

SelfMAD: Enhancing Generalization and Robustness in Morphing Attack Detection via Self-Supervised Learning

Marija Ivanovska¹, Leon Todorov², Naser Damer³, Deepak Kumar Jain⁴, Peter Peer², Vitomir Štruc¹

¹ Faculty of Electrical Engineering, University in Ljubljana, Slovenia

² Faculty of Computer and Information Science, University in Ljubljana, Slovenia

³ Dalian University of Technology, China

⁴ Fraunhofer Institute for Computer Graphics Research, Germany

Abstract—With the continuous advancement of generative models, face morphing attacks have become a significant challenge for existing face verification systems due to their potential use in identity fraud and other malicious activities. Contemporary Morphing Attack Detection (MAD) approaches frequently rely on supervised, discriminative models trained on examples of bona fide and morphed images. These models typically perform well with morphs generated with techniques seen during training, but often lead to sub-optimal performance when subjected to novel unseen morphing techniques. While unsupervised models have been shown to perform better in terms of generalizability, they typically result in higher error rates, as they struggle to effectively capture features of subtle artifacts. To address these shortcomings, we present SelfMAD, a novel self-supervised approach that simulates general morphing attack artifacts, allowing classifiers to learn generic and robust decision boundaries without overfitting to the specific artifacts induced by particular face morphing methods. Through extensive experiments on widely used datasets, we demonstrate that SelfMAD significantly outperforms current state-of-the-art MADs, reducing the detection error by more than 64% in terms of EER when compared to the strongest unsupervised competitor, and by more than 66%, when compared to the best performing discriminative MAD model, tested in cross-morph settings. The source code for SelfMAD is available at <https://github.com/LeonTodorov/SelfMAD>.

I. INTRODUCTION

Automatic face recognition systems (FRSs) [22] are commonly employed to verify an individual’s identity by matching their face image with the corresponding data stored in the system’s database. Although these systems are widely used and generally very accurate, they are vulnerable to certain types of attacks representing manipulated data. A notable example are *face morphing attacks* [19], [52], [54], created by blending the facial features of two or more individuals. The resulting morphed image can then be used to falsely authenticate any person whose facial attributes were utilized in the morphing process.

In the age of big data and major advancements in generative models, the widespread availability of open-source morphing techniques have made it nearly effortless to create realistic, high-quality morphed face images [11], [71]. The automatic detection of face morphing attacks is therefore

The research presented in this work was supported in parts by the ARRS Research Programmes P2-0250(B) “Metrology and Biometric Systems”, P2-0214 “Computer Vision”, and the ARRS Research Project J2-50065 “DeepFake DAD”.

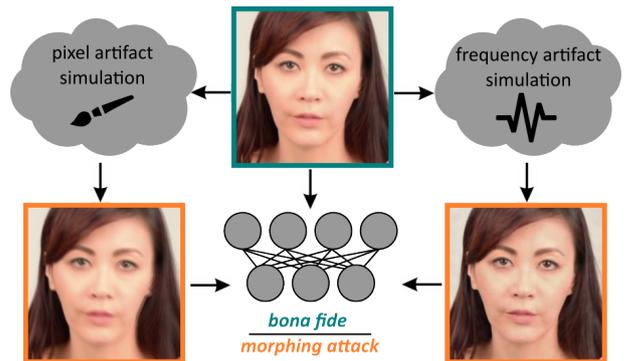


Fig. 1. We propose SelfMAD, a Self-supervised Morphing Attack Detection method, that learns to detect morphed faces by replicating common artifacts of various widely used morphing techniques. Pixel space manipulations simulate artifacts typical for image-level morphing techniques, while frequency space manipulations reproduce the fingerprints of latent-space morphing techniques. Self-MAD is robust and generalizes effectively, without overfitting to specific face morphing attack examples.

critical for preventing illegal activities [1] such as identity theft, personal document frauds, social engineering attacks etc. In recent years, the threat posed by morphed image attacks has been predominantly addressed with the development of powerful *morphing attack detection (MAD)* methods. In most cases, they rely on *supervised, contrastive learning*, optimizing models to differentiate between bona fide images and examples of known morphing attacks [5], [6], [48]. Although such an approach tends to be highly accurate when tested on attacks encountered during training, these methods (i) fail to detect morphs generated by unfamiliar or novel attacking techniques [11]; and (ii) their performance usually declines, when applied to data from unfamiliar sources, due to domain shifts [18].

To overcome the limitation in the generalization capabilities of supervised MADs, some researchers have investigated the use of *unsupervised, one-class models* [12], [26]. Instead of relying on contrastive learning, their objective is to capture only characteristics of bona fide training images and detect morphing attacks as out-of-distribution samples. While one-class models are expected to generalize better to previously unseen morphs, they frequently: (i) mistakenly classify legitimate samples as attacks, particularly when the bona fide training data does not fully capture the diversity of real-world scenarios [18]; (ii) struggle with high sensitivity

to the quality and quantity of the bona fide images used during training, reducing the models’ robustness to varying conditions [28]; and (iii) fail to learn distinctive bona fide features that are absent in manipulated data, which hinders the MAD models’ ability to identify morphing attacks as samples that do not fit the modeled patterns [12].

Recently, *self-supervised learning* has emerged as a promising alternative to the fully unsupervised training of one-class models. Unlike unsupervised methods, self-supervised learning typically utilizes synthetically generated samples or perturbations, to enhance the model’s ability to detect subtle or complex anomalies in various forms of manipulated data. In the context of image manipulation detection, self-supervised models have been successfully applied to tasks such as anomaly detection in industrial images [30], [33], [68], detection of adversarial attacks [16], [35], [40], presentation attack detection [2], [7], [29], deep-fake detection [34], [36], [59], detection of AI-synthesized images [8], [37], [64] etc. Despite these advancements, the full potential of self-supervised learning methods applied to the task of morphing attack detection has yet to be explored.

In this paper, we address the challenges of both supervised and unsupervised approaches by framing morphing attack detection as a self-supervised task. As a result, we make the following main contributions in this work:

- We propose **SelfMAD**, a novel *Self-supervised Morphing Attack Detection* method, designed to detect attacks by simulating common face morphing artifacts present across various types of manipulations (Fig. 1).
- Through extensive experiments, we demonstrate that our approach enhances the detection generalization to previously unseen morphing attacks, reducing the risk of overfitting to specific morphing techniques.
- We perform a rigorous comparison study against widely used and highly competitive MAD models, showing that SelfMAD offers a robust solution, effectively balancing the strengths of both supervised and unsupervised methods.
- We conduct an in-depth ablation study to assess the impact of SelfMAD components.

II. RELATED WORK

In this section, we first discuss different widely used techniques for generation of face morphing attacks. Next, we present a brief overview of the studies on morphing attack detection (MAD). Finally, we review self-supervised learning (SSL) methods applied for detection of out-of-distribution samples, providing a background of our research work. For a more in-depth coverage of these topics, readers are referred to some of the excellent surveys [24], [54], [61].

Morphing Attack Generation. Face morphing attack generation techniques are generally categorized into traditional and deep learning-based methods. Traditional approaches typically operate in the pixel space and involve three main steps: aligning corresponding face features from two images, warping them to match geometrically, and blending the warped images to merge color values [55]. The alignment relies on detected facial landmarks, and different methods may use various warping techniques [23], [39], [50],

[53]. However, this process can introduce misaligned pixels, leading to artifacts and ghost-like images that are often noticeable. To address these issues, post-processing steps such as image smoothing, sharpening, edge correction, and histogram equalization are commonly applied to enhance image quality and reduce artifacts [57], [65].

With the advent of advanced deep learning-based generative models, recent approaches have significantly improved the quality of morphed face images compared to traditional landmark-based methods. These modern techniques typically involve embedding two face images into the latent space of a generative model, performing vector interpolation to create a morphed image, and then decoding this vector back into pixel space. While Generative Adversarial Networks (GANs) are commonly used for this purpose [14], [62], [70], newer methods also incorporate diffusion-based networks to achieve even greater image quality and realism [3], [4], [11]. Despite their high level of realism and difficulty to detect, morphs produced by these generative models still exhibit detectable irregularities in texture or frequency patterns, which can reveal their manipulated nature [9].

Morphing Attack Detection. Existing morphing attack detection (MAD) models are in general categorized into single-image (S-MAD) and differential (D-MAD) approaches. S-MAD models analyze face morphs individually, without comparing them to other images, whereas D-MAD models compare manipulated samples to a reference image. D-MADs are typically highly accurate in closed-set scenarios, while S-MADs are primarily employed to detect attacks when there is no prior knowledge of the subjects’ identities. In this section, we focus our literature review exclusively on S-MADs, as they are most closely related to our work.

Regardless of the face morphing technique employed, the resulting morphs typically exhibit image irregularities such as noise, pixel discontinuities, distortions, spectrum discrepancies, and other visual artifacts. Early MAD approaches aimed to detect these irregularities using hand-crafted techniques, including photo-response non-uniformity (PRNU) noise analysis [51], reflection analysis [56], or texture-based descriptors like LBP [42], LPQ [43], or SURF [38]. While these methods produced promising results, their ability to generalize across different scenarios was limited [12].

More sophisticated MADs take advantage of the capabilities of data-driven, deep learning algorithms [25]. Raghavendra *et al.* [47], for instance, were amongst the first proposing pretrained deep models as a supervised approach to the detection of morphing attacks. In their work, morphs were detected with a simple, fully-connected binary classifier, fed with fused VGG19 and AlexNet features, pretrained on ImageNet. Similarly, Wandzik *et al.* [63], achieved high detection performance with features extracted from a general-purpose face recognition systems (FRSs), fed to an SVM, while another study of Ramachandra *et al.* [48] utilized Inception models for the same purpose. Damer *et al.* [15] on the other hand argued that pixel-wise supervision, where each pixel is classified as a bona fide or a morphing attack, is superior, when used in addition to the binary, image-level

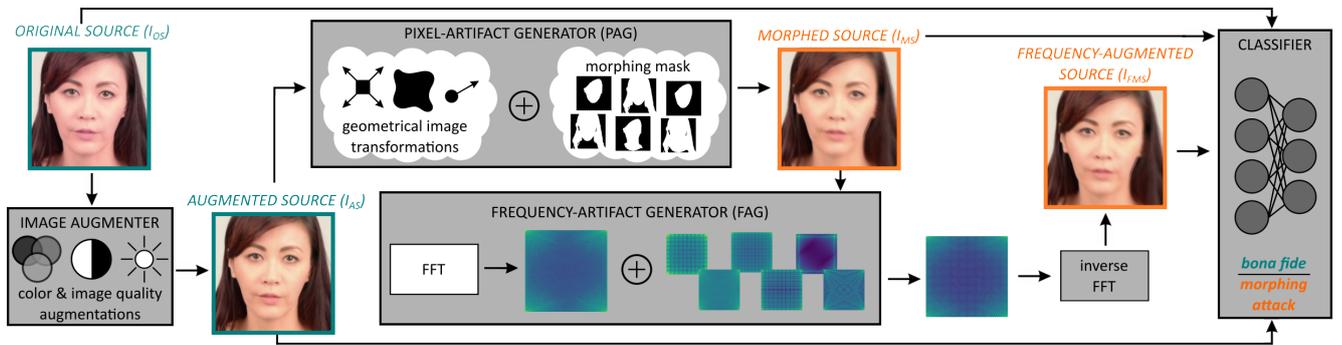


Fig. 2. **Overview of Self-MAD**, a self-supervised morphing attack detection method that learns to detect morphed faces using a set of bona fide images and simulated morphing attacks. The model consists of four key components: *i) pixel augmenter*, to simulate subtle visual variations in real-world images; *ii) pixel artifact generator* that mimics artifacts typical for image-level morphing methods; *iii) frequency artifact generator* which reproduces frequency fingerprints associated with advanced, latent-level morphing techniques; and *iv) classifier* differentiate between genuine and manipulated samples.

objective. MixFaceNet [5] by Boutros *et al.*, another highly efficient deep learning architecture inspired by mixed depth-wise convolutions, demonstrates even better results [13], by capturing different levels of attack cues through differently-sized convolutional kernels. A more complex approach is proposed by Neto *et al.*, who incorporate a regularization term during the training of their model OrthoMAD, whose goal is to disentangle identity features for more robust morphing attack detection. The follow-up method, IDistill further improves MAD results by introducing more efficient feature disentanglement and adding interpretability. All these supervised approaches however typically experience a significant decline in their performance, when applied to images that do not fit the training distribution.

In some more recent studies, researchers have tried to improve the generalization capabilities of the MAD techniques, by proposing unsupervised, one-class learning models trained exclusively on bona fide samples. Different from the supervised techniques discussed above, Damer *et al.* [12], for example, were among the first to achieve significant performance generalization on unseen attacks with two different one-class methods, i.e. a one-class support vector machine (OCSVM) and an isolation forest (ISF). Comparable generalization capabilities were later demonstrated in [26], where Ibsen *et al.* explored the use of a Gaussian Mixture Model (GMM), a Variational Autoencoder (VAE) and Single-Objective Generative Adversarial Active Learning (SO-GAAL) in addition to an OCSVM. A more advanced approach was proposed by Fang *et al.* [18] who enhance an unsupervised convolutional autoencoder with self-paced learning (SPL). With this approach, the model neglects suspicious unlabeled training data, widening the reconstruction error gap between bona fide samples and morphing attacks. Ivanovska *et al.* [28] achieve further improvement in detection performance by deploying diffusion models to learn the distribution of bona-fide images. Different from reconstruction techniques, Fu *et al.* [21] measure the authenticity of images through image quality estimation. Nevertheless, learning distinctive image features in a one-class setting is still a challenging task, frequently leading to misclassification of realistic morphs with subtle inconsistencies.

Self-Supervised Learning for Anomaly Detection. Self-

supervised learning has recently emerged, as a promising alternative that tackles the challenges of out-of-distribution sample detection, by combining the strengths of both supervised and unsupervised approaches. Unlike fully supervised methods, which depend heavily on labeled data, and unsupervised one-class models that in contrast focus solely on bona fide samples, self-supervised techniques utilize automatically generated labels derived from the data itself. For instance, in industrial anomaly detection, self-supervised models simulate typical defects by synthesizing data representing anomalous samples [30], [33], [68]. In adversarial attack detection, Deb *et al.* adopt a similar approach by artificially creating challenging and diverse adversarial attacks in the image space [16], while Li *et al.* and Naseer *et al.* apply data manipulation in the latent space. These strategies have been so far successfully applied to various biometric tasks [2], [7], [29]. Recently, a notable success has been achieved in tasks related to deepfake detection. Li *et al.* and Shiohara *et al.* [34], [59] for example, propose simulating typical image inconsistencies in deepfakes, where an original face has been either replaced or enhanced. Differently, Wang *et al.* choose to simulate frequency artifacts instead of focusing on local image irregularities [64]. More sophisticated algorithms simulating frequency-based irregularities have also been proposed by Coccomini *et al.* [8] and Lu *et al.* [37] who demonstrate remarkable generalizability of this approach.

Motivated by the success of the self-supervised paradigm, in this paper we explore its capabilities in the context of MAD (morphing attack detection). Unlike existing MAD methods, we augment normal data by simulating typical *morphing irregularities* that manifest in both pixel and frequency domains. This newly generated data is then used to extract general and robust features for effective attack detection.

III. METHODOLOGY

We introduce **SelfMAD**, a self-supervised morphing attack detection method designed to recognize manipulated face images by searching for general inconsistencies that are common across various types of face morphs, and at the same time independent of specific face identities. The proposed method utilizes a proxy task, which simulates morphing inconsistencies through a three-stage image pre-processing pipeline that consists of: *(i) image augmen-*

tion, (ii) pixel-artifact generation, and (iii) frequency-artifact generation steps. The manipulated images generated by the pipeline, along with the original, unprocessed images, are then used for the training of a neural discriminator, which learns to distinguish between bona fide and altered samples, without using actual face morphs. Note that for simulating morphing inconsistencies, we use various transformations and frequency masks that have been shown to be suitable for modeling artifacts and subtle cues induced into facial images by various (GAN, diffusion, or pixel-level) image-manipulation techniques. [10], [32], [59], [67] A high-level overview of SelfMAD is shown in Fig. 2.

Image Augmenter. The first component in our pipeline focuses on augmenting input images by applying a series of transformations that alter the visual appearance of the data, without changing its underlying structure. The primary goal of this processing stage is to *simulate subtle (global) visual variations* that may occur in real-world face images. Specifically, given an input bona fide image I_{OS} (original source), the Image Augmenter applies a set of transformations ψ , to generate an augmented image I_{AS} (augmented source):

$$I_{AS} = \psi(I_{OS}), \quad (1)$$

where $\psi \in \{\text{RGBShift, HueSaturationValue, RandomBrightnessContrast, RandomDownScale, Sharpen}\}$ comprises five basic (global) image transformations. Color-related variations in real-world image are simulated by applying the first two functions in ψ , which introduce slight adjustments to the RGB and HSV values of the input sample I_{OS} . The third function additionally alters the brightness and the contrast of the image. To mimic real-world variations in image quality, the augmenter also applies compression at varying intensities, using either downsampling or sharpen operation that are implemented with the last two functions in ψ . The parameters for all five transformation functions are randomly selected within predefined ranges. After applying ψ , the content of the input image I_{OS} remains unchanged, ensuring that the transformed output image I_{AS} is still considered bona fide.

Pixel-Artifact Generator. The second image processing component of SelfMAD *introduces pixel space artifacts* to simulate image irregularities created by traditional, landmark-based morphing techniques (see section II). To faithfully replicate the image-level morphing of two non-identical faces, the pixel artifact generator performs pixel space blending of the original source image I_{OS} and the geometrically transformed version of I_{AS} from the previous SelfMAD step. Formally, the pixel-artifact generator first creates $I_{AS'}$ by transforming I_{AS} with set of functions ζ :

$$I_{AS'} = \zeta(I_{AS}), \quad (2)$$

where $\zeta \in \{\text{Translation, ElasticTransform, Scaling}\}$ comprises three key geometrical image transformations. The first function performs only image translation in a randomly chosen direction. Therefore, it mimics the imperfect alignment of corresponding face landmarks, typical for the morphing of two different faces. The second function,

the elastic transformation, is a smooth, non-linear image deformation that locally adjusts pixel positions to simulate variations in the shapes of individual face parts. This function aims to reflect the natural distinctive features of different real-world face structures. The third and last function applies image scaling, consequently adjusting the dimensions of facial features to capture potential size discrepancies. Finally, the geometrically transformed image $I_{AS'}$ is blended with the original source image I_{OS} using a blending mask M :

$$I_{MS} = I_{AS'} \odot a \cdot M + I_{OS} \odot (1 - a) \cdot M, \quad (3)$$

where I_{MS} is the resulting morphed image source and a is the blending factor, uniformly sampled from a set of predefined discrete values. The blending mask M is a binary image, that allows the algorithm to selectively apply pixel artifacts. When all values of M equal 1, the pixel-artifact generator performs a classical pixel-wise morphing of $I_{AS'}$ and I_{OS} . Alternatively, the pixel-artifact generator can randomly choose a mask M that represents a combination of two or more facial parts, segmented by a pretrained face parser. This strategy helps SelfMAD emphasize pixel irregularities in different face regions.

Frequency-Artifact Generator. In the third image processing stage, the pipeline *simulates frequency artifacts* commonly created by generative models that perform face morphing in their latent spaces, i.e. on face template level (see section II). These morphs often appear flawless in the image space, maintaining correct semantics and visual consistency. However, they typically contain so-called frequency fingerprints that are not present in pristine images. These fingerprints can either represent structured geometrical artifacts frequently linked to generative adversarial networks (GANs) or abnormal densities in the frequency spectrum, often associated with diffusion-based models. The frequency artifact generator models these inconsistencies by first creating a random structured geometrical pattern Φ , then calculating its Fast Fourier Transform F_{Φ} :

$$F_{\Phi} = FFT(\Phi), \quad (4)$$

where the type of the pattern Φ is uniformly chosen to represent one of the following: a symmetrical grid, an asymmetrical grid, a square checkerboard, a circular checkerboard, randomly distributed squares, a set of random lines or a set of random stripes. The magnitudes $|F_{\Phi}|$ are then separately superimposed on the magnitudes of the Fourier transform of the image from the earlier step, i.e. the morphed source image I_{MS} :

$$|F_{MS}|_t = (1 - k) \cdot |FFT(I_{MS})| \oplus k \cdot |F_{\Phi}|, \quad (5)$$

where k is a predetermined constant of the pixel-wise summation. The resulting magnitude spectra is then transformed back to the image space, by applying the inverse FFT:

$$I_{FMS} = InverseFFT(|F_{MS}|_t, \theta(F_{MS})), \quad (6)$$

where θ is the phase spectra of the Fourier transform F_{MS} .

Classifier. The final component of SelfMAD is a classifier \mathcal{D} , optimized by minimizing the binary cross-entropy loss:

$$\mathcal{L}_{BCE} = -[y \log(\mathcal{D}(I)) + (1 - y) \log(1 - \mathcal{D}(I))], \quad (7)$$

where I are the input images generated at different pipeline stages (I_{OS} , I_{AS} , I_{MS} , and I_{FMS}), y is the corresponding image label, which equals 0 for bona fide images I_{OS} and I_{AS} and 1 for simulated morphing attacks I_{MS} , and I_{FMS} , while $\mathcal{D}(I)$ is the predicted probability of I being a bona fide or an attack. By leveraging a diverse set of synthetic artifacts, the classifier learns to identify general features indicative of morphing attacks rather than overfitting to characteristics of specific morphing techniques. Moreover, the usage of different blending masks during the pixel-artifact generation stage helps the classifier to focus on different regions of the image. This approach prevents the model from becoming biased toward facial parts where artifacts might be either more common or most obvious.

The pseudocode of SelfMAD is given in Algorithm 1.

Algorithm 1 Pseudocode of SelfMAD

- 1: **Input:** bona fide RGB images $I_{OS}^{H \times W \times 3}$, constant k
 - 2: **Output:** transformed images I_{AS} , I_{MS} , I_{FMS} , Classifier \mathcal{D}
 - 3: **Image Augmentation**
 - 4: **for** each image I_{OS} **do**
 - 5: $I_{AS} \leftarrow$ sequentially apply {RGBShift,
 - 6: HueSaturationValue, RandomBrightnessContrast,
 - 7: OneOf(RandomDownScale, Sharpen)} to I_{OS}
 - 8: **end for**
 - 9: **Pixel Artifact Generation**
 - 10: **for** each image I_{AS} **do**
 - 11: $I_{AS}' \leftarrow$ sequentially apply {Translation,
 - 12: ElasticTransform, Scaling} to I_{AS}
 - 13: $M^{H \times W} \leftarrow$ generate a binary blending mask
 - 14: $a \leftarrow$ Uniform({0.5, 0.5, 0.5, 0.375, 0.25, 0.125})
 - 15: $I_{MS} \leftarrow I_{AS}' \odot a \cdot M + I_{OS} \odot (1 - a) \cdot M$
 - 16: **end for**
 - 17: **Frequency Artifact Generation**
 - 18: **for** each image I_{MS} **do**
 - 19: $\Phi \leftarrow$ Uniform({*symmetrical grid, asymmetrical grid,*
 - 20: *random squares, random lines, stripes,*
 - 21: *square checkerboard, circular checkerboard*})
 - 22: $F_{\Phi} \leftarrow$ FourierTransform(Φ)
 - 23: $F_{MS} \leftarrow$ FourierTransform(I_{MS})
 - 24: $|F_{MS}'| \leftarrow (1 - k) \cdot |F_{MS}| \oplus k \cdot |F_{MS}|$
 - 25: $I_{MAS} \leftarrow$ InverseFourierTransform($|F_{MS}'|, \theta(F_{MS})$)
 - 26: **end for**
 - 27: **Classifier training**
 - 28: $\mathcal{L}_{BCE} \leftarrow \mathcal{D}(\theta, I_{OS}, I_{AS}, I_{MS}, I_{FMS})$
 - 29: $\omega' \text{ (updated weights of } \mathcal{D}) \leftarrow$ backpropagate $\mathcal{L}_{BCE}(\omega)$
-

IV. EXPERIMENTAL SETUP

In this section, we first describe the datasets used in our experiments. Next, we explain the evaluation metrics and provide details about the implementation of the models.

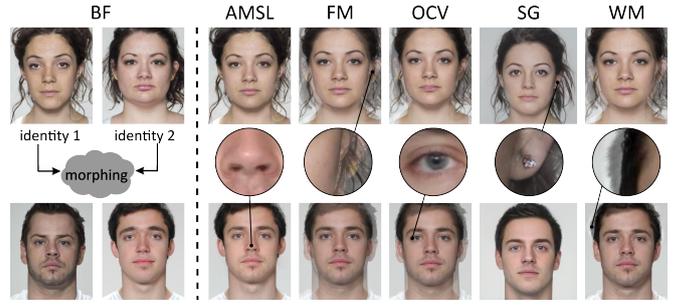


Fig. 3. Selected samples from FRL-Morphs [49], representing bona fide (BF) images (left) and morphing attacks generated with different morphing methods, i.e. AMSL, FaceMorpher (FM), OpenCV (OCV), StyleGAN2 (SG), and Webmorph (WM). Note typical ghosting artifacts of landmark-based morphing attacks and irregularities generated with StyleGAN2.

A. Datasets

Training data. In our experiments, we follow the protocols used in some recent MAD papers [6], [25] and train SelfMAD on SMDD [13]. SMDD is a large-scale synthetic dataset specifically developed for face morphing attack detection. The dataset comprises high-quality bona fide and morphed samples, generated with a privacy-preserving approach based on StyleGAN2, ensuring diversity by simulating a wide range of facial features and morphing artifacts while avoiding real-world identity data. The dataset comprises 80,000 images of size 1024×1024 pixels, including 50,000 bona fide samples and 30,000 morphing attacks, evenly split into training and testing subsets. For training SelfMAD, we use only the training subset, which contains 25,000 bona fide images. Raw images are preprocessed by detecting faces in each image using Dlib [31]. The detected square face regions are enlarged by a randomly selected margin between 4% and 20%, then cropped and resized to 384×384 pixels. Additionally, segmentations of individual facial parts are generated for each cropped face using a pretrained SegFormer [66]. Both, the cropped RGB images and their corresponding segmentations, are then fed into SelfMAD and processed according to the steps described in Section III.

In the experiments, we compare SelfMAD against both, unsupervised and discriminatively trained MAD methods. During the training of the unsupervised methods, we follow the protocols established by their respective authors and use the bona fide training data reported in the original papers. Similarly, the discriminative approaches are trained on the data utilized in their corresponding related works. Discriminatively trained MADs utilize bona fide images and morphing attacks from three publicly available datasets: LMA-DRD [15], MorGAN [14], and SMDD [13]. The LMA-DRD dataset includes morphs generated using OpenCV, with digital morphs labeled as "D" and re-digitized morphs (printed and scanned) labeled as "PS." The MorGAN dataset contains two types of morphs: LMA morphs created with OpenCV and GAN morphs produced using a DCGAN-based model. In the synthetic StyleGAN2-based SMDD dataset, both bona fide and morphing attack training subsets are utilized during training.

Testing data. All MAD methods are tested on six standard datasets, representing various face morphing genera-

TABLE I

COMPARISON OF SELF-MAD AGAINST ONE-CLASS MAD MODELS. OUR SELF-MAD IS THE TOP-PERFORMER ACROSS VARIOUS FACE MORPHING ATTACK TYPES, OUTPERFORMING THE RUNNER-UP SPL-MAD BY 10.12% IN TERMS OF EER, AND BY 17.31% AND 16.1% IN TERMS OF BPCER@APCER 5% AND 10%, RESPECTIVELY. WHILE SELF-MAD SHOWS SLIGHT UNDERPERFORMANCE WITH STYLEGAN2 (SG) MORPHS, IT SIGNIFICANTLY SURPASSES COMPETITIVE METHODS IN OTHER ADVANCED MORPHING ATTACKS, I.E. MORPH-PIPE, GREEDY-DiM AND MORCODE.

Model		FIQA-MagFace [21]			CNNIQA [21]			SPL-MAD [18]			MAD-DDPM [28]			SBI [59]			SelfMAD [Ours]		
		EER	BE@AE(%)		EER	BE@AE(%)		EER	BE@AE(%)		EER	BE@AE(%)		EER	BE@AE(%)		EER	BE@AE(%)	
			5.00	10.00		5.00	10.00		5.00	10.00		5.00	10.00		5.00	10.00		5.00	10.00
FRGC-M	FM	33.82	73.79	62.84	42.84	75.94	66.86	16.91	25.39	21.47	25.62	95.12	90.15	16.68	38.07	26.14	5.59	6.43	2.80
	OCV	33.30	74.71	62.52	43.15	74.64	66.35	20.75	32.50	25.42	28.22	95.12	90.15	15.32	36.31	25.10	2.59	1.14	0.41
	SG	14.21	26.46	17.60	36.51	70.34	57.93	16.80	26.13	21.09	9.02	95.12	90.15	52.90	97.10	94.40	15.84	45.23	25.52
FERET-M	FM	25.14	61.22	44.44	13.23	35.17	19.32	20.42	40.85	27.09	27.98	95.27	90.17	26.47	60.87	52.36	3.19	1.70	0.38
	OCV	26.14	61.50	43.95	20.45	58.60	37.23	25.71	57.45	45.60	31.38	95.27	90.17	28.73	70.08	60.61	1.13	0.57	0.38
	SG	12.67	24.63	15.71	33.84	79.55	66.17	25.33	62.06	49.72	32.14	95.27	90.17	41.83	90.55	82.42	18.14	46.12	32.33
FRLL-M	AMSL	30.94	77.94	66.18	21.61	60.29	39.22	3.26	0.50	0.50	27.13	94.94	90.02	11.76	24.23	16.78	0.99	0.05	0.05
	FM	27.99	73.04	57.35	19.97	57.84	36.76	1.03	0.99	0.99	10.40	95.19	90.38	13.73	36.99	26.10	0.00	0.26	0.17
	OCV	24.73	66.18	53.43	7.53	11.76	4.41	1.88	0.50	0.50	13.76	95.17	90.01	12.25	27.85	18.84	0.00	0.00	0.00
	SG	7.53	8.82	5.39	35.92	75.49	68.14	14.65	32.18	24.75	14.32	95.17	90.18	44.61	94.68	90.92	10.34	24.22	12.52
	WM	27.19	68.14	55.39	21.54	46.57	33.33	6.39	11.39	3.47	30.30	95.09	90.34	39.22	89.93	83.37	3.45	1.64	0.41
Morph-PIPE		49.62	91.54	84.18	66.54	98.83	97.35	18.88	33.62	25.66	13.88	95.14	90.14	30.75	92.41	77.69	5.89	12.44	0.53
Greedy-DiM		47.00	94.61	85.78	49.40	96.08	93.14	37.72	80.69	71.78	36.10	95.20	89.70	33.82	90.60	81.60	7.60	37.60	27.80
MorCode		23.60	53.21	40.63	99.17	100.00	100.00	10.77	19.09	11.86	32.93	95.14	90.06	17.67	58.80	33.46	4.08	3.64	1.21
Average		27.42	61.13	49.67	36.55	67.22	56.16	15.75	30.24	23.56	23.80	95.16	90.13	27.55	64.89	54.99	5.63	12.93	7.46

*FM: FaceMorpher, OCV: OpenCV, SG: StyleGAN2, WM: Webmorph, BE@AE: BPCER@APCER

tion techniques: FRGC-Morphs [49], FERET-Morphs [49], FRLL-Morphs [49], Morph-PIPE [69], Greedy-DiM [3], and MorCode [46]. FRGC-Morphs, FERET-Morphs, FRLL-Morphs, derived from respective source datasets FRGC [44], FERET [45] and FRLL [17], contain face morphing images generated by conventional landmark-based methods AMSL [41], FaceMorpher, OpenCV, and WebMorph. Additionally, they also contain a deep-learning subset generated with StyleGAN2. Similarly, MorCode features GAN-based images generated using a VQ-GAN-based generator, with spherical interpolation in its latent space to create a single morphed image from two input face identities. In contrast, Morph-PIPE and Greedy-DiM are based on the latest generation of generative methods, i.e., diffusion models. While Greedy-DiM is derived from FRLL, MorCode, and Morph-PIPE morphs both originate from FRGC, and fulfill the quality constraints laid down by the International Civil Aviation Organization (ICAO). Selected samples representing bona fide images and different morphing attacks are presented in Fig. 3. Faces in all testing images were detected by Dlib [31] and cropped out with a fixed margin of 12.5%.

B. Evaluation metrics

The model evaluation follows the testing protocol proposed in existing studies [18], [60] to enable consistent benchmarking and comparability. To ensure compliance with the ISO/IEC 30107-3 [27], we report the detection Equal Error Rate (EER), where the Attack Presentation Classification Error Rate (APCER), equals the Bona fide Presentation Classification Error Rate (BPCER). APCER here quantifies the proportion of attack samples misclassified as bona fide, while BPCER represents the proportion of bona fide samples misclassified as attacks. Besides EER, we also report the MAD performance at two operational points, i.e. BPCER at fixed APCER values of 5% and 10%.

C. Implementation details

The input to SelfMAD represents RGB bona fide face images. Face blending masks are generated by combining face segmentations created with a SegFormer face parser,

pretrained on the CelebAMask-HQ dataset¹. Once SelfMAD generates augmented face images with superimposed pixel and frequency artifacts, bona fide samples along with simulated morphing attacks are fed into a binary classifier. For the classifier we consider 4 different CNN architectures: EfficientNet-B4, ResNet-152, Swin-B and HRNet-W18. All of them were pretrained on ImageNet and finetuned until convergence on samples generated during earlier stages of SelfMAD. In this paper, we only report results obtained with HRNet as a backbone, as it achieved the best overall performance on considered testing datasets. A comparison of HRNet’s performance against other backbones is provided in the supplementary material. The classifier is optimized by SGD combined with SAM [20]. The radius ρ of SAM is set to 0.05 with a momentum of 0.9, while the learning rate equals 0.001. SelfMAD is implemented in Python 3.10 with PyTorch 2.4 and CUDA 12.5. Experiments were run on NVIDIA GeForce RTX 4090. The source code of SelfMAD is available at <https://github.com/LeonTodorov/SelfMAD>.

V. RESULTS

In this section, we first compare SelfMAD against competitive one-class and discriminatively trained MAD models. Next, we investigate the impact of different SelfMAD components on model’s performance. Finally, we visualize the regions of interest identified by SelfMAD when the model classifies a sample as a morphing attack.

Comparison against one-class MADs. SelfMAD is designed to be trained exclusively with bona fide input samples, so we first compare its performance against other highly competitive one-class methods. The results are summarized in Table I. As shown, SelfMAD achieves the best overall performance across various face morphing attack types, surpassing the runner-up SPL-MAD [18] by 10.12% in terms of EER, and by 17.31% and 16.1% in terms of BPCER@APCER 5% and 10%, respectively. We note, that SelfMAD is especially successful in the detection of more recent, advanced GAN and diffusion-based morphing attacks from Morph-PIPE, Greedy-DiM and MorCode, improving

¹<https://huggingface.co/jonathandinu/face-parsing>

TABLE II

COMPARISON OF SELF MAD AGAINST DISCRIMINATIVE MAD MODELS IN TERMS OF EER, BPCER@APCER (5%) AND BPCER@APCER (10%). SELF MAD SIGNIFICANTLY OUTPERFORMS COMPETING METHODS ACROSS ALL MORPHING ATTACK CATEGORIES, EXCEPT FOR FRGC AND FERET STYLEGAN2 (SG) MORPHS, WHERE OUR METHOD SHOWS A SLIGHT UNDERPERFORMANCE.

Model	Train data	Test data:	FRGC-M			FERET-M			FRLL-M				Morph-PIPE	Greedy-DiM	MorCode	Average			
			FM	OCV	SG	FM	OCV	SG	AMSL	FM	OCV	SG					WM		
Discriminatively trained	MixFaceNet [5]	D	EER	32.43	20.34	19.59	27.92	33.59	37.49	23.15	10.84	10.84	19.70	31.03	46.64	45.10	39.83	28.46	
			BE@AE(%)	5.00	85.89	51.97	64.63	69.38	79.40	91.87	71.40	23.71	39.97	54.58	85.26	98.41	93.40	92.79	71.62
			10.00	73.44	37.03	41.91	57.09	68.43	81.10	61.47	13.66	18.02	39.20	71.01	95.75	85.60	84.90	59.19	
				25.54	30.13	19.75	31.40	29.62	30.19	19.70	8.87	9.85	11.82	30.05	61.51	41.67	43.84	28.14	
				5.00	75.31	79.25	46.89	82.61	77.13	83.93	56.55	15.38	28.26	28.07	76.82	100.00	89.00	92.19	66.53
				10.00	55.19	64.83	33.82	73.53	66.16	68.43	45.06	8.33	14.00	17.76	65.36	99.62	78.80	84.45	55.38
				5.00	35.76	32.33	17.26	36.07	33.95	27.43	27.09	27.09	18.23	29.06	45.81	35.18	40.69	37.28	31.66
				10.00	85.58	85.06	49.38	84.12	89.41	76.37	84.64	69.93	44.64	83.39	94.35	78.15	93.40	83.84	78.73
				5.00	76.35	72.93	30.81	76.75	78.45	61.63	74.53	49.05	27.27	71.11	90.01	66.16	88.00	74.51	66.97
				10.00	26.25	29.91	33.40	44.08	44.72	48.19	45.32	46.80	39.41	39.41	44.83	56.62	48.04	54.42	42.96
				5.00	76.45	80.71	89.94	95.46	90.55	94.90	95.45	97.08	96.56	85.76	95.25	99.85	97.00	99.39	92.45
				10.00	60.89	68.78	79.36	90.74	83.55	88.66	88.83	92.35	93.94	78.97	92.55	99.24	93.40	95.75	86.21
				5.00	14.32	14.35	13.03	19.99	20.62	36.64	31.03	8.37	9.85	38.92	31.03	33.69	39.71	37.55	24.94
				10.00	33.40	35.27	26.97	53.69	53.88	87.33	65.56	11.51	15.89	84.94	78.38	85.66	90.20	86.12	57.77
				5.00	20.95	20.95	15.87	34.97	39.51	75.24	60.00	8.16	10.07	76.27	67.16	73.75	86.20	76.10	47.51
				10.00	36.44	38.64	36.70	14.60	11.55	40.61	35.96	14.78	11.33	35.96	35.96	13.36	17.16	17.87	25.78
				5.00	95.12	91.29	97.82	47.26	21.55	95.84	87.08	69.85	22.19	93.94	90.09	54.93	64.20	50.15	70.09
				10.00	86.00	83.51	93.67	25.33	13.99	88.09	71.95	33.51	12.04	80.61	75.68	22.76	36.60	33.92	54.12
				5.00	42.81	53.80	38.02	14.88	11.48	45.29	22.66	15.76	7.88	28.57	26.60	20.97	33.82	34.74	28.38
				10.00	94.92	98.86	66.29	52.93	17.39	97.16	53.75	34.28	12.29	78.15	66.67	65.63	83.40	85.13	64.77
				5.00	89.52	97.20	59.13	32.70	12.29	93.95	46.85	22.42	6.72	67.02	48.81	50.30	70.00	73.60	55.04
				10.00	21.53	12.93	19.08	18.50	24.66	26.44	18.23	10.34	12.81	19.21	39.41	17.27	17.16	21.15	19.91
				5.00	93.67	48.86	61.00	75.99	75.99	77.69	47.49	26.72	27.35	42.39	92.22	54.32	40.80	56.68	58.66
				10.00	66.29	20.12	39.21	48.58	65.97	60.87	33.47	12.54	19.66	30.20	83.70	33.69	28.60	39.98	41.63
			5.00	47.17	46.01	24.09	12.19	25.02	43.23	47.78	21.18	31.03	48.28	50.25	32.18	15.20	54.89	35.61	
			10.00	99.79	99.07	100.00	31.95	70.51	88.66	92.41	66.07	97.22	97.46	93.61	96.28	39.60	87.71	82.88	
			5.00	98.76	97.72	100.00	16.82	54.63	80.15	88.00	48.80	83.21	90.59	89.76	81.49	22.20	83.26	73.95	
			10.00	15.97	19.33	15.36	8.36	16.09	20.91	4.43	1.97	2.46	17.24	9.85	20.40	42.16	39.29	16.70	
			5.00	51.76	56.85	35.68	42.34	74.48	81.47	4.64	1.72	1.80	34.62	12.29	55.39	95.40	93.63	45.86	
			10.00	29.25	38.90	21.06	5.10	35.92	53.31	2.21	1.29	0.90	26.43	9.75	37.78	86.80	86.65	31.10	
			5.00	42.22	30.26	39.19	14.74	15.59	11.27	55.67	39.41	28.57	25.62	65.52	33.63	31.86	10.91	31.75	
			10.00	96.89	93.57	89.00	31.57	32.33	21.93	99.26	81.36	69.45	67.59	99.92	95.90	85.40	21.70	70.42	
			5.00	90.15	80.29	78.11	21.55	20.98	13.04	98.48	70.45	57.82	57.77	99.75	82.09	77.40	12.06	61.43	
			10.00	43.87	20.50	24.86	30.33	27.64	16.65	50.25	32.02	41.38	33.50	39.90	8.30	51.96	27.58	32.05	
			5.00	93.36	53.11	47.10	69.57	64.84	31.38	99.49	78.95	92.79	75.04	93.28	20.18	95.20	72.00	70.45	
			10.00	85.89	37.34	38.28	57.84	53.69	22.12	98.16	68.99	88.94	68.33	87.96	5.16	92.20	61.08	61.86	
			5.00	27.19	30.26	42.39	27.21	36.57	29.91	12.81	8.37	19.21	19.70	24.14	58.38	25.98	29.59	27.98	
			10.00	79.46	89.00	91.70	70.32	86.77	81.85	35.26	20.10	63.47	66.94	72.73	100.00	66.00	63.73	70.52	
			5.00	60.48	76.04	84.34	57.47	73.72	68.81	20.37	5.76	41.52	48.61	64.13	99.92	46.00	51.75	57.06	
			10.00	49.76	49.08	41.67	51.10	52.73	35.79	75.86	67.49	59.61	67.49	59.61	40.76	29.90	61.98	53.06	
			5.00	99.69	98.76	95.23	95.65	98.11	82.99	100.00	99.31	98.94	98.53	99.84	95.45	86.60	100.00	96.36	
			10.00	98.65	96.58	89.83	92.25	96.98	72.40	100.00	98.71	96.97	97.55	99.43	89.98	73.00	99.70	93.00	
			5.00	18.82	23.38	32.82	19.42	43.80	38.98	7.88	2.96	15.27	19.70	11.33	11.90	56.86	11.24	22.45	
			10.00	35.68	47.61	97.20	46.12	99.43	93.76	7.17	2.23	22.60	35.84	12.69	14.64	99.80	13.81	44.90	
			5.00	54.15	67.12	99.90	69.57	100.00	99.43	11.03	2.66	37.59	45.58	18.18	26.02	100.00	21.55	53.77	
			10.00	2.80	0.41	25.52	0.38	0.38	32.33	0.05	0.17	0.00	12.52	0.41	0.53	27.80	1.21	7.46	
			5.00	5.59	2.59	15.84	3.19	1.13	18.14	0.99	0.00	0.00	10.34	3.45	5.89	7.60	4.08	5.63	
			10.00	6.43	1.14	45.23	1.70	0.57	46.12	0.05	0.26	0.00	24.22	1.64	12.44	37.60	3.64	12.93	
		5.00	2.80	0.41	25.52	0.38	0.38	32.33	0.05	0.17	0.00	12.52	0.41	0.53	27.80	1.21	7.46		
		10.00	2.80	0.41	25.52	0.38	0.38	32.33	0.05	0.17	0.00	12.52	0.41	0.53	27.80	1.21	7.46		

*FM: FaceMorpher, OCV: OpenCV, SG: StyleGAN2, WM: Webmorph, BE@AE: BPCER@APCER

TABLE III

ABLATION STUDY EVALUATING DIFFERENT SELF MAD COMPONENTS. WHILE THE PIXEL ARTIFACT GENERATOR (PAG) ALONE ACHIEVES STRONG MORPHING DETECTION PERFORMANCE, COMBINING IT WITH THE FREQUENCY ARTIFACT GENERATOR (FAG) ENHANCES THE OVERALL SELF MAD PERFORMANCE ACROSS VARIOUS ATTACK CATEGORIES.

Model	Test data:	FRGC-M			FERET-M			FRLL-M				Morph-PIPE	Greedy-DiM	MorCode	Average		
		FM	OCV	SG	FM	OCV	SG	AMSL	FM	OCV	SG					WM	
PAG	EER	7.24	3.52	8.02	6.17	2.83	21.05	1.48	0.00	0.00	9.36	8.37	2.21	26.00	3.08	7.10	
	BE@AE(%)	5.00	11.31	2.39	18.88	7.37	1.51	48.02	0.28	0.26	0.00	13.09	12.37	0.08	95.40	1.59	15.18
		10.00	4.56	0.52	6.33	2.46	1.32	36.29	0.09	0.26	0.00	9.00	5.57	0.08	90.20	0.38	11.22
FAG	EER	66.41	47.88	36.05	54.36	54.78	23.39	83.25	82.27	73.40	46.80	80.79	97.93	78.40	97.79	65.96	
	BE@AE(%)	5.00	100.00	100.00	98.13	98.87	98.30	48.58	100.00	100.00	100.00	100.00	100.00	100.00	100.00	95.99	
		10.00	100.00	99.59	94.50	95.84	95.65	37.43	100.00	100.00	100.00	99.67	100.00	100.00	100.00	94.48	
SelfMAD	EER	5.59	2.59	15.84	3.19	1.13	18.14	0.99	0.00	0.00	10.34	3.45	5.89	7.60	4.08	5.63	
	BE@AE(%)	5.00	6.43	1.14	45.23	1.70	0.57	46.12	0.05	0.26	0.00	24.22	1.64	12.44	37.60	3.64	12.93
		10.00	2.80	0.41	25.52	0.38	0.38	32.33	0.05	0.17	0.00	12.52	0.41	0.53	27.80	1.21	7.46

*FM: FaceMorpher, OCV: OpenCV, SG: StyleGAN2, WM: Webmorph, BE@AE: BPCER@APCER

the average EER performance of the runner-ups in these three datasets, SBI [59] and MAD-DDPM [28], by more than 78%. However, while SelfMAD outperforms other models in most morphing attack categories, it exhibits slight underperformance in detecting StyleGAN2 (SG) morphs. In this specific category, the MAD approach based on face image quality assessment, FIQA-MagFace [21], demonstrates superior results despite having one of the lowest overall performance with an average EER of 27.42%. Notably, SBI, a self-supervised model, that is conceptually closest to SelfMAD, is outperformed by our model by a large margin of 21.92% EER, 51.96% BPCER@APCER 5% and

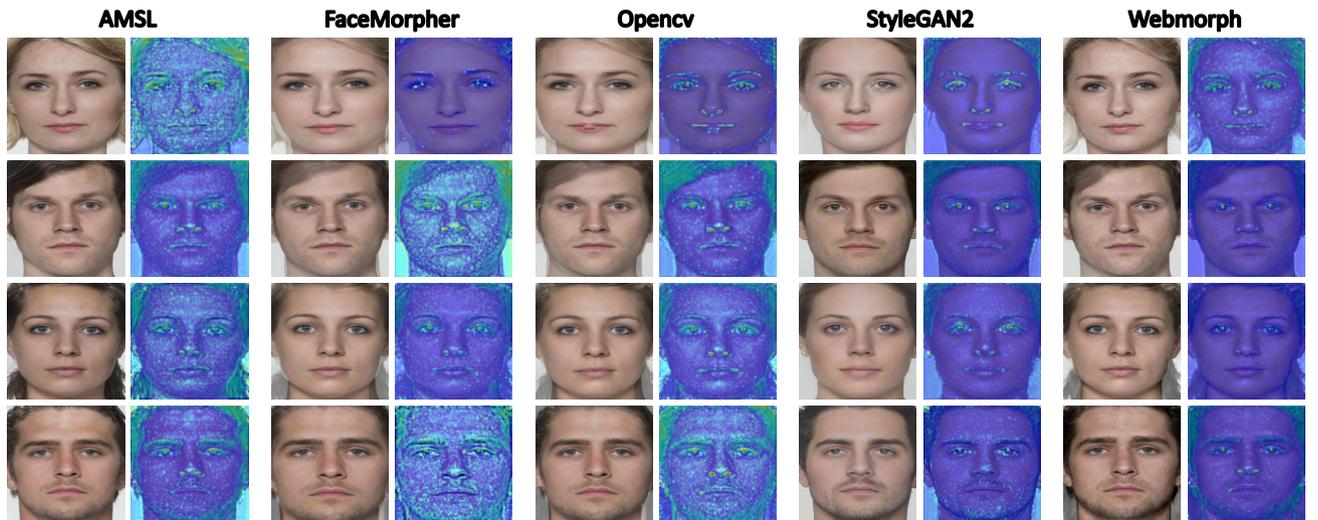


Fig. 4. **GradCAM visualizations** of SelfMAD, generated with HRNet-W18 as backbone. Heatmaps highlight eyes, nostrils, lip borders and hair, as areas where morphing pixel irregularities are most prominent. Note that StyleGAN2 images are also prone to asymmetries and blurry edges in these regions.

We run five training sessions, each time using a new, distinct type of face morphing attack as the training data and then test the trained models on testing splits from another datasets. The results, presented in Table II, show that SelfMAD is again the top-performer, surpassing the runner-up, PW-MAD trained on SMDD, by a large margin of 11.07% in terms of EER, 32.93% in terms of BPCER@APCER 5% and 23.64% in terms of BPCER@APCER 10%. However, we again note that SelfMAD shows a slight underperformance in FRGC and FERET StyleGANv2 (SG) morphs.

Ablation. To assess the contributions of individual components in SelfMAD, we conduct an ablation study, where we evaluate the model’s detection performance under three scenarios: (1) we generate only pixel-space artifacts, (2) we generate only frequency-space artifacts, and (3) both, pixel and frequency artifacts, are added to the original bona fide input image. The results of this analysis, conducted with the best-performing SelfMAD backbone, HRNet-W18, are summarized in Table III. Notably, the generation of pixel-space artifacts alone yields strong morphing attack detection performance across various morphing techniques, including some advanced approaches, such as StyleGAN2, Morph-PIPE and MorCode. In contrast, the generation of only frequency-space artifacts performs poorly, in terms of all three evaluation metrics. However, when combined with pixel-level artifacts, frequency-space artifacts significantly reduce the EER and BPCER@APCER performance in most of the morphing attack categories. This finding underscores the advantage of our method in leveraging both artifact generation strategies to enhance overall detection performance in diverse morphing attack scenarios.

Explainability Analysis. To gain deeper insight into SelfMAD, we generate Grad-CAM [58] heatmaps, that show the regions of interest identified by the model during the classification of bona fide and morphing attack samples. For this purpose, we use the best-performing SelfMAD backbone, HRNet-W18, and extract heatmaps at the fourth

stage of the classifier. Selected examples are illustrated in Fig. 4. We observe that SelfMAD primarily focuses on the eyes, nostrils and borders of the lips. These regions are in fact facial parts where pixel-space irregularities, such as ghosting artifacts, are most apparent. Similarly, StyleGAN2-generated face images are also known to exhibit inconsistencies in these areas, such as asymmetries, unnatural reflections, blurry edges or a lack of clear edge definition etc. Additionally, in some images the model also highlights hairy regions such as beard, mustache, eyebrows and hairstyles (Fig. 4, row 4).

VI. CONCLUSION

We presented SelfMAD, a self-supervised model that detects face morphing attacks by replicating typical artifacts of various common face morphing techniques. In extensive experiments across multiple widely used datasets, SelfMAD consistently outperformed state-of-the-art MAD models, reducing the equal error rate (EER) by a significant margin compared to leading discriminative and unsupervised MAD approaches. By leveraging both pixel and frequency-space artifacts generation, SelfMAD demonstrated superior generalization, particularly against unseen morphing techniques.

Limitations and Future Work. Despite its strong performance, SelfMAD has some limitations. While the model shows robustness across a range of face morphing techniques, including some very recent GAN and diffusion-based attacks, its detection performance slightly decreases when confronted with face morphs produced StyleGAN2. To address this issue, our future work will focus on enhancing the augmentation techniques used to simulate morphing artifacts, incorporating additional proxy tasks to further boost the model’s performance. Moreover, we plan to refine the feature extraction architecture and further expand our evaluation experiments by including additional recent face morphing attack methods to provide an even broader and more comprehensive assessment of our approach.

VII. ETHICAL IMPACT STATEMENT

The primary goal of this work is to enhance the robustness and generalization of Morphing Attack Detection (MAD) techniques in order to protect face recognition systems from identity fraud and other security threats. The presented research does not involve human subjects, and all datasets utilized are publicly available, ensuring compliance with standard ethical practices in data handling.

We acknowledge that the use of facial recognition technologies has raised concerns about privacy, surveillance, and potential misuse in societal contexts. However, the proposed SelfMAD method does not pose significant privacy concerns, as it is designed exclusively to detect manipulated face images (morphs) rather than perform face recognition. The development and deployment of MAD systems enhance the security of face recognition technologies, helping prevent identity theft and fraud by ensuring individuals cannot use morphed images to falsely authenticate themselves.

While MAD technology has clear benefits in securing digital identity systems, we acknowledge that there is a possibility for unintended misuse if deployed without adequate oversight. For example, in contexts where MAD tools are used, improper configurations or interpretations of results could lead to incorrect conclusions, such as misidentifying legitimate users as fraudulent due to false positives. To mitigate such risks, we recommend adherence to strict operational protocols, including continuous performance evaluation, calibration for different datasets, and clear guidelines on the application scope.

It is important to note that the standard public datasets used in this research, while widely accepted in the research community, may carry inherent biases, such as imbalanced demographics, which could influence the performance of the proposed model. In this work, we adhered to established protocols from prior studies to ensure consistency and comparability. However, we emphasize the need for future research to incorporate more diverse datasets and benchmarks that better reflect real-world scenarios, reducing potential biases and enhancing fairness in MAD systems.

REFERENCES

- [1] A. Agarwal and N. Ratha. Chapter 8 - Manipulating faces for identity theft via morphing and deepfake: Digital privacy. In *Deep Learning*, volume 48, pages 223–241. Elsevier, 2023. 1
- [2] Y. Baweja, P. Oza, P. Perera, and V. M. Patel. Anomaly Detection-Based Unknown Face Presentation Attack Detection. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–9, 2020. 2, 3
- [3] Z. W. Blasingame and C. Liu. Greedy-DiM: Greedy Algorithms for Unreasonably Effective Face Morphs. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–11, 2024. 2, 6
- [4] Z. W. Blasingame and C. Liu. Leveraging Diffusion for Strong and High Quality Face Morphing Attacks. *IEEE Transactions on Biometrics, Behavior, and Identity Science (TBIOM)*, 6(1):118–131, 2024. 2
- [5] F. Boutros, N. Damer, M. Fang, F. Kirchbuchner, and A. Kuijper. MixFaceNets: Extremely Efficient Face Recognition Networks. In *IEEE IJCB*, 2021. 1, 3, 7
- [6] E. Caldeira, P. C. Neto, T. Gonçalves, N. Damer, A. F. Sequeira, and J. S. Cardoso. Unveiling the Two-Faced Truth: Disentangling Morphed Identities for Face Morphing Detection. In *IEEE European Signal Processing Conference (EUSIPCO)*, 2023. 1, 5
- [7] T. Chugh and A. K. Jain. Fingerprint Spoof Detector Generalization. *IEEE Transactions on Information Forensics and Security (TIFS)*, 16:42–55, 2021. 2, 3
- [8] D. A. Cocomini, R. Caldelli, C. Gennaro, G. Fiameni, G. Amato, and F. Falchi. Deepfake Detection without Deepfakes: Generalization via Synthetic Frequency Patterns Injection, 2024. 2, 3
- [9] R. Corvi, D. Cozzolino, G. Poggi, K. Nagano, and L. Verdoliva. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 973–982, 2023. 2
- [10] R. Corvi, D. Cozzolino, G. Poggi, K. Nagano, and L. Verdoliva. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 973–982, 2023. 4
- [11] N. Damer, M. Fang, P. Siebke, J. N. Kolf, M. Huber, and F. Boutros. MorDIFF: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Diffusion Autoencoders. In *International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2023. 1, 2
- [12] N. Damer, J. H. Grebe, S. Zienert, F. Kirchbuchner, and A. Kuijper. On the Generalization of Detecting Face Morphing Attacks as Anomalies: Novelty vs. Outlier Detection. In *IEEE BTAS*, 2019. 1, 2, 3
- [13] N. Damer, C. A. F. López, M. Fang, N. Spiller, M. V. Pham, and F. Boutros. Privacy-Friendly Synthetic Data for the Development of Face Morphing Attack Detectors. *IEEE CVPRW*, 2022. 3, 5
- [14] N. Damer, A. M. Saladić, A. Braun, and A. Kuijper. MorGAN: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Generative Adversarial Network. In *IEEE BTAS*, 2018. 2, 5
- [15] N. Damer, N. Spiller, M. Fang, F. Boutros, F. Kirchbuchner, and A. Kuijper. PW-MAD: Pixel-Wise Supervision for Generalized Face Morphing Attack Detection. In *Springer AVC*, 2021. 2, 5, 7
- [16] D. Deb, X. Liu, and A. K. Jain. FaceGuard: A Self-Supervised Defense Against Adversarial Face Images. In *IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–8, 2023. 2, 3
- [17] L. DeBruine and B. Jones. Face Research Lab London Set, 2017. 6
- [18] M. Fang, F. Boutros, and N. Damer. Unsupervised Face Morphing Attack Detection via Self-paced Anomaly Detection. In *IJCB*, 2022. 1, 3, 6, 7
- [19] M. Ferrara, A. Franco, and D. Maltoni. *On the Effects of Image Alterations on Face Recognition Accuracy*, pages 195–222. Springer International Publishing, 2016. 1
- [20] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur. Sharpness-aware Minimization for Efficiently Improving Generalization. In *International Conference on Learning Representations (ICLR)*, 2021. 6
- [21] B. Fu and N. 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. 3, 6, 7
- [22] K. Grm, V. Štruc, A. Artiges, M. Caron, and H. K. Ekenel. Strengths and Weaknesses of Deep Learning Models for Face Recognition against Image Degradations. *IET Biometrics*, 2018. 1
- [23] M. Hildebrandt, T. Neubert, A. Makrushin, and J. Dittmann. Benchmarking face morphing forgery detection: Application of stirtrace for impact simulation of different processing steps. In *International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2017. 2
- [24] H. Hojjati, T. K. K. Ho, and N. Armanfard. Self-supervised anomaly detection in computer vision and beyond: A survey and outlook. *Neural Networks*, 172:106106, 2024. 2
- [25] M. Huber, F. Boutros, A. T. Luu, K. Raja, R. Ramachandra, N. Damer, P. C. Neto, T. Gonçalves, A. F. Sequeira, J. S. Cardoso, J. Tremoço, M. Lourenço, S. Serra, E. Cermeño, M. Ivanovska, B. Batagelj, A. Kronovšek, P. Peer, and V. Štruc. SYN-MAD 2022: Competition on Face Morphing Attack Detection Based on Privacy-aware Synthetic Training Data. In *IJCB*, 2022. 2, 5
- [26] M. Ibsen, L. J. Gonzalez-Soler, C. Rathgeb, P. Drozdowski, M. Gomez-Barrero, and C. Busch. Differential Anomaly Detection for Facial Images. In *IEEE WIFS*, 2021. 1, 3
- [27] International Organization for Standardization. ISO/IEC DIS 30107-3:2016: Information Technology – Biometric presentation attack detection – P. 3: Testing and reporting, 2017. 6
- [28] M. Ivanovska and V. Štruc. Face morphing attack detection with denoising diffusion probabilistic models. In *2023 11th International Workshop on Biometrics and Forensics (IWBF)*, 2023. 2, 3, 6, 7
- [29] H. Kim, X. Cui, M.-G. Kim, and T. H. B. Nguyen. Fingerprint

- Generation and Presentation Attack Detection using Deep Neural Networks. In *IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 375–378, 2019. 2, 3
- [30] Y. Kim, T. Lee, Y. Hyun, E. Coatanea, S. Mika, J. Mo, and Y. Yoo. Self-supervised representation learning anomaly detection methodology based on boosting algorithms enhanced by data augmentation using StyleGAN for manufacturing imbalanced data. *Computers in Industry*, 153:104024, 2023. 2, 3
- [31] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009. 5, 6
- [32] N. Larue, N.-S. Vu, V. Struc, P. Peer, and V. Christophides. Seeable: Soft discrepancies and bounded contrastive learning for exposing deepfakes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21011–21021, 2023. 4
- [33] C.-L. Li, K. Sohn, J. Yoon, and T. Pfister. CutPaste: Self-Supervised Learning for Anomaly Detection and Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9664–9674, June 2021. 2, 3
- [34] L. Li, J. Bao, T. Zhang, H. Yang, D. Chen, F. Wen, and B. Guo. Face X-Ray for More General Face Forgery Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5000–5009, 2020. 2, 3
- [35] Y. Li, P. Angelov, and N. Suri. Self-Supervised Representation Learning for Adversarial Attack Detection, 2024. 2
- [36] Y. Li and S. Lyu. Exposing DeepFake Videos By Detecting Face Warping Artifacts. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019. 2
- [37] Y. Lu and T. Ebrahimi. Towards the Detection of AI-Synthesized Human Face Images, 2024. 2, 3
- [38] A. Makrushin, C. Kraetzer, J. Dittmann, C. Seibold, A. Hilsmann, and P. Eisert. Dempster-Shafer Theory for Fusing Face Morphing Detectors. In *EUSIPCO*, 2019. 2
- [39] A. Makrushin, T. Neubert, and J. Dittmann. Automatic Generation and Detection of Visually Faultless Facial Morphs. In *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*, pages 39–50, 2017. 2
- [40] M. Naseer, S. Khan, M. Hayat, F. Khan, and F. Porikli. A Self-supervised Approach for Adversarial Robustness. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 259–268, 2020. 2
- [41] T. Neubert, A. Makrushin, M. Hildebrandt, C. Kraetzer, and J. Dittmann. Extended StirTrace benchmarking of biometric and forensic qualities of morphed face images. *IET Biometrics*, 2018. 6
- [42] T. Ojala, M. Pietikäinen, and D. Harwood. A Comparative Study of Texture Measures With Classification Based on Featured Distributions. *Pattern Recognition*, 1996. 2
- [43] V. Ojansivu and J. Heikkilä. Blur Insensitive Texture Classification Using Local Phase Quantization. In *Img. and Sig. Processing*, 2008. 2
- [44] 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 *CVPR*, volume 1, pages 947–954 vol. 1, 2005. 6
- [45] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss. The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing*, 16(5):295–306, 1998. 6
- [46] A. R. PN, R. Ramachandra, S. Venkatesh, K. S. Rao, P. Mitra, and R. Krishna. MorCode: Face Morphing Attack Generation using Generative Codebooks. In *35th British Machine Vision Conference Workshop Proceedings (BMVCW)*, 2024. 6
- [47] R. Raghavendra, K. B. Raja, S. Venkatesh, and C. Busch. Transferable Deep-CNN Features for Detecting Digital and Print-Scanned Morphed Face Images. In *IEEE CVPRW*, 2017. 2
- [48] R. Ramachandra, S. Venkatesh, K. Raja, and C. Busch. Detecting Face Morphing Attacks with Collaborative Representation of Steerable Features. In *CVIP*, 2020. 1, 2, 7
- [49] E. Sarkar, P. Korshunov, L. Colbois, and S. Marcel. Vulnerability Analysis of Face Morphing Attacks from Landmarks and Generative Adversarial Networks. *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2020. 5, 6
- [50] S. Schaefer, T. McPhail, and J. Warren. Image deformation using moving least squares. *ACM Trans. Graph.*, 25(3):533–540, 2006. 2
- [51] U. Scherhag, L. Debiasi, C. Rathgeb, C. Busch, and A. Uhl. Detection of Face Morphing Attacks Based on PRNU Analysis. *IEEE TBBIS*, 2019. 2
- [52] 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 *IWBF*, 2017. 1
- [53] 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 *International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2017. 2
- [54] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch. Face Recognition Systems Under Morphing Attacks: A Survey. *IEEE Access*, 7:23012–23026, 2019. 1, 2
- [55] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch. Face Recognition Systems Under Morphing Attacks: A Survey. *IEEE Access*, 2019. 2
- [56] C. Seibold, A. Hilsmann, and P. Eisert. Reflection Analysis for Face Morphing Attack Detection. In *EUSIPCO*, 2018. 2
- [57] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert. Detection of Face Morphing Attacks by Deep Learning. In *International Workshop on Digital Forensics and Watermarking (IWDW)*, pages 107–120, 2017. 2
- [58] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017. 8
- [59] K. Shiohara and T. Yamasaki. Detecting Deepfakes With Self-Blended Images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18720–18729, 2022. 2, 3, 4, 6, 7
- [60] J. Tapia and C. Busch. AlphaNet: Single Morphing Attack Detection Using Multiple Contributors. In *IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6, 2023. 6
- [61] S. Venkatesh, R. Ramachandra, K. Raja, and C. Busch. Face Morphing Attack Generation and Detection: A Comprehensive Survey. *IEEE Transactions on Technology and Society*, 2(3):128–145, 2021. 2
- [62] 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 *8th International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2020. 2
- [63] L. Wandzik, G. Kaeding, and R. V. Garcia. Morphing Detection Using a General-Purpose Face Recognition System. In *EUSIPCO*, 2018. 2
- [64] H. Wang, J. Fei, Y. Dai, L. Leng, and Z. Xia. General GAN-generated Image Detection by Data Augmentation in Fingerprint Domain. In *IEEE International Conference on Multimedia and Expo (ICME)*, pages 1187–1192, 2023. 2, 3
- [65] Y. Weng, L. Wang, X. Li, M. Chai, and K. Zhou. Hair Interpolation for Portrait Morphing. *Computer Graphics Forum*, 2013. 2
- [66] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In *Neural Information Processing Systems (NIPS)*, 2021. 5
- [67] T. Yang, J. Cao, D. Wang, and C. Xu. Fingerprints of generative models in the frequency domain. *arXiv preprint arXiv:2307.15977*, 2023. 4
- [68] V. Zavrtnik, M. Kristan, and D. Skočaj. DRAEM - A Discriminatively Trained Reconstruction Embedding for Surface Anomaly Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8330–8339, October 2021. 2, 3
- [69] H. Zhang, R. Ramachandra, K. Raja, and C. Busch. Morph-PIPE: Plugging in Identity Prior to Enhance Face Morphing Attack Based on Diffusion Model. In *Proceedings of Norwegian Information Security Conference (NISK)*, 2023. 6
- [70] 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 TBBIS*, 3(3):365–383, 2021. 2
- [71] K. Zhang, Y. Zhou, X. Xu, B. Dai, and X. Pan. DiffMorpher: Unleashing the Capability of Diffusion Models for Image Morphing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7912–7921, 2024. 1