

SRL-MAD: Structured Residual Latents for One-Class Morphing Attack Detection

Diogo J. Paulo^{1,2,3}, Hugo Proença^{1,2}, João C. Neves^{1,3}

¹University of Beira Interior, Portugal ²IT: Instituto de Telecomunicações ³NOVA LINCS
diogo.paulo@ubi.pt

Abstract—Face morphing attacks represent a significant threat to biometric systems as they allow multiple identities to be combined into a single face. While supervised morphing attack detection (MAD) methods have shown promising performance, their reliance on attack-labeled data limits generalization to unseen morphing attacks. This has motivated increasing interest in one-class MAD, where models are trained exclusively on bona fide samples and are expected to detect unseen attacks as deviations from the normal facial structure. In this context, we introduce SRL-MAD, a one-class single-image MAD that uses structured residual Fourier representations for open-set morphing attack detection. Starting from a residual frequency map that suppresses image-specific spectral trends, we preserve the two-dimensional organization of the Fourier domain through a ring-based representation and replace azimuthal averaging with a learnable ring-wise spectral projection. To further encode domain knowledge about where morphing artifacts arise, we impose a frequency-informed inductive bias by organizing spectral evidence into low, mid, and high-frequency bands and learning cross-band interactions. These structured spectral features are mapped into a latent space designed for direct scoring, avoiding the reliance on reconstruction errors. Extensive evaluation on FERET-Morph, FRLL-Morph, and MorDIFF demonstrates that SRL-MAD consistently outperforms recent one-class and supervised MAD models. Overall, our results show that learning frequency-aware projections provides a more discriminative alternative to azimuthal spectral summarization for one-class morphing attack detection.

Index Terms—Morphing Attack Detection, Frequency-Aware Representation Learning, Open-Set Anomaly Detection

I. INTRODUCTION

Face morphing attacks represent a significant threat to biometric authentication systems since they allow adversaries to bypass face recognition by blending facial features from multiple individuals. As a result, single-image morphing attack detection (S-MAD) has been investigated to enhance the security of biometric pipelines. However, despite progress in supervised S-MAD, real-world deployment remains difficult due to the diversity of morphing techniques and the frequent occurrence of previously unseen attack distributions as novel morphing methods are developed. This gap has motivated the interest in one-class and anomaly-based S-MAD, in which models are trained exclusively on bona fide samples and are expected to detect previously unseen morphing attacks as deviations from normal patterns [1]–[3]. These methods reduce the dependence on specific attack models and are better aligned with open-set and cross-dataset evaluation protocols.




Fig. 1. **Illustration of SRL-MAD.** From a single face image, residual Fourier features are extracted and divided into low, mid, and high-frequency bands. These are projected into a compact 1D latent space where morphing attacks appear well-separated from bona fides.

Nonetheless, the development of such methods that are both interpretable and robust remains a challenge, particularly when dealing with high-dimensional facial representations. Moreover, existing one-class models are sensitive to the quality and size of the training data [1], struggling to learn representations that explicitly separate bona fide images from morph-specific artifacts [4].

Motivated by the current challenges, we propose SRL-MAD, a one-class method that transforms structured residual Fourier features, extracted from facial images, into an interpretable discriminative score. Rather than relying on reconstruction errors, which have been shown to be unstable and attack-dependent [2], the proposed method learns a frequency-structured latent representation that captures bona fide residual patterns. Therefore, morphing attacks are detected as deviations from this learned manifold. Our approach starts by obtaining a residual Fourier map, from a single face image, by removing the image-specific power-law spectrum $f^{-\alpha}$, emphasizing deviations associated with morphing artifacts [5]. Rather than collapsing this residual map using the azimuthal average technique (i.e., radially averaging frequency magnitudes and

discarding angular structure), we preserve its two-dimensional organization by dividing it into concentric rings and introducing a learnable ring-wise spectral projection mechanism, which enables adaptive weighting of frequency components within each spatial ring and eases more expressive modeling of residual statistics while preserving the geometric structure. Furthermore, we introduce a structured frequency prior by grouping rings into low, mid, and high-frequency bands, which are projected into a compact representation that captures cross-band interactions. Projecting all bands into these three bands allows the model to jointly reason about global coherence (low bands) and localized morphing traces that typically manifest in the mid/high bands. In contrast to the azimuthal average, these bands are learned, creating a latent space that is enforced by mapping this representation to a single latent coordinate, serving as a discriminative score where morphing attacks appear well-separated from bona fides (Figure 1). Anomaly scoring is performed in the latent space rather than the reconstruction space, thereby improving stability across morphing techniques and datasets. We evaluate the proposed method on three common benchmarks, including FERET-Morph, FRLL-Morph [6], and MorDIFF [7], which cover landmark-based, GAN-based, and diffusion-based morphing attacks.

Our method achieves an average Equal Error Rate (EER) of 4.76% across all test sets, outperforming recent one-class approaches such as SPL-MAD [2] and MAD-DDPM [1], and remains competitive with SelfMAD++ [8]. Additionally, our approach surpasses several supervised methods, including MADation [9], achieving a lower average EER. These results show that using structured frequency-domain residuals and enforcing a compact latent space provide a strong foundation for open-set morphing attack detection, advancing the state of the art in one-class S-MAD.

In summary, our contributions are three-fold:

- **Frequency-aware modeling for one-class morphing attack detection.** We introduce SRL-MAD that transforms residual Fourier features into a compact latent manifold explicitly separating bona fide images from morphs.
- **Learnable geometry-preserving spectral aggregation.** We replace azimuthal averaging with learnable ring-wise spectral projection, enabling adaptive weighting of frequency components while preserving the two-dimensional organization of the Fourier magnitude map.
- **Frequency-domain inductive bias with cross-band interaction modeling.** We encode prior knowledge on where morphing artifacts manifest by organizing residual spectra into low, mid, and high-frequency groups and learn a joint projection that captures dependencies between global structure (low) and localized blending/texture inconsistencies (mid/high).

II. RELATED WORK

Morphing Attack Detection. Early approaches of face morphing attack detection used handcrafted features such as LBP, LPQ, BSIF, or SIFT descriptors to capture subtle morphing artifacts [10], [11], but showed limited generalization

to unseen morphing techniques. In contrast, deep learning-based S-MAD [9], [12]–[16] approaches have achieved higher performance since they learned hierarchical feature representations. Recently, foundation model-based approaches such as MADation [9] adapted CLIP-like vision-language architectures [17] for MAD tasks, exploiting large-scale pretraining to enhance zero-shot generalization to unseen morphing techniques. However, most of these supervised MAD methods still rely on labeled morphing datasets and exhibit limited cross-domain robustness. To address generalization gaps, SelfMAD [18] introduced a self-supervised approach to simulate morphing artifacts in both pixel and frequency domains. Its successor, SelfMAD++ [8], further integrates foundation models with local feature enhancement and self-supervised pretraining, improving generalization to unseen morphing types through localized artifact reasoning and representation learning. Additionally, recent work has explored alternative input representations that use residual images or frequency-domain representations [2], [15], [19]–[21] instead of spatial-domain RGB images. These methods aim to emphasize subtle artifacts introduced by morphing while suppressing identity-related information. Our work follows this line by operating on structured residual spectra, but differs in its focus on anomaly detection rather than supervised classification.

Anomaly-Based and One-Class MAD. Anomaly detection has emerged as a promising direction for MAD, modeling the distribution of bona fide samples and treating morphs as outliers. In [2], the authors introduced self-paced learning to separate bona fide and morphed samples through unsupervised training, improving generalization and reconstruction separability, and showed that morphing attacks can be easier to reconstruct than bona fide images, as evidenced by lower reconstruction errors. Furthermore, [1] proposed MAD-DDPM, a diffusion-based one-class model that learns bona fide image distributions using denoising diffusion probabilistic models [22]. Despite these advances, these approaches typically rely on reconstruction errors as anomaly scores. However, reconstruction errors are often sensitive to architectural choices and may fail when models generalize well to both bona fide and attack samples [2]. Additionally, [3] explored attack-agnostic features extracted by large vision models pretrained only on bona fide data. These features, modeled with a Gaussian mixture model, outperform traditional one-class detectors across unseen morph scenarios. Our method aligns with this line of work but explicitly enforces a compact latent space in a one-dimensional space, making deviations from the bona fide distribution directly interpretable and robust. The anomaly score is derived from the bona fide latent distribution, avoiding reliance on reconstruction error at test time.

Frequency-Domain cues in MAD. Frequency-domain analysis has proven effective in face morphing detection, as morphing operations often introduce characteristic artifacts in specific frequency bands [21], [23]. Prior works have employed Fourier analysis, wavelet decompositions, and radial or azimuthal averages to summarize frequency information [5],




Fig. 2. **Overview of the proposed method (SRL-MAD).** We extract the Residual Fourier Map from a single face image by subtracting a power-law baseline to highlight morphing artifacts. The map is divided into concentric rings to form a structured 2D feature representation, which is then processed through a learnable ring-wise spectral projection mechanism. These features are refined into low, mid, and high-frequency bands and projected by an autoencoder into a compact 1D latent space. The final 1D latent coordinate serves as a discriminative score.

[24]. While the azimuthal average provides a compact representation, it assumes uniform importance across frequency components and spatial regions. Building on these insights, we improve frequency-based MAD by replacing azimuthal average with learnable ring-wise spectral projections, allowing the network to weight frequency components differently for each spatial ring. Furthermore, by combining rings into low, mid, and high-frequency groups, the model captures structured spectral dependencies while maintaining low dimensionality.

III. METHODOLOGY

Figure 2 presents an overview of the proposed method. Our approach operates in the frequency domain and is based on residual Fourier analysis to extract features, which are subsequently modeled using an autoencoder. While inspired by prior work on global Fourier residual representations [5], SRL-MAD goes beyond feature extraction, preserving the two-dimensional frequency structure through a learnable ring-wise spectral projection that replaces the azimuthal average and introduces frequency priors to guide the latent representation used for the score.

A. Residual Fourier Map

Given an input facial image I , we obtain the two-dimensional discrete Fourier transform and keep the magnitude spectrum. To reduce dynamic range, a logarithmic compression is applied, obtaining:

$$S(u, v) = \log(1 + |\mathcal{F}\{I(x, y)\}|), \quad (1)$$

where $\mathcal{F}(\cdot)$ denotes the 2D DFT mapping the spatial signal $I(x, y)$ to frequency coordinates (u, v) . Natural images approximately follow a power-law decay in the frequency

domain. To suppress this dominant trend and highlight deviations associated with morph artifacts, we estimate a power-law baseline from the image itself. First, the frequency plane is divided into concentric radial bands, and a one-dimensional radial profile is obtained by averaging $S(u, v)$ within each band. A linear model is then fitted in log-log space:

$$\log S(f) \approx a + b \log f, \quad (2)$$

where f denotes the radial frequency associated with each band. This fitted baseline is subsequently mapped back to the two-dimensional frequency plane, producing a baseline map S' that depends only on radial distance.

The final residual Fourier map $\mathcal{R}$ is obtained as the difference between the log-magnitude spectrum and the mapped baseline. The DC component is set to zero. It is worth mentioning that this residual is obtained per-image and does not involve subtracting statistics estimated from the training set.

B. Ring-Based 2D Feature Representation

The residual Fourier map is divided into R concentric rings centered at the DC component. Each ring collects frequency coefficients with a similar radial distance. Unlike earlier approaches that collapse each ring into a single value through azimuthal averaging [5], i.e., averaging magnitudes over all angular directions at a fixed radius to produce a 1D radial profile, we keep all residual coefficients within each ring.

Since the number of frequency samples increases with radius, the number of coefficients per ring also varies. Let K_r denote the number of coefficients in ring r , and define $K_{\max} = \max_r K_r$. We construct a fixed-size matrix $\mathbf{X} \in \mathbb{R}^{R \times K_{\max}}$, where each row corresponds to one ring

and contains the residual coefficients ordered clockwise. Rings with fewer than $K_{\max}$ coefficients are padded with zeros and a corresponding binary mask $\mathbf{M}$ is saved to identify valid coefficients and ensure that padded positions do not influence learning. This representation preserves fine-grained frequency information within each band, avoiding azimuthal averaging and preserving intra-band structure that can be informative for morphing artifact detection.

C. Learnable Ring Projection

The ring matrix $\mathbf{X}$ is first processed by a learnable ring-wise projection that combines frequency coefficients within each ring. Specifically, we learn a set of projection kernels that operate along the angular dimension, allowing the model to emphasize artifact-relevant spectral patterns while suppressing nuisance variability. This operation is performed as a depth-wise one-dimensional convolution with kernel size $K_{\max}$, yielding a compact ring descriptor $\mathbf{x}_{\text{ring}} \in \mathbb{R}^R$. To prevent padded coefficients from affecting the learned projection weights, the binary mask is enforced throughout training, ensuring that gradients do not propagate through invalid positions. Compared to azimuthal averaging or linear dimensionality reduction (e.g., PCA), this formulation allows the model to adaptively emphasize discriminative frequency patterns.

D. Frequency-Band Refinement

To encode a frequency-domain inductive bias informed by how morphing artifacts manifest across bands, the ring descriptor x_{ring} is divided into three contiguous frequency bands corresponding to low, mid, and high frequencies, which are defined by equally dividing the radial index range. For each group, we learn a band-wise linear projection that combines the ring responses within the band. This results in a three-dimensional refined representation:

$$x_{\text{ref}} = [f_{\text{low}}(x_{\text{ring}}^{(1)}), f_{\text{mid}}(x_{\text{ring}}^{(2)}), f_{\text{high}}(x_{\text{ring}}^{(3)})]. \quad (3)$$

This step explicitly summarizes residual evidence across frequency ranges, preserving band-specific priors while reducing dimensionality before encoding.

E. Autoencoder with Template Reconstruction

The refined vector x_{ref} is encoded using an autoencoder with a one-dimensional latent representation. The encoder applies a small non-linear mapping to obtain the latent code $\hat{z}$, while the decoder maps $\hat{z}$ back to a ring-level reconstruction $\hat{x}_{\text{ring}} \in \mathbb{R}^R$. To recover the full ring matrix, we introduce a learnable template matrix $\mathbf{T} \in \mathbb{R}^{R \times K_{\max}}$. The reconstructed matrix is obtained through a factorized formulation:

$$\hat{\mathbf{X}} = \hat{x}_{\text{ring}} \odot \mathbf{T}, \quad (4)$$

where $\hat{x}_{\text{ring}}$ is broadcast along the frequency dimension. This structure constrains the reconstruction to respect the ring-wise organization while allowing frequency-specific patterns to be learned through the matrix $\mathbf{T}$.

F. Training Objective

The network is trained by minimizing two reconstruction losses. The first enforces accurate reconstruction of the ring-level descriptor, while the second term encourages consistency between the reconstructed and original matrices and is applied only to valid coefficients.

$$\mathcal{L} = \underbrace{\|\hat{\mathbf{x}}_{\text{ring}} - \mathbf{x}_{\text{ring}}\|_2^2}_{\text{ring-level reconstruction loss}} + \underbrace{\frac{\|(\hat{\mathbf{X}} - \mathbf{X}) \odot \mathbf{M}\|_2^2}{\sum \mathbf{M}}}_{\text{masked matrix reconstruction loss}}. \quad (5)$$

Optimizing these terms of the loss jointly is essential since the ring-level loss stabilizes the low-dimensional representation, while the full reconstruction loss ensures that gradients propagate through the learnable projection and template components, preventing degenerate solutions.

IV. EXPERIMENTAL SETUP

Datasets. Although our approach could be adapted to arbitrary bona fide face datasets since our approach is one-class unsupervised, we strictly follow the protocols used in some recent MAD papers [9], [18], [26] and train our method on SMDD [25] to ensure comparability with prior work. No morph samples are used during training, reflecting realistic deployment scenarios where attack samples may not be available during system development. Nevertheless, training on larger bona fide datasets (e.g., CASIA-WebFace) is an interesting direction to assess potential gains in generalization. To assess the generalization capability of the proposed method, as test data, we rely on three morphing datasets, FRLL-Morph, FERET-Morph [6], and MorDIFF [7]. SMDD is a synthetic StyleGAN2-based development dataset with OpenCV morphs, whereas FRLL-Morph is an FRLL-based benchmark containing landmark and GAN-based morphs being identity-disjoint from SMDD. MorDIFF [7] provides additional FRLL-based diffusion morphs, enabling evaluation under stronger and more realistic attacks.

Evaluation metrics. All experiments follow the ISO/IEC 30107-3 [27] standard for biometric presentation attack detection. Based on the anomaly scores, we report the Bona Fide Presentation Classification Error Rate (BPCER) at fixed Attack Presentation Classification Error Rate (APCER) thresholds of 5% and 10%. Additionally, the overall detection accuracy is reported as the Equal Error Rate (EER), where APCER equals BPCER.

Implementation Details. We conduct the experiments on NVIDIA GeForce RTX 5090 containing 32GB of VRAM, and the framework utilized was Pytorch. To ensure consistency across evaluations, Fourier features were extracted using images of uniform size 500×500, as previous works [5], [8], [15] suggest that higher-resolution images typically preserve more morphing artifacts, specifically in the Fourier domain. We optimize the model of our method with Adam with a learning rate of 5×10^{-5} using a batch size of 128 for training.

TABLE I

COMPARISON WITH STATE-OF-THE-ART ONE-CLASS AND SELF-SUPERVISED S-MAD METHODS TRAINED ON SMDD. WE USE METRICS SUCH AS EER (%), AND BPCER@APCER (%), REPORTED AT APCER = 5% AND 10%.

Test data		SelfMAD++ [8]			SPL-MAD [2]			MAD-DDPM [1]			SRL-MAD (Ours)		
		EER (%)	BPCER@APCER (%)		EER (%)	BPCER@APCER (%)		EER (%)	BPCER@APCER (%)		EER (%)	BPCER@APCER (%)	
FERET-Morph [6]	FaceMorpher	0.43	0.38	0.19	20.42	65.39	40.85	27.98	99.24	95.27	16.47	37.89	25.62
	OpenCV	0.43	0.19	0.19	25.71	77.02	57.45	31.38	99.24	95.27	5.51	8.47	1.57
	StyleGAN2	1.49	2.08	0.19	25.33	82.34	62.06	32.14	99.24	95.27	1.92	0.60	0.45
FRLL-Morph [6]	AMSL	0.99	0.05	0.05	3.26	36.14	0.50	27.13	98.99	94.94	0.00	0.49	0.49
	FaceMorpher	0.00	0.26	0.26	1.03	0.99	0.99	10.40	99.14	95.19	0.12	0.49	0.49
	OpenCV	0.00	0.00	0.00	1.88	5.45	0.50	13.76	99.02	95.17	5.78	6.86	3.43
	StyleGAN2	1.48	2.13	0.25	14.65	64.85	32.18	14.32	99.10	95.17	13.01	25.98	15.20
	Webmorph	3.45	1.64	3.94	6.39	30.69	11.39	30.30	99.10	95.09	0.00	0.49	0.49
MorDIFF [7]		0.00	0.00	0.00	9.78	54.95	23.27	2.80	99.20	96.40	0.00	0.00	0.00
Average		1.03	1.20	0.62	14.27	46.20	28.91	23.58	99.12	95.60	<u>4.76</u>	<u>9.03</u>	<u>5.30</u>

TABLE II

COMPARISON WITH STATE-OF-THE-ART DISCRIMINATIVE S-MAD METHODS TRAINED ON SMDD. WE USE METRICS SUCH AS EER (%), AND BPCER@APCER (%), REPORTED AT FIXED APCER = 5% AND APCER = 10%.

Test data		MixFaceNet [25]			PW-MAD [13]			MADatation [9]			SRL-MAD (Ours)		
		EER (%)	BPCER@APCER (%)		EER (%)	BPCER@APCER (%)		EER (%)	BPCER@APCER (%)		EER (%)	BPCER@APCER (%)	
FERET-Morph [6]	FaceMorpher	19.99	78.26	53.69	8.36	89.79	42.34	7.08	15.69	9.26	16.47	37.89	25.62
	OpenCV	20.62	85.26	53.88	16.09	97.92	74.48	15.00	29.36	21.40	5.51	8.47	1.57
	StyleGAN2	36.64	96.79	87.33	20.91	98.30	81.47	21.66	46.31	36.86	1.92	0.60	0.45
FRLL-Morph [6]	AMSL	31.03	75.72	65.56	4.43	26.53	4.64	27.94	77.89	61.43	0.00	0.49	0.49
	FaceMorpher	8.37	23.63	11.51	1.97	3.18	1.72	1.47	1.31	0.74	0.12	0.49	0.49
	OpenCV	9.85	65.11	15.89	2.46	4.10	1.80	1.47	2.38	1.39	5.78	6.86	3.43
	StyleGAN2	38.92	94.44	84.94	17.24	49.51	34.62	19.61	47.55	32.57	13.01	25.98	15.20
	Webmorph	31.03	96.72	78.38	9.85	23.59	12.29	24.51	56.76	42.67	0.00	0.49	0.49
MorDIFF [7]		5.88	8.80	6.00	11.27	43.40	18.40	24.02	50.20	38.00	0.00	0.00	0.00
Average		22.48	69.41	50.80	<u>10.29</u>	48.48	30.20	15.86	<u>36.38</u>	<u>27.15</u>	4.76	9.03	5.30

V. RESULTS

We decided to benchmark our method against both one-class and supervised MAD models. For consistency with recent MAD benchmarking studies and to ensure fair comparison across methods, we report baseline results following the evaluation protocol adopted in [8]. It is worth mentioning that unsupervised one-class methods operate under an open-set assumption, where only bona fide samples are available during training. This setting is well aligned with real deployments, where the space of possible morph generation pipelines is effectively unbounded. The main difficulty in such open-set scenarios is therefore the lack of knowledge of the attack types during training. In contrast, supervised approaches rely on explicit supervision on attack samples but are often biased toward the morph type present in training and degrade under cross-morph or dataset shifts. Despite this difference, our method achieves competitive results with one-class approaches and, in several comparisons against supervised approaches, achieves superior performance, confirming its potential for open-set morphing attack detection.

Comparison to One-Class and Self-Supervised MAD Models. We compare our approach with SPL-MAD [2] and MAD-DDPM [1], which represent recent baselines in one-class S-MAD, as well as SelfMAD++ [8], which is a self-supervised S-MAD method. While these methods differ in their modeling strategies and input representations, all approaches are trained under the same one-class constraint on SMDD and evaluated on the same target benchmarks using standard metrics. Results presented in Table I show that our method outperforms SPL-

MAD [2] and MAD-DDPM [1] across all datasets, achieving a lower average EER of 4.76% compared to 14.27% and 23.58%, respectively. In particular, MAD-DDPM [1] exhibits near-saturated BPCER values across most test sets, which indicates limited performance when forced to operate at a low APCER. Regarding SelfMAD++ [8], it achieves low error rates, which can be attributed to its use of explicit self-supervised artifact generation, which provides access to attack-related cues during training. Our method achieves its strongest gains on OpenCV and StyleGAN2 attacks from FERET-Morph [6], obtaining EERs of 5.51% and 1.92%, respectively; however, performance degrades in more challenging scenarios, such as FRLL StyleGAN2. Nevertheless, our method remains competitive and consistently improves upon SPL-MAD [2] and MAD-DDPM [1]. Overall, these results highlight the effectiveness of modeling frequency-domain residuals as a domain-agnostic cue for morph detection.

Comparison to Supervised MAD Models. We further compare our method with state-of-the-art supervised approaches, including MixFaceNet [25], PW-MAD [13], and MADatation [9], as reported in Table II. These methods rely on supervised training with morph samples, which can improve performance when the training and testing morph distributions are similar. However, such models may also overfit to the morph generation process observed during training and degrade under cross-dataset evaluation. Despite this, our method achieves the best overall performance, with an average EER of 4.76% compared to 10.29% for PW-MAD [13] and 15.86% for MADatation [9], and dramatically lower BPCER values at both

APCER operating points (9.03% and 5.30%, respectively). On AMSL and Webmorph attacks from FRLL-Morph [6], our method achieves EERs of 0.00%, while supervised approaches exhibit significantly higher EERs. This observation supports the advantage of one-class models, which avoid overfitting to specific morph attacks. Overall, the proposed approach advances the state of the art among one-class MAD methods and surpasses several supervised models, despite relying solely on bona fide training data.

VI. CONCLUSION

This paper describes a novel one-class method for detecting face morphing attacks. Our method uses residual Fourier features and learns a compact latent representation for anomaly scoring. Through extensive benchmarking across multiple datasets and comparisons with other one-class and supervised approaches, we show that our method achieves competitive performance and even surpasses common supervised approaches, showing the potential of using residual Fourier features as a proxy for morphing artifacts. Future work includes training on larger datasets, such as CASIA-WebFace, to capture more bona fide patterns and to further improve open-set robustness.

ACKNOWLEDGMENTS

This work was funded by the Portuguese Foundation for Science and Technology (FCT) under the PhD grant 2025.03420.BD. This work is supported by NOVA LINS (UID/04516/2025) with the financial support of FCT.IP (DOI: <https://doi.org/10.54499/UID/04516/2025>). It was also supported, when eligible, by national funds through FCT under project UID/50008/2025 – Instituto de Telecomunicações (DOI: <https://doi.org/10.54499/UID/50008/2025>).

REFERENCES

- [1] M. Ivanovska and V. Štruc, “Face morphing attack detection with denoising diffusion probabilistic models,” in *11th International Workshop on Biometrics and Forensics (IWBF)*, pp. 1–6, IEEE, 2023.
- [2] M. Fang, F. Boutros, and N. Damer, “Unsupervised face morphing attack detection via self-paced anomaly detection,” in *IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–11, IEEE, 2022.
- [3] L. Colbois and S. Marcel, “Evaluating the effectiveness of attack-agnostic features for morphing attack detection,” in *IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–9, IEEE, 2024.
- [4] 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 *2019 IEEE 10th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pp. 1–5, IEEE, 2019.
- [5] D. J. Paulo, H. Proença, and J. C. Neves, “FD-MAD: Frequency-Domain Residual Analysis for Face Morphing Attack Detection,” *arXiv preprint arXiv:2601.20656*, 2026.
- [6] E. Sarkar, P. Korshunov, L. Colbois, and S. Marcel, “Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks,” *arXiv preprint arXiv:2012.05344*, 2020.
- [7] 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 *11th International Workshop on Biometrics and Forensics, IWBF*, pp. 1–6, IEEE, 2023.
- [8] M. Ivanovska Preskar, L. Todorov, P. Peer, and V. Štruc, “Selfmad++: self-supervised foundation model with local feature enhancement for generalized morphing attack detection,” *Information fusion*, vol. 127, no. Part C, 103921, 2026.
- [9] E. Caldeira, G. Ozgur, T. Chettaoui, M. Ivanovska, P. Peer, F. Boutros, V. Struc, and N. Damer, “Madation: Face morphing attack detection with foundation models,” in *Proceedings of the Winter Conference on Applications of Computer Vision*, pp. 1650–1660, 2025.
- [10] U. Scherhag, C. Rathgeb, and C. Busch, “Morph detection from single face image: A multi-algorithm fusion approach,” in *Proceedings of the 2018 2nd International Conference on Biometric Engineering and Applications*, pp. 6–12, 2018.
- [11] 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 *5th International Workshop on Biometrics and Forensics*, pp. 1–6, IEEE, 2017.
- [12] P. C. Neto, T. Gonçalves, M. Huber, N. Damer, A. F. Sequeira, and J. S. Cardoso, “Orthomad: Morphing attack detection through orthogonal identity disentanglement,” in *International Conference of the Biometrics Special Interest Group (BIOSIG)*, pp. 1–5, IEEE, 2022.
- [13] 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 *International Symposium on Visual Computing*, pp. 291–304, Springer, 2021.
- [14] 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 *31st European Signal Processing Conference (EUSIPCO)*, pp. 955–959, IEEE, 2023.
- [15] L. Dargaud, M. Ibsen, J. Tapia, and C. Busch, “A principal component analysis-based approach for single morphing attack detection,” in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 683–692, 2023.
- [16] H. Zhang, R. Ramachandra, K. Raja, and C. Busch, “Generalized single-image-based morphing attack detection using deep representations from vision transformer,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1510–1518, 2024.
- [17] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al., “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning*, pp. 8748–8763, 2021.
- [18] M. Ivanovska, L. Todorov, N. Damer, D. K. Jain, P. Peer, and V. Štruc, “Selfmad: Enhancing generalization and robustness in morphing attack detection via self-supervised learning,” in *IEEE 19th International Conference on Automatic Face and Gesture Recognition (FG)*, 2025.
- [19] R. Ramachandra, S. Venkatesh, K. Raja, and C. Busch, “Detecting face morphing attacks with collaborative representation of steerable features,” in *Proceedings of 3rd International Conference on Computer Vision and Image Processing: CVIP 2018, Volume 1*, pp. 255–265, Springer, 2019.
- [20] S. Venkatesh, R. Ramachandra, K. Raja, L. Spreeuwiers, R. Veldhuis, and C. Busch, “Morphed face detection based on deep color residual noise,” in *9th International Conference on Image Processing Theory, Tools and Applications (IPTA)*, pp. 1–6, IEEE, 2019.
- [21] M. Huber, P. C. Neto, A. F. Sequeira, and N. Damer, “FX-MAD: Frequency-Domain Explainability and Explainability-Driven Unsupervised Detection of Face Morphing Attacks,” in *Proceedings of the Winter Conference on Applications of Computer Vision*, pp. 766–776, 2025.
- [22] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [23] J. E. Tapia and C. Busch, “Face feature visualisation of single morphing attack detection,” in *11th International Workshop on Biometrics and Forensics (IWBF)*, pp. 1–6, 2023.
- [24] I. Avcibas, “Morphed face detection with wavelet-based co-occurrence matrices,” *IEEE Signal Processing Letters*, vol. 31, pp. 1344–1348, 2024.
- [25] 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,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1606–1617, 2022.
- [26] M. Huber, F. Boutros, A. T. Luu, K. Raja, R. Ramachandra, N. Damer, P. C. Neto, T. Gonçalves, A. F. Sequeira, J. S. Cardoso, et al., “Syn-mad 2022: Competition on face morphing attack detection based on privacy-aware synthetic training data,” in *IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–10, IEEE, 2022.
- [27] “Information technology — biometric presentation attack detection — part 3: Testing and reporting,” International Standard ISO/IEC 30107-3:2023, 2023.