifacts induced by stochastic interpolation.

## 4 EXPERIMENTAL SETUP

To evaluate its effectiveness, we tested the proposed morphing attack on three datasets against two different state-of-the-art FR systems. All training, optimization, and evaluation was conducted on a system with dual Intel Xeon Silver 4114 CPUs and an NVIDIA Tesla V100 32GB GPU with CUDA version 10.1 and CUDNN version 8.4. The proposed morphing attack, MAD algorithm, and the FR systems are implemented in PyTorch [40].

### 4.1 Face Recognition Systems

In order to evaluate the strength of the proposed morphing attack, three publicly available FR systems are used, specifically, the FaceNet<sup>1</sup>, VGGFace2<sup>2</sup>, and ArcFace<sup>3</sup> models. These models are representative recognition systems with state-of-the-art face verification performance [41], [42], [43] and hence are widely used. For both models the last fully connected layer is used to provide a rich feature representation of the input image. Then for a presented face, its feature vector is compared with that of the feature vector belonging to the target face. If the distance between these two representations is sufficiently “small”, the presented face is then said to have the same identity as the target face. The VGGFace2 model improves upon acclaimed VGGFace [44] by using an improved training dataset, also called VGGFace2. Following the introduction of Squeeze and Excitation Network (SENet) by Hu *et al.* [45], Cao *et al.* [41] presented the SENet architecture as the optimal choice when used with the VGGFace2 dataset. Google’s FaceNet model consists of an Inception-ResNet V1 architecture which is pre-trained on the VGGFace2 dataset [42]. A new state-of-the-art model, ArcFace [43], uses a novel loss function during training to improve the embeddings of faces. This loss function is known as Additive Angular Margin Loss or ArcFace Loss. The particular ArcFace model used for evaluation consists

1. <https://pypi.org/project/facenet-pytorch>
2. [https://github.com/ox-vgg/vgg\\_face2](https://github.com/ox-vgg/vgg_face2)
3. <https://github.com/deepinsight/insightface>

of a 100-layer Improved ResNet [46] trained on the Glint360K dataset<sup>4</sup> [47].

Additionally, the three FR systems use different pre-processing pipelines. Across all datasets the images and generated morphs are cropped as to be appropriate for passport photos; consequently, a face extractor such as MTCNN [48] is omitted from the verification pipeline. The FaceNet model resizes the image such that the short side of the image is 180 pixels long and then the image is cropped to a  $160 \times 160$  resolution. Lastly, the images are normalized to  $[-1, 1]$ . The VGGFace2 model resizes the image such that the short side of the image is 256 pixels long and then crops the image to  $224 \times 224$  pixels. The mean RGB vector<sup>5</sup> is subtracted from the cropped image to normalize the image. The ArcFace model resizes images such that the short side of the image is 112 pixels long which is then cropped to a  $112 \times 112$  resolution. The image is then normalized to have values in  $[-1, 1]$ .

## 4.2 Datasets

In this work, the FERET [49], FRLL [25], and FRGC v2.0 [50] datasets were used to evaluate the proposed technique, as they are commonly used in MAD with a large number of different identities [18], [27]. Notably, the FRLL dataset consists of high quality close-up frontal images at a  $1350 \times 1350$  resolution with 189 facial landmarks—a large number of landmarks. The StyleGAN2, MIPGAN-II, and diffusion models were all trained on the FFHQ dataset, which contains 70,000 images at a  $1024 \times 1024$  resolution [34]. Morphs using OpenCV, FaceMorpher, and StyleGAN2 were created by Sarkar *et al.* [18] on the FRLL, FERET, and FRGC datasets. Additionally, Zhang *et al.* [23] created morphs via MIPGAN-II on the three datasets.

In order to create a morphed face, two component identities are needed. Naturally, if the two component identities are disparate, the resulting morph is likely to be very weak. To rectify this and for evaluation purposes, the component identity pairs were selected by following the existing protocol offered by Sarkar *et al.* [18]. These pairings resulted in 1222 unique morphs on FRLL, 964 on FRGC, and 529 on FERET.

## 5 RESULTS

The proposed morphing attack is compared to state-of-the-art techniques drawing from both GAN-based and Landmark-based methods. The effectiveness of the proposed method is quantitatively measured on three fronts, being the visual fidelity of the generated morphed images, the vulnerability of state-of-the-art FR systems to the morphing attack, and the detection potential of the morphing attack, respectively. Furthermore, an exploration of interpolation techniques for the stochastic latent code is provided.

### 5.1 Evaluation of Visual Fidelity

The visual fidelity of the Diffusion attack is compared against other morphing attacks. Whilst on first glance it may appear that the ability to deceive an FR system should imply a high level of visual fidelity, this is not a simple assertion. We posit two reasons for this discrepancy:

4. The Glint360K dataset consists of 17,091,657 imagees of 360,232 individuals. Models trained on this dataset can achieve state-of-the-art performance.

5. The mean vector is specifically  $\langle 131.0912, 103.8827, 91.4953 \rangle$  for the red, green, and blue channels.

TABLE 1: FID across different morphing attacks. Lower is better.

Morphing Attack	FRLL	FRGC	FERET
StyleGAN2	45.19	86.41	<b>41.91</b>
FaceMorpher	91.97	88.14	79.58
OpenCV	85.71	100.02	91.94
MIPGAN-II	66.41	115.96	70.88
Diffusion	<b>42.63</b>	<b>64.16</b>	50.45

- 1) The image pre-processing pipeline for an FR system may crop out a significant portion of artifacts in the original morphed image.
- 2) The FR system can be vulnerable to certain adversarial morphing attacks even though they inject noticeable artifacts into the morphed image.

This can lead to a situation wherein the FR system is fooled by a morphed image; however, it would be trivial for a human agent to notice the artifacts present in the image. Moreover, a deep learning system could be specifically trained to notice such artifacts, greatly reducing the potential of such an attack to go undetected.

To quantitatively assess the visual fidelity of the generated images, the Fréchet Inception Distance (FID) is employed, as it has shown to correlate well with human assessment of fidelity [51]. The FID is a measure of distance between the generated and target distributions. Therefore, the lower the FID metric, the more similar the generated distribution is to the target distribution, which correlates well with visual fidelity. The metric is defined as the Fréchet distance, or 2-Wasserstein metric<sup>6</sup>, between two Gaussian distributions, each representing the activations the deepest layer of an Inception v3 network<sup>7</sup> induced by images from the generated and target distributions.

Table 1 shows the FID metric between the generated images from different morphing attacks and bona fide samples from the dataset the morphing attack is drawn from. The FID metric is calculated using `pytorch-fid` [53]. What we can see is that morphed images generated from the Diffusion attack generally has the lowest FID, with the StyleGAN2-based attack following closely behind. Both Landmark-based morphs—OpenCV and FaceMorpher—has noticeably higher FIDs than the deep learning-based morphs. These results correlate well with visual inspection of the morphs as Figures 4b, 4f, 4i and 4m exhibit prominent artifacts outside the central face region. Likewise, the MIPGAN-II attack seems to struggle with some distortion outside the central face region, as can be seen in Figures 4e and 4l. Interestingly, on the FRLL dataset the StyleGAN2 morphing pipeline consistently darkens morphs relative to its component images; however, the visual fidelity is relatively high albeit noticeable darkening. Importantly, stochastic details such as hair seem to be modelled well by the Diffusion attack, while other attacks distort such details for the cases of OpenCV, FaceMorpher, and MIPGAN-II attacks, or present details that have little similarity to both identities for the case of StyleGAN2 attack. While exhibiting far less visual artifacts than other morphing techniques, the Diffusion attack tends to

6. The 2-Wasserstein metric between two probability measures  $\mu, \nu$  with finite moments on  $\mathbb{R}^n$  is defined as

$$W_2(\mu, \nu) = \left( \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^n \times \mathbb{R}^n} \|x - y\|_2^2 \, d\pi(x, y) \right)^{\frac{1}{2}}$$

where  $\Pi(\mu, \nu)$  is the set of all distributions with marginals  $\mu$  and  $\nu$ .

7. The Inception v3 network is trained on the publically available ImageNet dataset [52].

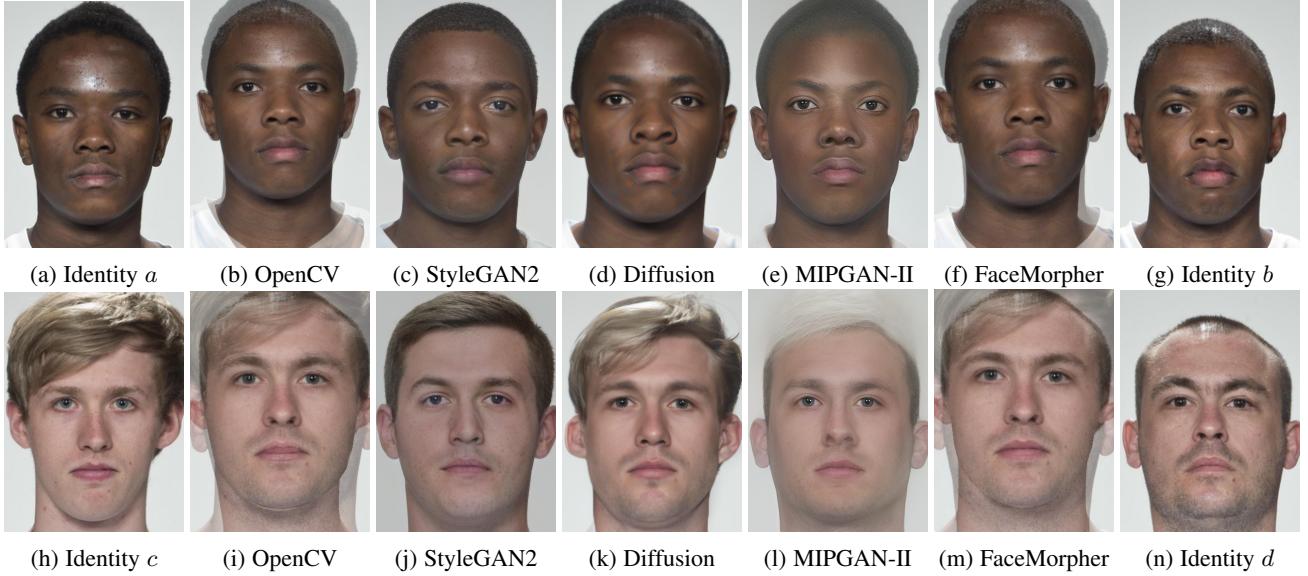


Fig. 4: Comparison across different morphing algorithms of two identity pairs from the FRLL dataset.

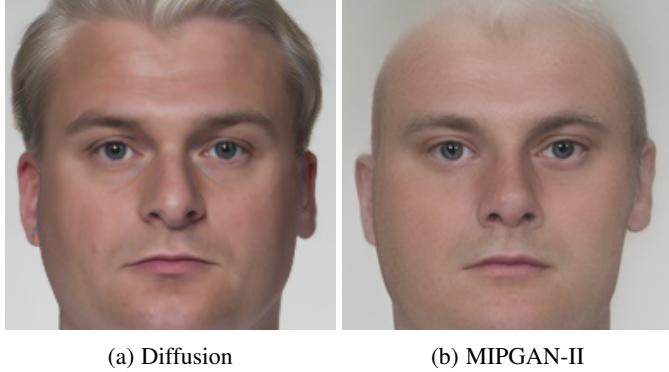


Fig. 5: Comparison of Diffusion and MIPGAN-II morphed faces on FRLL. Images are resized to  $256 \times 256$  and cropped to  $224 \times 224$  to match VGGFace2 pre-processing pipeline.

slightly smooth out the skin texture. Overall, the Diffusion attack exhibits the highest consistent visual fidelity among all presented attacks.

Interestingly, the MIPGAN-II morphs exhibit a much higher FID than the StyleGAN2 morphs, despite the two approach being based on the same deep learning backbone, *i.e.*, the StyleGAN2 model. Figure 5 compares two morphed faces generated by the Diffusion and MIPGAN-II attacks on the FRLL dataset. High frequency artifacts can be easily observed in Figure 5b, particularly, near the hairline and the transition region between hair and the background. Comparing Figure 5a and Figure 5b, the hair generated by the MIPGAN-II attack looks unnatural with a strange texture as though an image sharpening filter has been applied to the image, greatly enhancing the magnitude of high frequency content, which aligns with the observation in Figure 4. Moreover, the MIPGAN-II images seem to be de-saturated when compared to images produced by other attacks, leading to a washed-out appearance. One possible explanation for the low visual fidelity is the identity loss overpowering the perceptual quality loss, leading to morphed images with low visual fidelity but high effectiveness.

against FR systems.

## 5.2 Vulnerability of FR Systems

The strength of the proposed face morphing algorithm is further evaluated by measuring the ability of the morph to deceive an FR system. The attack success is quantitatively verified against two state-of-the-art FR systems. To ensure a valid comparison across five different morphing attacks, the same pairs of component identities were used in evaluating every morphing attack, *i.e.*, for every pair of component identities a morphed image was created for each of the five attacks. For the FaceNet, VGGFace2, and ArcFace FR systems, the False Match Rate (FMR) is set at 0.1% following the guidelines of Frontex [54]. Additionally, the distance between faces is measured using the  $L^2$  distance between the outputs of the FR model. All measurements were collected by embedding the bona fide, imposter, and morphed images from all morphing attacks using the three FR systems. Any derivative metrics of performance were calculated using the embeddings we collected.

The vulnerability of FR systems to morphing attacks is assessed by comparing the error rates in detection. Specifically, the Attack Presentation Classification Error Rate (APCER)<sup>8</sup> is measured at specific Bona fide Presentation Classification Error Rate (BPCER)<sup>9</sup> values. The FR systems are treated as differential morph detectors, *i.e.*, two images are sent to the FR model, one of which is bona fide and the other is unknown. The  $L^2$  distance between the embeddings of bona fide image and the unknown image is calculated. A threshold is set to obtain various desired BPCER values. If the distance is less than this threshold the unknown sample is classified as bona fide; otherwise, it is not. In Table 2 the APCER values for the five different morphing attacks are presented across all three datasets evaluated on three different BPCER values of 0.1%, 1%, and 5%, respectively. Due to a variety of factors—such as image quality and number of

8. APCER is the proportion of attack presentations incorrectly classified as bona fide presentations.

9. BPCER is the proportion of bona fide presentations incorrectly classified as attack presentations.

TABLE 2: The APCER at specific BPCER values. Higher is better.

Dataset	FR System	Morphing Attack	APCER @ BPCER = 0.1%	APCER @ BPCER = 1%	APCER @ BPCER = 5%
FRLL	FaceNet	StyleGAN2	0.99	0.05	0
	FaceNet	FaceMorpher	2.25	0.14	0.05
	FaceNet	OpenCV	3.24	0.33	0
	FaceNet	MIPGAN-II	<b>8.87</b>	0.47	0.09
	FaceNet	Diffusion	8.83	<b>0.99</b>	<b>0.23</b>
	VGGFace2	StyleGAN2	0.05	0.05	0
	VGGFace2	FaceMorpher	1.36	1.08	0.23
	VGGFace2	OpenCV	2.35	<b>2.11</b>	0.28
	VGGFace2	MIPGAN-II	1.31	0.99	0.23
	VGGFace2	Diffusion	<b>2.68</b>	2.07	<b>0.52</b>
FRGC	ArcFace	StyleGAN2	0	0	0
	ArcFace	FaceMorpher	0	0	0
	ArcFace	OpenCV	0	0	0
	ArcFace	MIPGAN-II	0.05	0	0
	ArcFace	Diffusion	<b>3.33</b>	<b>0.28</b>	<b>0.09</b>
	FaceNet	StyleGAN2	74.04	36.69	17.73
	FaceNet	FaceMorpher	87.9	38.85	14.9
	FaceNet	OpenCV	84.1	31.43	11.36
	FaceNet	MIPGAN-II	<b>96.54</b>	<b>61.9</b>	<b>33.48</b>
	FaceNet	Diffusion	91.73	48.86	24.95
FERET	VGGFace2	StyleGAN2	81.42	46.22	26.7
	VGGFace2	FaceMorpher	95.42	63.65	38.18
	VGGFace2	OpenCV	<b>95.31</b>	<b>64.62</b>	<b>39.4</b>
	VGGFace2	MIPGAN-II	91.92	57.8	30.84
	VGGFace2	Diffusion	93.71	58.25	32.18
	ArcFace	StyleGAN2	18.44	6.26	0.93
	ArcFace	FaceMorpher	26.11	5.44	0.37
	ArcFace	OpenCV	2.12	0.07	0
	ArcFace	MIPGAN-II	30.39	12.18	2.79
	ArcFace	Diffusion	<b>40.6</b>	<b>20.41</b>	<b>7.49</b>
FRGC	FaceNet	StyleGAN2	15.65	9.35	2.49
	FaceNet	FaceMorpher	10.71	5.1	0.91
	FaceNet	OpenCV	8.79	3.06	0.17
	FaceNet	MIPGAN-II	21.03	10.77	2.21
	FaceNet	Diffusion	<b>24.04</b>	<b>13.95</b>	<b>4.99</b>
	VGGFace2	StyleGAN2	54.08	18.42	5.73
	VGGFace2	FaceMorpher	80.5	32.65	12.7
	VGGFace2	OpenCV	<b>81.01</b>	32.6	12.87
	VGGFace2	MIPGAN-II	66.5	18.14	5.84
	VGGFace2	Diffusion	80.9	<b>35.2</b>	<b>14.34</b>
FERET	ArcFace	StyleGAN2	0.51	0	0
	ArcFace	FaceMorpher	0.17	0	0
	ArcFace	OpenCV	0	0	0
	ArcFace	MIPGAN-II	0.06	0	0
	ArcFace	Diffusion	<b>9.69</b>	<b>2.27</b>	<b>0.06</b>

bona fide images per identity—the results vary across different datasets; while there is some variance across different FR systems, they tend to agree more closely. Noticeably, all attacks perform rather poorly on the FRLL dataset, although the Diffusion attack performs the best among them, which could be attributed to the limited number of bona fide images per identity; for FRLL dataset there are only two bona fide images per identity: a neutral face (used to create the morph) and a smiling face. All FR systems were more vulnerable to the different morphing attacks when evaluated on the FRGC dataset. The MIPGAN-II attack performed very well against FaceNet on the FRGC dataset, which makes sense as this technique was refined on the FRGC dataset in particular [23]. However, this attack was not as strong against VGGFace2. Instead, that FR system was more vulnerable to OpenCV, FaceMorpher, and Diffusion. This could be attributed to the different pre-processing pipelines. The Diffusion-based morphing attack generally performs close to the top performer

on all FR systems. As with FRGC, on the FERET dataset VGGFace2 is more vulnerable to landmark-based attacks, OpenCV and FaceMorpher, than FaceNet. Diffusion-based morphs pose the greatest threat on FERET, consistently having high APCER values. In general, the following observations can be drawn from Table 2:

- Among the five different attacks, FR systems are most vulnerable to Diffusion attacks. Moreover, Diffusion attacks always rank in the top three in terms of performance.
- FR systems are the least vulnerable to the StyleGAN2 attack. The StyleGAN2 attack is consistently outperformed by its successor, MIPGAN-II, and the other deep learning-based attack, Diffusion, while often falling behind landmark-based attacks.
- Even though the state-of-the-art ArcFace model is very resilient to morphing attacks, it presents a significant vulnerability to Diffusion morphs, compared to other mor-

TABLE 3: MMPMR at FMR = 0.1% across different morphing attacks. Higher is better.  $\dagger$  the geometric mean.

Morphing Attack	FRLL			FRGC			FERET			Mean $\dagger$
	FaceNet	VGGFace2	ArcFace	FaceNet	VGGFace2	ArcFace	FaceNet	VGGFace2	ArcFace	
StyleGAN2	4.69	6.05	19.89	0.18	0.85	5.49	0.54	0.76	4.95	2.15
FaceMorpher	11.26	36.4	45.03	0.51	9.15	41.28	2.3	10.78	<b>60.73</b>	12.05
OpenCV	17.34	<b>40.93</b>	47.7	0.14	<b>12.16</b>	3.99	1.69	11.12	4.61	6.47
MIPGAN-II	<b>30.96</b>	26.74	56.52	<b>3.12</b>	7.94	33.54	6	5.39	18.19	14.16
Diffusion	28.14	35.37	<b>88.09</b>	2.68	8.47	<b>46.74</b>	<b>6.47</b>	<b>13.03</b>	59.75	<b>19.80</b>

phing attacks.

In addition to using the error rates to assess the vulnerability of FR systems, the Mated Morphed Presentation Match Rate (MMPMR) [55] is used as a measure of vulnerability. Scherhag *et al.* [55] proposed two variants of the MMPMR metric for the scenario in which multiple bona fide images of an identity were used in morph process, excluding the image used to create the morph, called the MinMax-MMPMR and ProdAvg-MMPMR. The MinMax-MMPMR metric is likely to increase the number of accepted morphs as the number of bona fide images per identity increases. Therefore, the ProdAvg-MMPMR is the specific MMPMR variant used to assess the vulnerability of FR systems. Any mention hereafter to MMPMR refers specifically to ProdAvg-MMPMR unless stated otherwise.

Let  $\mathbb{P}_M \in \mathcal{P}(\mathcal{X})$  be the distribution of morphed images such that for some  $x_{ab} \sim \mathbb{P}_M$ ,  $x_{ab}$  denotes a morphed image made from identities  $a, b$ , where  $\mathcal{P}(\mathcal{X})$  denotes the set of all probability measures on  $\mathcal{X}$ . Let  $\mathbb{P}_k \in \mathcal{P}(\mathcal{X})$  denote the distribution of bona fide images of identity  $k$ . Then with abuse of notation  $\mathbb{P}_{k \setminus x_{ab}}$  is the distribution of bona fide images of identity  $k$  excluding those images used in creating the morph  $x_{ab}$ . The MMPMR metric for a particular threshold,  $\gamma > 0$ , equipped with FR system  $F : \mathcal{X} \rightarrow \mathcal{V}$  is then defined as

$$M(\gamma) = \mathbb{E}_{x_{ab} \sim \mathbb{P}_M} \left[ \prod_{k \in \{a, b\}} \mathbb{E}_{x \sim \mathbb{P}_{k \setminus x_{ab}}} \left[ \|F(x_{ab}) - F(x)\|_2 < \gamma \right] \right]$$

i.e., the expected success rate of the morphing attack to deceive the FR system. The product term is the joint probability of successful verification of both identities.

Table 3 presents the MMPMR metric when the FMR is set at 0.1% for all datasets and FR systems. Interestingly, the FRLL dataset has the highest overall MMPMR metrics in contrast to the results from Table 2. This can likely be attributed to limited number of bona fide images per identity in the FRLL dataset, in contrast with other datasets, as the particular choice of MMPMR metric heavily punishes failed verifications for either identity. Therefore, with FRLL only having one possible bona fide image per identity the MMPMR metric could be skewed higher relative to the other datasets. On average the Diffusion attack greatly outperforms the other attacks; conversely, the Landmark-based attacks on average exhibit mediocre performance. In agreement with Table 2 the StyleGAN2 attack shows abysmal performance in comparison with the other attacks. It is noteworthy the state-of-the-art ArcFace model seems more vulnerable than the other FR systems, and is particular vulnerable to the Diffusion attack.

### 5.2.1 The Effect of Pre-processing on an FR System

Here the impact of the pre-processing pipeline on the vulnerability of an FR system is examined. In particular the cropping process is further explored. To study this an additional margin size is added

TABLE 4: APCER at FMR = 0.1% across different margin sizes on the FaceNet FR system. Higher is Better.

Dataset	Morphing Attack	Margin Size			
		0	20	40	80
FRLL	MIPGAN-II	54.84	56.53	57.18	58.03
	StyleGAN2	15.12	15.57	17.14	25.11
	FaceMorpher	74.48	75.26	73.86	47.91
	OpenCV	76.21	75.82	74.96	48.4
	Diffusion	51.25	54.47	57.07	59.02
FERET	MIPGAN-II	19.33	21.71	22.11	26.36
	StyleGAN2	14.12	17.57	18.14	22.17
	FaceMorpher	36.11	36.45	36.28	17.8
	OpenCV	36.11	38.44	37.64	13.89
	Diffusion	23.24	26.02	25.91	30.39
FRGC	MIPGAN-II	12.33	14.3	16.2	20.97
	StyleGAN2	7.18	8.71	9.74	14.42
	FaceMorpher	17.5	18.87	20.39	9.07
	OpenCV	17.02	18.47	19.91	6.6
	Diffusion	9.7	10.71	12.27	15.62

to the image after an initial face extraction and cropping performed by MTCNN, such that a margin size  $N$  adds back at most  $N$  pixels to the cropped image in both dimensions. Therefore, the larger  $N$  is the less tightly cropped the image passed to the FR system is. Table 4 illustrates the impact of the margin size on the APCER metric on the FaceNet FR system. Generally, as the margin size increases, the performance of Landmark-based attacks decreases while the performance of deep learning-based attacks increases. As illustrated in Figure 4, the Landmark-based attacks have noticeable artifacts outside the central face region; conversely, the deep learning-based morphs have less artifacts in the outside regions and generally look more realistic to a human observer. This observation aligns with the visual fidelity results from Table 1. Therefore, a MAD algorithm or FR system which uses less tightly cropped faces would be more resilient against attacks with visual artifacts outside the core face region.

### 5.2.2 General Remarks on the Vulnerability Study

The poor performance of the StyleGAN2 attack could be attributed to the darkening of images with light backgrounds, see Figure 4, and the aliasing effects latent to the StyleGAN2 generation pipeline, which was addressed by Karras *et al.* [56]. Moreover, the structure of the StyleGAN2 latent space can make exploration in the space difficult, which could possibly explain the poor performance in attacking the FR system compared to other attacks. MIPGAN-II, on the other hand, likely avoids these pitfalls due to its explicit latent optimization process for deceiving an FR system. The Diffusion attack utilizes an entirely different latent representation scheme, which seems to yield an advantage in the task of generating morphed faces. The pre-processing pipeline of the FR system seems to mostly mitigate the artifacts latent to the

TABLE 5: Ablation study on the impact of morphing attack on validation accuracy.

Dataset	Training Attack					Validation Attack				
	Diffusion	FaceMorpher	MIPGAN-II	OpenCV	StyleGAN2	Diffusion	FaceMorpher	MIPGAN-II	OpenCV	StyleGAN2
FERET	✗	✓	✓	✓	✓	72.73	99.23	100	99.95	99.33
	✓	✗	✓	✓	✓	99.9	76.39	100	99.85	99.64
	✓	✓	✗	✓	✓	99.69	99.38	100	99.95	99.54
	✓	✓	✓	✗	✓	99.74	99.48	100	99.74	99.43
	✓	✓	✓	✓	✗	99.74	98.56	99.9	99.74	87.89
FRGC	✗	✓	✓	✓	✓	75.89	99.98	99.97	99.9	99.93
	✓	✗	✓	✓	✓	99.95	99.48	100	99.9	99.95
	✓	✓	✗	✓	✓	99.83	99.85	99.82	99.8	99.85
	✓	✓	✓	✗	✓	99.93	100	100	99.23	99.93
	✓	✓	✓	✓	✗	99.93	99.93	99.94	99.88	97.83
FRL	✗	✓	✓	✓	✓	13.96	99.58	99.32	99.65	99.65
	✓	✗	✓	✓	✓	99.23	99.09	98.91	99.37	99.44
	✓	✓	✗	✓	✓	99.09	98.95	98.24	99.02	99.09
	✓	✓	✓	✗	✓	99.51	99.44	99.19	99.16	99.58
	✓	✓	✓	✓	✗	99.93	99.86	99.86	99.93	95.02

Landmark-based attacks; however, such artifacts could easily be detected by a human observer. Lastly, the state-of-the-art ArcFace FR system seems especially vulnerable to the Diffusion attack.

### 5.3 Detectability of Morphing Attacks

The performance of the proposed attack is further evaluated by the ability of Morphing Attack Detection (MAD) algorithms trained against other attacks to detect an unseen attack. To quantitatively assess the detectability of a particular morphing attack, a SE-ResNeXt101-32x4d model pre-trained on ImageNet [57] by NVIDIA is trained to detect morphing attacks. SE-ResNeXt101-32x4d is a state-of-the-art image recognition model based on the ResNeXt101-32x4d model [58] with the addition of the Squeeze-and-Excitation architecture [45]. For all experiments a 5-fold stratified  $k$ -fold cross validation strategy is employed, thus preserving the class balance between morphed and bona fide images in each fold. The model is fine-tuned on a collection of morphing attacks for 5 training epochs using exponential learning rate scheduler with differential learning rates in order to mitigate overfitting of the model.

#### 5.3.1 Ablation Study

To study the impact of a particular morphing attack on the ability of a MAD algorithm to detect morphing attacks, an ablation study was conducted where the SE-ResNeXt101-32x4d model was trained on all the morphing attacks except for one holdout. Table 5 shows the validation accuracy of each morphing attack when different morphing attacks were withheld from the training process. Due to the similar natures between the OpenCV and FaceMorpher attacks, the absence of one of these attacks does not greatly impact the validation accuracy. Interestingly, the absence of MIPGAN-II does not significantly change the validation accuracy of the attacks; however, the omission of StyleGAN2 during training does decrease the performance of the StyleGAN2 during validation, despite the presence of MIPGAN-II. Notably, the Diffusion attack is very difficult to detect as a novel attack, which can at least be partially attributed to its unique morph generation process in contrast with other morphing attacks.

#### 5.3.2 A Metric For Relative Strength

In this section we introduce a new metric to measure the strength of one morphing attack relative to another. We say a morph  $\alpha$

is “strong” relative to a morph  $\beta$  if the following conditions are satisfied:

- 1) It is easy to detect  $\beta$  when a detector is trained on  $\alpha$ , *i.e.*, high transferability.
- 2) It is hard to detect  $\alpha$  when a detector is trained on  $\beta$ , *i.e.*, low detectability.

Additionally, the relative strength metric,  $\Delta(\alpha||\beta)$ , should be positive when  $\alpha$  is stronger than  $\beta$  and negative when  $\alpha$  is weaker. A relative strength of 0 would denote that the two morphing attacks are equally strong.

As some of the morphing attacks are not deterministic but probabilistic, we choose to represent a morphing attack  $\alpha$  by the random variable  $X^\alpha : \Omega \rightarrow \mathcal{X}$  such that  $P(X^\alpha|x_a, x_b)$  denotes the distribution of morphs generated from images  $x_a, x_b$ . Moreover, we suppose there exists a detector  $f^\alpha : \mathcal{X} \rightarrow \{0, 1\}$  trained to distinguish between bona fide presentations and morphed presentations generated by  $\alpha$ ; wherein 0 denotes a bona fide presentation and 1 denotes a morphed presentation. The transferability of a morphing attack  $\alpha$  to  $\beta$  is defined as the probability the detector  $f^\alpha$  is able to detect the attack  $\beta$  given the probability  $f^\alpha$  detects  $\alpha$ , *i.e.*,  $T(\alpha, \beta) = P(f^\alpha(X^\beta) = 1 | f^\alpha(X^\alpha) = 1)$ . This metric can be represented as a ratio of expectations taken over the pairs of component bona fide images:

$$T(\alpha, \beta) = \frac{P(f^\alpha(X^\beta) = 1, f^\alpha(X^\alpha) = 1)}{P(f^\alpha(X^\alpha) = 1)} = \frac{\mathbb{E}_{x_a, x_b}[P(f^\alpha(X^\beta) = 1, f^\alpha(X^\alpha) = 1 | x_a, x_b)]}{\mathbb{E}_{x_a, x_b}[P(f^\alpha(X^\alpha) = 1 | x_a, x_b)]} \quad (12)$$

Let  $\{x_i^\alpha\}_{i=1}^N$  denote a collection of  $N$  samples drawn from  $P(X^\alpha|x_a, x_b)$  such that  $x_i^\alpha$  denotes the morph generated from  $i$ -th pair of bona fide identities  $(a_i, b_i)$ , and likewise for  $\beta$ . Then the metric in Equation (12) can be closely approximated by

$$T(\alpha, \beta) \approx \frac{\sum_{i=1}^N [f^\alpha(x_i^\beta) = 1 \wedge f^\alpha(x_i^\alpha) = 1]}{\sum_{i=1}^N [f^\alpha(x_i^\alpha) = 1]} \quad (13)$$

*i.e.*, the number of morphs from both  $\alpha$  and  $\beta$  detected over the number of morphs detected from  $\alpha$ .

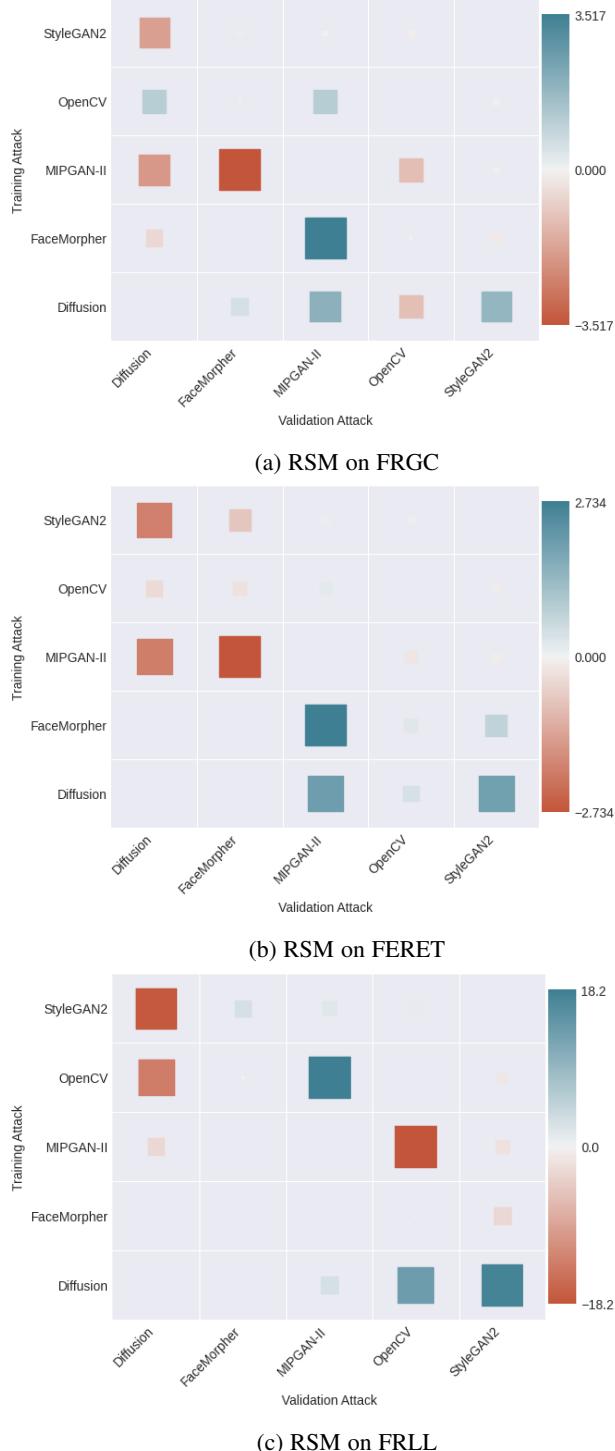


Fig. 6: Blue indicates higher strength and red indicates weak strength.

The relative strength metric (RSM) from  $\alpha$  to  $\beta$  is defined as the log ratio of the transferability metrics between the two morphing attacks:

$$\Delta(\alpha\|\beta) = \log\left(\frac{T(\alpha, \beta)}{T(\beta, \alpha)}\right) \quad (14)$$

The log of the ratio is chosen such that the RSM takes positive values when  $\alpha$  is “stronger” than  $\beta$  and negative values when

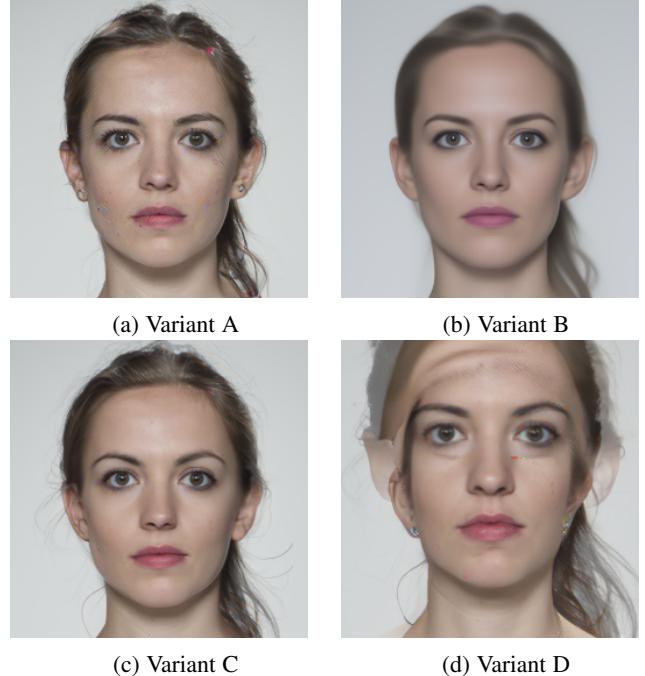


Fig. 7: Morphed image generated by different Diffusion attack variants on FRLL.

weaker—with a value of zero denoting equal strength. Additionally, there is an antisymmetry such that  $\Delta(\alpha\|\beta) = -\Delta(\beta\|\alpha)$ .

In contrast to the ablation study, the SE-ResNeXt101-32x4d model is only trained on a single attack per  $k$ -fold. The RSM is calculated between all attacks with the results shown in Figure 6. It can be observed that the RSM between the Landmark-based morphs and the RSM between the StyleGAN-based morphs is very small. As these attacks have similar morph generation pipelines, it makes sense that the transferability between the attacks is near identical. In general, the Landmark-based attacks seem to be stronger than the StyleGAN-based attacks, in particular the FaceMorpher attack. The MIPGAN-II attack is generally weaker than the other attacks. Overall, the Diffusion attack is the least detectable among the attacks along with generally being the strongest attack across the three datasets.

The results from Figure 6 corroborate with the results from Table 5, demonstrating the difficulty in detecting Diffusion attacks. From the perspective of training a MAD system, including samples from the FaceMorpher, StyleGAN, and Diffusion attacks would greatly increase the ability for the system to detect unknown attacks. Additionally, Table 5 and Figure 6 demonstrates a particular vulnerability existing MAD systems may have towards the emerging Diffusion attack.

#### 5.4 Study of the Diffusion-based Morphing Process

The diffusion morphing algorithm leverages both a stochastic and semantic representation of an image. While the semantic representation contains many of the key “identifying” features, the stochastic representation contains many of the details necessary for high visual fidelity. Due to the importance of the stochastic code for high fidelity, we investigated several methods for finding the morphed stochastic latent code,  $\mathbf{x}_T^{(ab)}$ . The first variant, variant A, is the baseline implementation with  $\ell_Z$  using linear interpolation,  $\ell_X$  using spherical linear interpolation, and

TABLE 6: MMPMR at FMR = 0.1% across different configurations. Higher is better.  $\dagger$  indicates our default choices.  $\ddagger$  the geometric mean.

Variant	$\ell_{\mathcal{X}}$	$\xi(x, y)$	FRLL			FRGC			FERET			
			FaceNet	VGGFace2	ArcFace	FaceNet	VGGFace2	ArcFace	FaceNet	VGGFace2	ArcFace	Mean $\ddagger$
A	slerp	$x, y \mapsto x$	<b>32.97</b>	34.71	<b>88.27</b>	<b>3.2</b>	<b>9.59</b>	44.83	<b>7.17</b>	<b>11.54</b>	58.69	<b>20.62</b>
B	lerp	$x, y \mapsto x$	10.81	11	68.94	1.17	2.17	38.85	2.33	4.69	35.41	8.79
C $\dagger$	slerp	$x, y \mapsto \frac{1}{2}(x + y)$	28.14	<b>35.37</b>	88.09	2.68	8.47	<b>46.74</b>	6.47	13.03	<b>59.75</b>	19.80
D	slerp	$x, y \mapsto \text{OpenCV}(x, y)$	9.14	9.34	29.84	0	1.37	1.85	0.14	1.42	1.65	0.0

TABLE 7: FID across different configurations. Lower is better.  $\dagger$  indicates our default choices.

Variant	$\ell_{\mathcal{X}}$	$\xi(x, y)$	FRLL	FRGC	FERET
A	slerp	$x, y \mapsto x$	48.13	<b>52.97</b>	55.66
B	lerp	$x, y \mapsto x$	82.05	119.33	97.75
C $\dagger$	slerp	$x, y \mapsto \frac{1}{2}(x + y)$	<b>42.63</b>	64.16	<b>50.45</b>
D	slerp	$x, y \mapsto \text{OpenCV}(x, y)$	93.85	84.51	108.49

$\xi$  does not perform any “pre-morphing”. Conversely, in variant B the stochastic codes are interpolated via linear interpolation. In variants C and D, instead of using the original image to calculate the stochastic code, the function  $\xi$  is used to construct the “pre-morph” passed to the stochastic encoder. Specifically, in variant C the two images are averaged pixel-wise and presented to the stochastic encoder; in contrast, in variant D the OpenCV morph is presented to the stochastic encoder. For each variant, we generated the same number of morphs across the three datasets as shown in Section 4.2, resulting over 10,000 new morphs across four variants.

In Table 7 the FID is calculated between the generated morphs and the bona fide samples for each particular dataset. Variant C generally presents the lowest FID score, closely followed by variant A. Both variants B and D exhibit clear degradation in performance when compared to variants A and C. Furthermore, the FID score seems to correlate well with human assessment of the generated samples, see Figure 7. Noticeably, the linear interpolation in variant B results in an overly smoothed face and generally darker image, greatly degrading visual fidelity. Variant D has prominent visual artifacts, similar to the artifacts found in the OpenCV morphs. Moreover, the poor performance seems to be aided by an issue of differing alignment strategies between the OpenCV and diffusion pipeline.

Notably, variant C often removes many of the high frequency artifacts found in variant A. This is likely due to the difficulty in smoothly interpolating between points in the stochastic latent space, in contrast with the semantic latent space. As such, variant C, which performs a pixel-wise average of the two source images before using the stochastic encoder, seems to greatly improve the ability to smoothly interpolate between different stochastic representations. This appears to be the primary reason variant C has a generally lower FID when compared to variant A. Both Figure 7 and Table 7 demonstrate the large importance that the stochastic code plays in creating high fidelity morphed images. Due to the high fidelity exhibited by variant C, this particular diffusion process was used in evaluation against other morphing attacks.

The MMPMR metric is calculated for each variant in Table 6. Variant A is slightly stronger than variant C, with variants B and D falling far behind, likely due to the high number of visual distortions. These results stand in contrast to the assessment of

visual fidelity wherein variant C outperforms variant A. This, again, illustrates a trade-off between visual fidelity and ability to fool the FR system; however, in this case the trade-off effectiveness against the FR system is relatively small in comparison to the gains in visual fidelity. Due to its excellent visual fidelity and strong MMPMR results, variant C was still chosen to be the default configuration for the Diffusion attack.

## 6 CONCLUSIONS

By addressing some of the key limitations of prior deep-learning based morphing attacks, namely, the trade-off between visual fidelity and effectiveness against FR systems, we have proposed a novel morphing attack using Diffusion-based methods for the generative process. The proposed attack consistently generates realistic morphed images with high visual fidelity, while also being able to strongly threaten FR systems. To evaluate the attack potential of the proposed method, we evaluated the vulnerability of three FR systems over three distinct datasets across four variants of the Diffusion morph, with the strongest variant achieving state-of-the-art performance. We also studied the impact of the pre-processing pipeline on the vulnerability of an FR system to morphing attacks. A novel metric to assess the strength of one morphing attack relative to another has been introduced. Moreover, the proposed attack was evaluated by its detection performance against a state-of-the-art MAD system. The Diffusion attack was shown to be very difficult to detect if not specifically trained against, showing the proposed attack can greatly threaten existing FR systems. Overall, the images generated by the Diffusion attack possess high visual fidelity, can deceive state-of-the-art FR systems, and are difficult for MAD mechanisms to detect. It is our belief that developing a stronger face morphing attack is essential to the design and testing of stronger MAD methods. As far as future work is concerned, the proposed method could be extended to higher resolutions using diffusion-based super resolution methods or latent diffusion methods.

## ACKNOWLEDGMENTS

This material is based upon work supported by the Center for Identification Technology Research and National Science Foundation under Grant #1650503.

## REFERENCES

- [1] L. J. Spreeuwers, A. J. Hendrikse, and K. J. Gerritsen, “Evaluation of automatic face recognition for automatic border control on actual data recorded of travellers at schiphol airport,” in *Proc. of the Int’l Conf. of Biometrics Special Interest Group (BIOSIG)*, 2012, pp. 1–6.
- [2] M. Ferrara, A. Franco, and D. Maltoni, *On the Effects of Image Alterations on Face Recognition Accuracy*. Cham: Springer International Publishing, 2016, pp. 195–222. [Online]. Available: [https://doi.org/10.1007/978-3-319-28501-6\\_9](https://doi.org/10.1007/978-3-319-28501-6_9)

[3] D. J. Robertson, R. S. S. Kramer, and A. M. Burton, "Fraudulent id using face morphs: Experiments on human and automatic recognition," *PLOS ONE*, vol. 12, no. 3, pp. 1–12, 03 2017. [Online]. Available: <https://doi.org/10.1371/journal.pone.0173319>

[4] R. Raghavendra, K. B. Raja, and C. Busch, "Detecting morphed face images," in *IEEE 8th Int'l Conf. on Biometrics Theory, Applications and Systems (BTAS)*, 2016, pp. 1–7.

[5] J. T. A. Andrews, T. Tanay, and L. D. Griffin, "Multiple-identity image attacks against face-based identity verification," *CoRR*, vol. abs/1906.08507, 2019. [Online]. Available: <http://arxiv.org/abs/1906.08507>

[6] R. Raghavendra, K. Raja, S. Venkatesh, and C. Busch, "Face morphing versus face averaging: Vulnerability and detection," *IEEE Int'l Joint Conf. on Biometrics (IJCB)*, pp. 555–563, 2017.

[7] C. Burt, "Face morphing threat to biometric identity credentials' trustworthiness a growing problem," *BiometricUpdate.com*.

[8] S. Venkatesh, R. Ramachandra, K. Raja, L. Spreeuwers, R. Veldhuis, and C. Busch, "Morphed face detection based on deep color residual noise," in *the 9th Int'l Conf. on Image Pr. Theory, Tools and Applications (IPTA)*, 2019, pp. 1–6.

[9] U. Scherhag, R. Raghavendra, K. B. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch, "On the vulnerability of face recognition systems towards morphed face attacks," in *2017 5th International Workshop on Biometrics and Forensics (IWBF)*, 2017, pp. 1–6.

[10] U. Scherhag, L. Debiasi, C. Rathgeb, C. Busch, and A. Uhl, "Detection of face morphing attacks based on prnu analysis," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 1, no. 4, pp. 302–317, 2019.

[11] Z. Blasingame and C. Liu, "Leveraging adversarial learning for the detection of morphing attacks," *2021 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–8, 2021.

[12] S. Venkatesh, R. Ramachandra, K. Raja, and C. Busch, "Face Morphing Attack Generation & Detection: A Comprehensive Survey," *arXiv e-prints arXiv:2011.02045*, Nov. 2020.

[13] T. Ojala, M. Pietikäinen, and D. Harwood, "A comparative study of texture measures with classification based on featured distributions," *Pattern Recognit.*, vol. 29, pp. 51–59, 1996.

[14] J. Kannala and E. Rahtu, "Bsif: Binarized statistical image features," in *Proceedings of the 21st Int'l Conf. on Pattern Recognition (ICPR2012)*, 2012, pp. 1363–1366.

[15] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert, "Detection of face morphing attacks by deep learning," in *Digital Forensics and Watermarking*, C. Kraetzer, Y.-Q. Shi, J. Dittmann, and H. J. Kim, Eds., 2017, pp. 107–120.

[16] U. Scherhag, C. Rathgeb, J. Merkle, and C. Busch, "Deep face representations for differential morphing attack detection," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3625–3639, 2020.

[17] M. Ngan, P. Grother, K. Hanaoka, and J. Kuo, "Face recognition vendor test (frvt) part 4: Morph - performance of automated face morph detection," 2020-03-06 2020.

[18] E. Sarkar, P. Korshunov, L. Colbois, and S. Marcel, "Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks," *ArXiv*, vol. abs/2012.05344, 2020.

[19] K. O'Haire, S. Soleymani, B. Chaudhary, P. Aghdaie, J. Dawson, and N. M. Nasrabadi, "Adversarially perturbed wavelet-based morphed face generation," in *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, 2021, pp. 01–05.

[20] 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 Int'l Conf. on Biometrics Theory, Applications and Systems (BTAS)*, 2018, pp. 1–10.

[21] S. Venkatesh, H. Zhang, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Can gan generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection," in *2020 8th International Workshop on Biometrics and Forensics (IWBF)*, 2020, pp. 1–6.

[22] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 8107–8116.

[23] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Mipgan—generating strong and high quality morphing attacks using identity prior driven gan," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 3, pp. 365–383, 2021.

[24] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 8780–8794. [Online]. Available: <https://proceedings.neurips.cc/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf>

[25] L. DeBruine and B. Jones, "Face research lab london set."

[26] S. Venkatesh, H. Zhang, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Can gan generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection," 2020.

[27] E. Sarkar, P. Korshunov, L. Colbois, and S. Marcel, "Are gan-based morphs threatening face recognition?" in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 2959–2963.

[28] A. Quek, "Facemorpher," 2019.

[29] S. Milborrow and F. Nicolls, "Active shape models with sift descriptors and mars," in *2014 International Conference on Computer Vision Theory and Applications (VISAPP)*, vol. 2, 2014, pp. 380–387.

[30] U. Scherhag, J. Kunze, C. Rathgeb, and C. Busch, "Face morph detection for unknown morphing algorithms and image sources: a multi-scale block local binary pattern fusion approach," *IET Biometrics*, vol. 9, pp. 278–289, 11 2020.

[31] D. E. King, "Dlib-ml: A machine learning toolkit," *J. Mach. Learn. Res.*, vol. 10, p. 1755–1758, Dec. 2009.

[32] L. Colbois and S. Marcel, "On the detection of morphing attacks generated by gans," in *2022 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2022, pp. 1–5.

[33] D. Roich, R. Mokady, A. H. Bermano, and D. Cohen-Or, "Pivotal tuning for latent-based editing of real images," *ACM Transactions on Graphics (TOG)*, 2021.

[34] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4396–4405.

[35] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=B1xsqj09Fm>

[36] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>

[37] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.

[38] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=St1giarCHLP>

[39] K. Preechakul, N. Chatthee, S. Wizadwongsu, and S. Suwajanakorn, "Diffusion autoencoders: Toward a meaningful and decodable representation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 10619–10629.

[40] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.

[41] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, 2018, pp. 67–74.

[42] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 815–823.

[43] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.

[44] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *BMVC*, 2015.

[45] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141.

[46] I. C. Duta, L. Liu, F. Zhu, and L. Shao, "Improved residual networks for image and video recognition," *arXiv preprint arXiv:2004.04989*, 2020.

[47] X. An, X. Zhu, Y. Gao, Y. Xiao, Y. Zhao, Z. Feng, L. Wu, B. Qin, M. Zhang, D. Zhang, and Y. Fu, "Partial fc: Training 10 million identities on a single machine," in *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 2021, pp. 1445–1449.

- [48] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [49] P. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, "The feret database and evaluation procedure for face-recognition algorithms," *Image Vis. Comput.*, vol. 16, pp. 295–306, 1998.
- [50] P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. 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)*, vol. 1, 2005, pp. 947–954 vol. 1.
- [51] M. Lucic, K. Kurach, M. Michalski, O. Bousquet, and S. Gelly, "Are gans created equal? a large-scale study," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, ser. NIPS'18. Red Hook, NY, USA: Curran Associates Inc., 2018, p. 698–707.
- [52] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [53] M. Seitzer, "pytorch-fid: FID Score for PyTorch," <https://github.com/mseitzer/pytorch-fid>, August 2020, version 0.2.1.
- [54] "Best practice technical guidelines for automated border control abc systems," 2015.
- [55] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. J. Veldhuis, L. Spreeuwers, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Ramachandra, 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)*, 2017, pp. 1–7.
- [56] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, "Alias-free generative adversarial networks," in *Proc. NeurIPS*, 2021.
- [57] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE conf. on computer vision and pattern recognition*, 2009, pp. 248–255.
- [58] S. Xie, R. B. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5987–5995, 2017.

**Zander Blasingame** received the B.Sc. degree in Computer Engineering from Clarkson University in 2018. He is currently pursuing a Ph.D. degree in Electrical and Computer Engineering with Clarkson University, Potsdam, New York. His research interests in biometrics include face morphing, deep fakes, latent representations of faces, and latent manipulations of faces.

**Chen Liu** [SM'21] received his Ph.D. in Electrical and Computer Engineering from the University of California, Irvine in 2008. Currently he is an associate professor in the Department of Electrical and Computer Engineering at Clarkson University, Potsdam, New York. His research interests in biometrics include face in video recognition, face morphing, and hardware acceleration for biometric workloads.