

FD-MAD: Frequency-Domain Residual Analysis for Face Morphing Attack Detection

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

¹ University of Beira Interior, Portugal ² NOVA LINCS ³ IT: Instituto de Telecomunicações

Abstract—Face morphing attacks present a significant threat to face recognition systems used in electronic identity enrolment and border control, particularly in single-image morphing attack detection (S-MAD) scenarios where no trusted reference is available. In spite of the vast amount of research on this problem, morph detection systems struggle in cross-dataset scenarios. To address this problem, we introduce a region-aware frequency-based morph detection strategy that drastically improves over strong baseline methods in challenging cross-dataset and cross-morph settings using a lightweight approach. Having observed the separability of bona fide and morph samples in the frequency domain of different facial parts, our approach 1) introduces the concept of residual frequency domain, where the frequency of the signal is decoupled from the natural spectral decay to easily discriminate between morph and bona fide data; 2) additionally, we reason in a global and local manner by combining the evidence from different facial regions in a Markov Random Field, which infers a globally consistent decision. The proposed method, trained exclusively on the synthetic morphing attack detection development dataset (SMDD), is evaluated in challenging cross-dataset and cross-morph settings on FRLM-Morph and MAD22 sets. Our approach achieves an average equal error rate (EER) of 1.85% on FRLM-Morph and ranks second on MAD22 with an average EER of 6.12%, while also obtaining a good bona fide presentation classification error rate (BPCER) at a low attack presentation classification error rate (APCER) using only spectral features. These findings indicate that Fourier-domain residual modeling with structured regional fusion offers a competitive alternative to deep S-MAD architectures.

I. INTRODUCTION

Face morphing attacks represent a significant threat to face recognition systems used in electronic identification (eID) enrolment and border control. When the facial images of two individuals are blended into a single plausible face, malicious actors can obtain documents that are accepted under both identities. These attacks exploit the tolerance of modern recognition systems to intra-class variation and have been demonstrated through classical landmark-based warping [21], [33], GAN-based synthesis [38], [14], [40], and, more recently, diffusion models [12], [5], [3], [4], [41]. This threat has driven extensive research on single-image morphing attack detection (S-MAD), where the goal is to determine whether a given face image is bona fide or morphed, in line with emerging standards such as ISO/IEC 30107-3 [1]. The earliest S-MAD approaches used handcrafted texture and color features in combination with classical classifiers [34], [32]. More recent research [13], [24], [28], [7] has focused on developing deep learning architectures trained on extensive datasets of bona fide and morphed images. Additionally, transformer-based and foundational models have been used,




Fig. 1: **Comparison between frequency-domain data of a bona fide face and a morph sample.** The key observation is that the morphs typically differ from the bona fide samples in the mid-high frequencies, which is used as the insight for the proposed method.

including [39] and [8]. Despite substantial progress, these methods often require complex architectures, which increase computational demands and complicate deployment in large-scale and time-sensitive environments. More recent advances include unsupervised and self-supervised S-MAD methods, which reduce the reliance on labelled data and aim to improve robustness to unseen attacks by modelling bona fide data distributions or leveraging self-supervised pretext tasks [19], [26], [25]. However, achieving strong cross-morph and cross-dataset generalization remains particularly challenging. Models trained on a given morphing pipeline often suffer performance drops when evaluated on unseen morphing attacks. Recent work has begun to exploit frequency-domain information and frequency-aware data augmentation to address this limitation. For example, Huber *et al.* [25] operate in the differential morphing attack detection (D-MAD) setting, while other approaches mainly regularize deep learning models by creating synthetic frequency perturbations [26]. Overall, frequency-domain cues have so far received little attention as primary features for S-MAD, despite their potential to capture global and local spectral regularities in the Fourier domain. Hence, we propose a residual frequency-domain S-MAD framework that directly targets cross-morph generalization by obtaining invariant global-local features in the Fourier domain, without relying on heavy architectures prone to overfitting and to capturing

dataset or morph-specific biases.

In this work, we introduce a global–local framework that explicitly models deviations from the natural power-law spectral decay. Our design is motivated by extensive evidence that natural images exhibit characteristic magnitude spectra that approximately follow a $f^{-\alpha}$ law [37], [6], [22], while synthetic and manipulated imagery often deviates from this behaviour [18], [9], [11], [35]. Based on these observations, we represent each face image in a frequency domain residual space that highlights deviations from the natural spectral trend (Figure 1). At the global level, we obtain a radial log-magnitude spectrum, estimate an image-specific spectral baseline, and treat the difference as a compact residual descriptor, which is then fed to a classifier. At the local level, we apply the same residual construction to four semantic regions (left and right eye, nose, and mouth), obtaining region-wise descriptors. To capture structured dependencies between these regions, such as local artefacts affecting only a subset of them, we use a pairwise Markov Random Field (MRF) over the region labels, to encourage coherent decisions across neighboring regions and avoid scattered or inconsistent labels. This approach provides a probabilistic global–local morphing score that is both computationally efficient and interpretable.

We evaluate the proposed method in a challenging cross-dataset and cross-morph scenario. Models are trained on privacy-friendly synthetic data from SMDD [13] and tested on FRLL-Morph [33] and MAD22 [24], which include a broad range of landmark-based and GAN-based attacks [40], [27], [38], [30]. Our approach consistently outperforms or matches several recent supervised methods on FRLL-Morph [33], and achieves competitive results on MAD22 [24] while remaining significantly more lightweight than transformer-based and multi-task architectures. Ablation studies demonstrate that global and local Fourier residuals provide complementary information and that MRF-based fusion is essential for addressing both classical and GAN-based morphs. t-SNE visualizations further reveal well-structured, attack-sensitive embeddings in the frequency domain, reinforcing the interpretability of our approach.

In summary, our main contributions are as follows:

- A novel S-MAD framework that operates entirely in the Fourier domain, modelling deviations from the natural $f^{-\alpha}$ spectral law using compact global and local residual features, that amplify morphing artefacts.
- An MRF-based fusion of semantic facial regions, which captures structured dependencies between local decisions and improves robustness to heterogeneous, region-specific morphing artefacts, including those from GAN-based attacks.
- Extensive cross-morph experiments on FRLL-Morph [33] and MAD22 [24], demonstrating that our approach achieves state-of-the-art or highly competitive performance compared to recent supervised S-MAD methods, while remaining computationally efficient, making it suitable for real-world deployment.

The code will be made publicly available.

II. RELATED WORK

Face Morphing. Previous face morphing techniques demonstrated that blending the geometry and texture of two identities can produce a synthetic face image that matches both contributors in automatic face verification systems. [20] was among the first to show the vulnerability of passport-style biometric systems to such attacks, motivating substantial research into morphing generation and detection. Subsequent works improved the realism and controllability of morphs, ranging from landmark-based image warping [21] to more advanced methods which interpolate the latent representations like GAN-based morphs [40], [10] and more recently diffusion-based morphs [3], [4], [41], [12], [5].

Morphing Attack Detection. Morphing Attack Detection (MAD) is typically categorized into single-image MAD (S-MAD) and differential MAD (D-MAD). D-MAD, which compares a live image with a trusted reference, achieves strong accuracy but is less relevant for passport issuance scenarios where no live comparison is available. S-MAD is therefore a more challenging and widely studied setting. Hence, several studies have developed S-MAD systems [7], [32], [28], [15], [39], [34], [8], [24], [19], [26], in order to detect whether the investigated image is a morph based only on its features.

Early S-MAD approaches relied on handcrafted texture descriptors [34] such as LBP, BSIF, or SIFT histograms, exploiting inconsistencies in facial regions introduced by the morphing warp. However, as morphing attacks became more diverse, these features lacked robustness against unseen morphing methods. Deep learning–based S-MAD approaches have significantly improved performance on controlled benchmarks; however, they exhibit limited robustness under cross-morphing scenarios. Although approaches such as IDistill [7] or MixFaceNet-MAD [13] achieve high accuracy on SMDD, they degrade significantly on FRLL-morphs such as AMSL, indicating a strong dependence on the texture statistics of the training dataset. Even MADation [8], which uses CLIP’s large-scale pretraining and lightweight LoRA [23] fine-tuning, illustrates these limitations. Even though MADation [8] achieves good performance in a large portion of the MAD22 and MorDIFF evaluation settings, the authors report notable performance drops for some unseen morphing methods (e.g., MIPGAN-I/II when trained on SMDD). This reinforces that even foundation models used in S-MAD methods remain sensitive to cross-morph scenarios. These findings evidence that current deep learning S-MAD systems, while powerful, are inherently vulnerable to dataset bias, texture bias overfitting, and poor cross-morph generalization, especially when the training data come from a single dataset such as SMDD [13].

In [16], the authors proposed analyzing texture artifacts using PCA projections, highlighting the utility of interpretable, low-dimensional texture representations rather than features derived from black-box models. However, this approach relies only on global image descriptors. It does not model region-specific inconsistencies or dependencies among facial

areas—artifacts that are often amplified by the morphing process. Hence, our method extends this family by combining global descriptors with local spectral cues and structured probabilistic reasoning.

Frequency Artifacts in Generated Morphs. Extensive research in image forensics demonstrates that generative models and image warping techniques leave characteristic traces in the frequency domain [18], [11], [9], [35]. Furthermore, these frequency irregularities are not limited to deepfakes or GAN/diffusion synthesis; face morphing operations also distort natural spectral statistics, and these irregularities have been studied in some prior works [25], [36], [2]. It is worth mentioning that [25] studies frequency cues in the D-MAD setting, which assumes access to a trusted reference image; therefore, it is not directly comparable to our S-MAD scenario and is cited here primarily to support the relevance of frequency domain artifacts. Face morphing combines geometric warping with pixel-level blending, both of which disrupt local textures, smooth natural edges, and attenuate high-frequency content—particularly around manipulated regions such as eyes, nose, or mouth. Recent works in morphing forensics further demonstrate that, because morphs introduce artifacts at high frequencies, the altered images deviate from the natural $f^{-\alpha}$ power-law behavior [37], [6], [22]. These observations motivate the use of frequency domain features, which can generalize more robustly across morphing pipelines than spatial CNN features, which are prone to overfitting dataset-specific textures.

III. METHODOLOGY

In order to enhance the detection of face morphing attacks, we propose a lightweight and interpretable approach based on the Fourier frequency domain. Our method is based on the observation that natural images follow a predictable spectral decay, typically following a $f^{-\alpha}$ distribution of magnitude across spatial frequencies [37], [6], [22], where α is the spectral decay exponent (i.e., the slope of the radial spectrum in log-log space), which varies with image content but is typically stable for natural images. In contrast, face morphing attacks tend to distort this natural behavior by introducing artifacts and irregularities in the frequency spectrum, particularly in the mid and high-frequency bands. We obtain both global and local (region-wise) residual Fourier magnitude features, removing the estimated natural $f^{-\alpha}$ spectral trend to emphasize morphing artifacts, and train a classifier on the resulting residuals. To integrate region-specific evidence while capturing structured region-level dependencies, we introduce a Markov Random Field (MRF) model whose unary potentials derive from local classifiers and whose pairwise potentials penalize label disagreements, encouraging globally consistent region labels while still allowing the unary terms to capture localized morphing artifacts. The final morphing score is obtained by aggregating the global classifier output and the MRF local decisions within a unified decision function, thereby using complementary frequency cues. An overview of the proposed framework is illustrated in Figure 2.

A. Global Fourier Residual Features

Given an RGB image $I \in \mathbb{R}^{H \times W \times 3}$ we obtain, for each channel, the two-dimensional discrete Fourier transform (DFT):

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

where $\mathcal{F}(\cdot)$ denotes the 2D DFT mapping the spatial signal $I(x, y)$ to frequency coordinates (u, v) . We then compute its centered log-magnitude spectrum as follows:

$$S(u, v) = \log(1 + |F(u, v)|). \quad (2)$$

To obtain a compact and rotation-invariant representation, we split the magnitude spectrum into K concentric radial frequency bands. Let R_k denote the set of frequency coordinates in the k -th ring. We calculate the azimuthally-averaged band energy as:

$$b_k = \frac{1}{|R_k|} \sum_{(u, v) \in R_k} S(u, v), \quad (3)$$

forming a 1D global spectral profile $b = [b_1, \dots, b_K]$. The number of radial frequency bands K is determined by the maximum attainable radius in the Fourier domain, that is, $K = \frac{N}{2}\sqrt{2}$, for an image of size $N \times N$. Additionally, we decided to keep the entire spectrum (no low or high-frequency bands are discarded) to avoid discarding potentially discriminative frequency content.

Power-Law Spectral Baseline Removal. Natural images approximately follow a power-law spectral decay $S(f) \propto f^{-\alpha}$, which is linear in log-log space. We therefore fit a line to the radial log-magnitude profile in log-frequency space:

$$\log S(f) = a + b \log f, \quad (4)$$

where a and b are the least squares intercept and slope of the fitted power-law baseline. That is, under the assumption $S(f) = C f^{-\alpha}$, we have $a = \log C$ and $b = -\alpha$. We define $f \in \{f_1, \dots, f_K\}$ as the discrete radial frequencies associated with the K concentric rings. This yields the estimated log baseline $\widehat{\log S(f)} = a + b \log f$. The residuals are then obtained as:

$$r(f) = \log S(f) - \widehat{\log S(f)}. \quad (5)$$

In the residual space, morphs are expected to distinguish from bona fide images, due to the irregularities of synthetic images in the mid and high frequencies (Figure 3).

To further remove redundancy, the RGB residual profiles are concatenated, standardized, and reduced using PCA, obtaining x_r . Subsequently, these features are then provided to a classifier to distinguish between bona fide and morphs, which outputs a global score $s_{\text{global}} = p(z_r = 1 \mid x_r)$, where $z_r \in \{0, 1\}$ denotes morph (0) or bona fide (1).

B. Local Region-Wise Spectral Descriptors

Global Fourier residual features mostly capture global trend inconsistencies, while morphing artifacts often appear in localized and heterogeneous regions. For instance, warping boundaries and interpolation errors introduce fine-grained irregularities around semantic regions such as the




Fig. 2: **Overview of the proposed Fourier-domain morphing detector.** The input face is transformed to the frequency domain and summarized into (top) global azimuthally averaged FFT magnitude profiles used by an SVM classifier to obtain a global score (s_{global}), and (bottom) local profiles extracted from the left eye, right eye, nose, and mouth, whose region-wise classifiers are fused by a Markov Random Field (MRF) to produce a local score (s_{local}). The global and local scores are then combined through a weighted fusion to obtain the final morphing detection score.




Fig. 3: **Residual radial-frequency profiles of a bona fide image and its corresponding morph.** Compared to the bona fide image, the morph exhibits systematic deviations, particularly in the mid-to-high frequency bands, consistent with blending and interpolation artifacts introduced during morphing.

eyes, nose, and mouth. Thus, to complement the global representation and target localized distortions, we extract region-wise Fourier residual features from facial subregions, yielding more subtle morphing cues that may be hidden or averaged in the global features.

For this, a face landmark detector is used to localize key facial areas and crop semantic region patches. The resulting patches are resized to a fixed resolution and processed with

the same pipeline as in Section III-A to obtain a region-specific Fourier feature vector x_r .

A logistic regression C_r is then trained per region, yielding posterior probabilities:

$$p_r(z_r = 1 | x_r). \quad (6)$$

We convert these into unary potential terms,

$$\psi_r(z_r) = -\log p_r(z_r | x_r). \quad (7)$$

These unary potentials form the basis of our structured fusion model, serving as input to the MRF; they are obtained from the region-wise classifiers. We chose to use the log-domain instead of raw probabilities for two reasons. First, multiplying many small probabilities can lead to numerical underflow, whereas log-space converts these products into simple sums that remain well-conditioned. Second, the MRF energy function is additive, and log-potentials allow the unary and pairwise terms to be combined in a coherent linear form, enabling exact inference and an appropriately weighted contribution from each region. Thus, the log-domain formulation ensures both computational stability and a principled probabilistic interpretation within the structured fusion model.

C. Structured Fusion with Markov Random Field

We use a pairwise MRF to obtain a consensus across different region estimates, which would not be possible by simply averaging local predictions. This would assume conditional independence across regions. In practice, bona fide faces tend to produce globally consistent decisions across

regions, whereas morphs often exhibit heterogeneous, region-level cues: for example, local blending artifacts may strongly affect the mouth and one eye while leaving other areas visually plausible. To capture this behavior, we use pairwise potentials that penalize label disagreements, encouraging globally consistent region labels while still allowing the unary terms to capture local morphing artifacts. As a result, the MRF serves as a context-aware fusion mechanism that exploits dependencies between regions rather than treating them independently.

MRF Formulation. Let $z = (z_1, \dots, z_R)$ denote the unknown binary labels of R facial regions. We build a fully connected MRF with a potential function defined as:

$$E(z) = \sum_{r=1}^R \psi_r(z_r) + \sum_{(i,j) \in \mathcal{E}} \psi_{ij}(z_i, z_j), \quad (8)$$

where $\mathcal{E}$ denotes the set of region pairs, $\mathcal{E} = \{(i, j) \mid 1 \leq i < j \leq R\}$.

The unary potentials $\psi_r(z_r)$ are defined in Sec. III-B and reflect the region-specific spectral evidence. Additionally, to model structured inconsistencies between facial regions, we also employ an Ising pairwise term. For each pair of connected regions (i, j) , the pairwise potential is defined as

$$\psi_{ij}(z_i, z_j) = \beta |z_i - z_j|, \quad \beta > 0, \quad (9)$$

where β controls the strength of the agreement constraint between regions. This pairwise potential penalizes label disagreements and thus encourages smooth, globally consistent labelings across regions. This allows the unary terms to express local morphing evidence while the MRF regularizes overly fragmented decisions.

Exact Inference. Since the number of regions R is limited when analyzing human faces, exact inference is tractable. Hence, we enumerate all 2^R configurations and calculate

$$P(z|x) = \frac{\exp(-E(z))}{\sum_{z' \in \{0,1\}^R} \exp(-E(z'))}. \quad (10)$$

As illustrated in Figure 4, although the resulting inference time grows exponentially with R , this does not pose a practical limitation for face morphing detection, since faces admit only a small number of meaningful semantic regions and region-based methods, including ours, operate in this setting. Finally, we output the expected fraction of bona fide labels:

$$s_{\text{local}} = \mathbb{E}_{P(z|x)} \left[\frac{1}{R} \sum_{r=1}^R z_r \right] = \sum_{z \in \{0,1\}^R} \left(\frac{1}{R} \sum_{r=1}^R z_r \right) P(z|x). \quad (11)$$

This constitutes the local morphing detection score.

D. Global-Local Fusion

The global score s_{global} (holistic spectral deviations) and the MRF score s_{local} (structured region-level evidence) provide complementary information. Therefore, we combine




Fig. 4: **Exact MRF inference time versus number of facial regions R .** While the complexity scales exponentially with R due to full enumeration, faces admit only a limited number of meaningful regions, and inference remains tractable even for conservative upper bounds on R .

them using a weighted fusion:

$$s_{\text{fused}} = \lambda s_{\text{global}} + (1 - \lambda) s_{\text{local}}, \quad (12)$$

where λ is tuned on the SMDD validation set. This fusion leverages both global robustness and region-specific structural coherence, yielding improved cross-morph generalization.

IV. EXPERIMENTAL SETUP

To assess the generalization capability of the proposed method, we rely on three morphing datasets, SMDD [13], FRLL-Morph [33], and MAD22 [24]. SMDD is a synthetic StyleGAN2-based development dataset with OpenCV morphs, while FRLL-Morph and MAD22 are FRLL-based benchmarks containing landmark and GAN-based morphs and are identity-disjoint from SMDD. Furthermore, we use MorDIFF [12], an extension of MAD22, which provides additional FRLL-based diffusion morphs, enabling evaluation under stronger and more realistic attacks.

A. Evaluation Metrics

All experiments follow the ISO/IEC 30107-3 [1] standard for biometric presentation attack detection. We report the Attack Presentation Classification Error Rate (APCER), defined as the proportion of morphing attacks incorrectly classified as bona fide, and the Bona Fide Presentation Classification Error Rate (BPCER), defined as the proportion of bona fide presentations incorrectly classified as attacks. We also report the Equal Error Rate (EER), corresponding to the operating point at which APCER equals BPCER. In line with the ISO/IEC 30107-3 [1] recommendations and prior S-MAD work [15], [7], [28], we additionally report BPCER at fixed APCER thresholds of 1% and 20%. Reporting BPCER@APCER is particularly informative from a security perspective, as it quantifies the false rejection rate under a controlled false acceptance rate. Since minimizing APCER is critical to prevent morphing attacks from being

falsely accepted, BPCER@APCER directly reflects the cost of enforcing a given security level.

B. Evaluation Protocols

The MAD22 benchmark [24] and SMDD [13] provide standardized experimental setups for single morphing attack detection, which include multiple morphing tools. Building on these protocols, we adopt standard cross-dataset and cross-morph settings commonly used in previous S-MAD studies [15], [7], [28], [13], [8]. For the first protocol, we train on SMDD [13] and test on FRLL-Morph [33], which yields a combined cross-dataset and cross-morph scenario (OpenCV-based synthetic training vs. multiple morphing tools on FRLL). For the second protocol, we also train on SMDD and test on the MAD22 benchmark [24], which includes both landmark and GAN-based morphs. For completeness, we further report results on the MorDIFF [12] extension, which provide challenging diffusion-based morphs built from the same FRLL identities. Unless stated otherwise, all results are obtained from the officially provided aligned images and reported using the same measures, enabling direct comparison with previously published results. Overall, this experimental design covers both cross-dataset and cross-morph generalization.

V. EXPERIMENTS AND RESULTS

In this section, we analyze the performance of our method across different benchmark datasets. We compare our method against state-of-the-art supervised methods and highlight its robust performance. We also provide ablations to quantify the contribution of global and local Fourier features.

Implementation Details. We conduct the experiments on NVIDIA GeForce RTX 5090 with 32GB of VRAM, and the framework utilized was Pytorch [29]. To ensure consistency across experiments, Fourier features were extracted using images of uniform size 500x500, as previous works [16], [26] suggest that higher-resolution images typically preserve more morphing artifacts. For Fourier feature extraction, we use a batch size of 128, and for the global features, a support vector machine (SVM) with an RBF kernel was trained. In order to obtain the semantic face regions, we use the RETINAFace [17] implementation, which provides 106 facial landmarks per face. For each region r , we select a fixed subset of landmark indices and crop a bounding box around the corresponding points with a small margin. Regarding the pairwise interactions in the MRF, we use a fixed parameter $\beta = 0.9$, tuned on SMDD and held constant across evaluations, and set the final fusion weight to $\lambda = 0.6$. Regarding the dataset proportion of positive and negative samples, we use the same number of bona fide images and morphs for the training phase.

A. Baseline Comparison

We compare our method with previous S-MAD works on both evaluation protocols.

Table I reports results for the first protocol. Across all attacks, our approach achieves the lowest average EER

(1.85%) and clearly improves over strong supervised baselines such as IDistill [7] and MADation (ViT-B) [8]. The gains are particularly pronounced for classical warping-based morphs (OpenCV, FaceMorph, WebMorph, AMSL), where our method achieves almost perfect detection (EER below 0.2% and BPCER=0% at both operating points). This behaviour is consistent with how these morph sets are constructed and with prior observations on landmark-based morphing pipelines. FRLL-Morphs [33] builds upon the AMSL protocols and provides predefined morph pairs together with classical landmark-based morphing tools, which are known to introduce characteristic artefacts, such as ghosting and local blending inconsistencies. As a result, these attacks are comparatively easier for S-MAD methods to exploit, even though they remain highly effective at attacking face recognition systems. This interpretation is further supported by our ablation study in Tables III and IV. When using only the local branch, the EER on FRLL-Morphs [33] drops to 0% for all landmark-based attacks, while remaining higher for the StyleGAN2 attack. GAN-based morphs tend to suppress local artefacts, making them inherently more challenging to separate from bona fide images, which explains the gap in performance between the landmark-based and StyleGAN2 attacks. In contrast, IDistill [7] explicitly targets identity entanglement by estimating whether two identities are present in the image (using $id_1 \cdot id_2$), making it less dependent on low-level artefacts and therefore particularly effective for StyleGAN2 morphs. Nevertheless, our global-local Fourier model is competitive and still yields lower average EER and BPCER than deep learning baselines, despite using only lightweight spectral descriptors rather than full CNN backbones.

In contrast, the MAD22 [24] benchmark is explicitly designed to yield challenging and realistic morphs. For each FRLL attack, the authors first obtain the embeddings with a strong face recognition model and then select only the most similar identity pairs as candidates for morphing. These pairs are then morphed using multiple pipelines, including both landmark-based approaches and high-quality GAN-based methods, which substantially reduce visible artefacts while preserving the identity blend. This protocol construction results in a more difficult detection scenario, explaining the higher error rates on MAD22 [24] presented in Table II compared to FRLL-Morphs [33]. Beyond the original MAD22 attacks, we additionally report results on MorDIFF [12], and, since diffusion-based morph attacks are recent, some earlier baselines do not report MorDIFF results; therefore, MorDIFF is excluded from the average to avoid unfair comparisons, but we deliberately include it to assess robustness against modern and more realistic morph generation pipelines. Table II shows that the D-FW [31] variant, which relies on a multi-task framework with additional 3D facial cues, obtains the best average EER. Despite this, our approach ranks second in terms of average EER (6.12% vs. 5.26% for D-FW-CDCN [31]), achieving a BPCER@APCER=1% of 40.10% and BPCER@APCER=20% of 4.22%.

In particular, the proposed method is highly effective

TABLE I: **Results comparison between our method and previous S-MAD solutions when tested on FRLL-Morph and trained on SMDD.** We report Equal Error Rate (EER %) and BPCER (%) at fixed APCER levels of 1% and 20% for different morphing attacks. Specific values unavailable in the original papers are marked with ”-”. The proposed method achieves the lowest average EER and demonstrates strong performance across different morphs.

Method	OpenCV			FaceMorph			WebMorph			AMSL			StyleGan2			Avg.	
	EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER
		1%	20%		1%	20%		1%	20%		1%	20%		1%	20%		
OrthoMAD [28]	0.73	13.74	3.76	0.98	0.98	0.08	15.23	15.23	9.50	14.80	65.05	10.89	6.54	13.74	3.76	7.66	21.75
IDistill [7]	2.46	6.14	0.16	2.05	4.26	0.16	4.01	14.41	0.33	4.00	21.10	2.85	1.96	8.51	0.08	2.90	10.88
MixFaceNet-MAD [13]	4.39	26.47	1.47	3.87	23.53	0.49	12.35	80.39	13.24	15.18	49.51	11.76	8.99	42.16	4.41	8.96	44.41
PW-MAD [15]	2.42	22.06	0.49	2.20	26.47	0.00	16.65	80.39	13.24	15.18	96.57	5.88	16.64	80.39	13.24	10.62	61.18
WB-Avcivas [2]	7.91	-	-	17.11	-	-	19.32	-	-	18.23	-	-	14.87	-	-	15.49	-
Inception-MAD [13]	5.38	38.73	0.98	3.17	30.39	0.49	9.86	53.92	2.94	10.79	72.06	4.90	11.37	72.06	6.86	8.11	53.43
MADation (ViT-B) [8]	2.97	4.41	0.49	1.35	1.47	0.00	3.42	5.88	0.00	3.85	12.07	2.41	17.21	54.50	13.10	5.76	15.67
Ours	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	9.25	24.02	2.45	1.85	4.80

TABLE II: **Results comparison between our method and previous S-MAD solutions when tested on MAD22 [24] and trained on SMDD [13].** Specific values unavailable in the original papers are marked with ”-”. While D-FW-CDCN [31] achieves the lowest average EER, our approach achieves the second-best average EER without relying on heavy 3D-aware multi-task architectures.

Method	OpenCV			FaceMorph			WebMorph			MIPGAN-I			MIPGAN-II			MorDIFF			Avg. [†]		
	EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER		EER	BPCER@APCER	
		1%	20%		1%	20%		1%	20%		1%	20%		1%	20%		1%	20%			
MixFaceNet-MAD [13]	8.33	44.31	1.96	4.60	55.88	0.00	18.20	92.86	15.20	16.70	89.21	11.76	20.62	91.64	20.10	8.50	62.67	2.94	13.69	74.78	9.80
Inception-MAD [13]	7.52	-	-	0.00	-	-	18.00	-	-	10.90	-	-	16.22	-	-	5.30	-	-	-	10.53	-
MorphHRNet [24]	5.69	33.82	1.47	5.90	48.04	1.47	9.80	56.86	3.92	15.30	75.98	11.27	10.41	61.27	2.94	-	-	-	9.42	55.19	4.21
Con-Text Net A [24]	17.48	74.02	16.18	0.00	0.00	0.00	26.20	93.14	31.86	12.30	59.31	6.37	12.91	59.31	5.88	-	-	-	13.78	57.16	12.06
D-FW-MixFaceNet [31]	13.72	-	-	0.10	-	-	10.80	-	-	6.70	-	-	6.61	-	-	-	-	-	7.59	-	-
D-FW-CDCN [31]	0.30	-	-	0.00	-	-	0.00	-	-	11.90	-	-	14.11	-	-	-	-	-	5.26	-	-
MADation (ViT-B) [8]	3.85	23.53	0.00	0.00	0.00	0.00	10.80	51.47	4.41	33.37	94.12	52.94	22.21	84.80	26.47	1.10	1.94	0	14.05	50.78	16.76
Ours	2.54	7.35	1.47	0.00	0.00	0.00	18.40	86.76	17.16	2.45	6.37	0.98	7.21	100.00	1.47	0.00	0.00	0.00	6.12	40.10	4.22

[†] Avg. is calculated across all morphing attacks excluding MorDIFF due to unavailable values which would lead to unfair comparisons.

on the FaceMorph, MIPGAN-I/II and MorDIFF [12] attacks, while the WebMorph attack remains the most difficult across methods. Our method achieves perfect separation on MorDIFF [12], outperforming MADation [8], the strongest competing approach with available results. Beyond quantitative results, Figure 5 shows that diffusion-based morphs form a compact and well-separated cluster away from bona fide samples in the feature space, indicating that the proposed method capture distinctive artefacts introduced by the diffusion process. These results demonstrate that our approach not only generalizes across classical attacks, but also remains effective against emerging diffusion-based morphing techniques. It is worth mentioning that our model achieves this performance without any auxiliary 3D supervision or heavy backbones, highlighting that Fourier-domain residuals are a competitive alternative for S-MAD under strict computational constraints. Although 3D-aware multi-task architectures such as D-FW [31] achieve slightly lower EERs, their higher complexity results in a significant increase in computational demands, making them less suitable for resource-constrained deployments. Moreover, integrating such cues into our framework is an interesting direction for future work, but currently beyond the scope of this paper.

B. Ablation Studies

In order to evaluate the impact of the isolated global and local features, we report results that show the impact of each component.

TABLE III: **Impact of global and local Fourier residual features on S-MAD performance.** EER (%) is reported, and the fusion of global and local features improves S-MAD performance, confirming the complementarity of global and regional cues.

Tested on FRLL-Morph [33]	Global Features (G)	Local Features (L)	G + L
FRLL-OpenCV	0.98	0.00	0.00
FRLL-FaceMorph	0.49	0.00	0.00
FRLL-WebMorph	0.41	0.08	0.00
FRLL-AMSL	2.16	0.00	0.00
FRLL-StyleGan2	6.37	18.09	9.25
Average	2.08	3.63	1.85

Impact of global and local features. Tables III and IV analyze the contribution of global Fourier residuals, local region-wise features, and their fusion.

On the first protocol (Table III), global features alone already provide strong performance (average EER 2.08%), while local features alone are particularly effective for most attacks but struggle on StyleGAN2 (average EER 3.63% dominated by this case). Fusing global and local information consistently improves the average EER to 1.85%, and either matches or outperforms the best single component on four out of five attacks. This confirms that global and local features capture complementary cues: global features are more robust to GAN-based morphs, whereas local residuals

TABLE IV: **Impact of global and local Fourier residual features on S-MAD performance.** EER (%) is reported, and the fusion of global with local features reduces the average EER, demonstrating that combining holistic and structured regional evidence helps achieve robust generalization across diverse morphing attacks.

Tested on MAD22 [24]	Global Features (G)	Local Features (L)	G + L
OpenCV	10.37	6.37	2.54
FaceMorph	9.20	0.00	0.00
WebMorph	45.80	21.57	18.40
MIPGAN-I	1.30	58.70	2.45
MIPGAN-II	33.03	13.24	7.21
MorDIFF	0.00	0.00	0.00
Average	16.62	16.65	5.10

excel at detecting classical morphing pipelines.

On MAD22 [24] (Table IV), the benefit of fusion is even more evident. Global and local features alone yield similar but relatively high average EERs (16.62% and 16.65%), with each component specializing in different attacks (e.g. global features handle MIPGAN-I well, while local features are better for WebMorph and MIPGAN-II). The fused model reduces the average EER to 5.10%, substantially outperforming both individual streams and achieving close-to-best performance for every attack. From this ablation, we conclude that combining general and region-wise Fourier residuals within the MRF is essential for robustness across diverse morphing mechanisms.

C. Feature Visualization

To better understand the behaviour of the Fourier-domain features, we visualize them using t-SNE. Figure 5 shows the t-SNE embeddings of FRLM-Morph [33] and MAD22 [24] global features. The different morphing attacks form well-separated clusters, and bona fide samples occupy a distinct region of the space. Regarding the FRLM-Morph [33] dataset, the AMSL and WebMorph attacks partially overlap, indicating that these pipelines produce similar spectral artefacts, which explains why they are harder to distinguish from each other. Nevertheless, bona fide samples remain clearly separated from all morph clusters, supporting the high detection rates obtained on this dataset.

As for the MAD22 [24] dataset, the distributions of some attacks, particularly WebMorph, overlap more with bona fide samples, revealing a more challenging decision boundary. This overlap correlates with the higher EERs observed for these attacks and highlights MAD22 as a significantly harder benchmark. Overall, the visualizations confirm that the proposed frequency-domain residuals encode discriminative, attack-sensitive structure while still enabling a compact and interpretable representation for S-MAD.

VI. CONCLUSION

This work introduced a novel frequency-domain method for detecting single-image face morphing attacks (S-MAD).




Fig. 5: **t-SNE visualization of frequency-domain global features obtained from different morphing attacks of FRLM-Morph [33] (left) and MAD22 [24] (right).** For MAD22, the WebMorph attack follows a similar distribution, making it a hard one.

Our approach combines global Fourier residual features with region-based spectral features, using a Markov Random Field. By removing the power-law natural trend and thereby obtaining residual features, our approach differs from traditional deep learning approaches, remains interpretable and reduces reliance on dataset-specific biases that hamper generalization across different morphing techniques. Our experiments were carried out on FRLM-Morph and MAD22, and show that combining global and local features helps our method generalize well from the SMDD training data to new morphing attacks. On FRLM-Morph, our approach outperforms or matches recent supervised S-MAD methods, and it performs competitively on MAD22, even though our models are much simpler than multi-task or 3D-aware ones. Then, the ablation studies show that global and local Fourier residuals provide complementary discriminative cues, and that using an MRF to maintain facial region consistency is important for handling both landmark-based and GAN-based morphs. As future work, we plan to expand our framework by adding new cues, such as 3D shape information, within the same frequency-domain and MRF setup, making the method even more robust to new morph types. Furthermore, we could expand it by exploring contrastive learning on Fourier residual embeddings to strengthen bona fide clusters and increase separation from morph samples, potentially improving robustness against challenging attack types that produce less consistently localized artefacts. Our spectral residual modeling might also work well for other biometric systems or attack-detection tasks, especially when computational resources are limited.

REFERENCES

- [1] Information technology — biometric presentation attack detection — part 3: Testing and reporting. International Standard ISO/IEC 30107-3:2023, 2023.
- [2] I. Avcibas. Morphed face detection with wavelet-based co-occurrence matrices. *IEEE Signal Processing Letters*, 31:1344–1348, 2024.
- [3] Z. W. Blasingame and C. Liu. Fast-dim: Towards fast diffusion morphs. *IEEE Security & Privacy*, 22(4):103–114, 2024.
- [4] 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. IEEE, 2024.
- [5] Z. W. Blasingame and C. Liu. Leveraging diffusion for strong and high quality face morphing attacks. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 6(1):118–131, 2024.
- [6] G. J. Burton and I. R. Moorhead. Color and spatial structure in natural scenes. *Applied optics*, 26(1):157–170, 1987.

- [7] 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)*, pages 955–959. IEEE, 2023.
- [8] 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*, pages 1650–1660, 2025.
- [9] K. Chandrasegaran, N.-T. Tran, and N.-M. Cheung. A closer look at fourier spectrum discrepancies for cnn-generated images detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7200–7209, 2021.
- [10] L. Colbois, H. O. Shahreza, and S. Marcel. Approximating optimal morphing attacks using template inversion. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–9. IEEE, 2023.
- [11] 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*, pages 973–982, 2023.
- [12] 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*, pages 1–6. IEEE, 2023.
- [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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1606–1617, 2022.
- [14] N. Damer, A. M. Saladie, A. Braun, and A. Kuijper. Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network. In *IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–10. IEEE, 2018.
- [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 *International Symposium on Visual Computing*, pages 291–304. Springer, 2021.
- [16] 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*, pages 683–692, 2023.
- [17] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5203–5212, 2020.
- [18] T. Dzanic, K. Shah, and F. Witherden. Fourier spectrum discrepancies in deep network generated images. *Advances in Neural Information Processing Systems*, 33:3022–3032, 2020.
- [19] 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)*, pages 1–11. IEEE, 2022.
- [20] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. In *IEEE International Joint Conference on Biometrics*, pages 1–7, 2014.
- [21] M. Ferrara, A. Franco, and D. Maltoni. Decoupling texture blending and shape warping in face morphing. In *International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5. IEEE, 2019.
- [22] D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4(12):2379–2394, 1987.
- [23] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- [24] 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. Synmad 2022: Competition on face morphing attack detection based on privacy-aware synthetic training data. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2022.
- [25] 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*, pages 766–776, 2025.
- [26] M. Ivanovska Preskar, L. Todorov, P. Peer, and V. Struc. Selfmad++: self-supervised foundation model with local feature enhancement for generalized morphing attack detection. *Information fusion*, 127(Part C, 103921), 2026.
- [27] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33:12104–12114, 2020.
- [28] 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)*, pages 1–5. IEEE, 2022.
- [29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019.
- [30] L. Qin, F. Peng, S. Venkatesh, R. Ramachandra, M. Long, and C. Busch. Low visual distortion and robust morphing attacks based on partial face image manipulation. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(1):72–88, 2020.
- [31] H. Rachalwar, M. Fang, N. Damer, and A. Das. Depth-guided robust face morphing attack detection. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–9. IEEE, 2023.
- [32] 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*, pages 255–265. Springer, 2019.
- [33] 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.
- [34] 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*, pages 6–12, 2018.
- [35] C. Tan, Y. Zhao, S. Wei, G. Gu, P. Liu, and Y. Wei. Frequency-aware deepfake detection: Improving generalizability through frequency space domain learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 5052–5060, 2024.
- [36] J. E. Tapia and C. Busch. Face feature visualisation of single morphing attack detection. *11th International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2023.
- [37] v. A. Van der Schaaf and J. v. van Hateren. Modelling the power spectra of natural images: statistics and information. *Vision research*, 36(17):2759–2770, 1996.
- [38] 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. *8th International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2020.
- [39] 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*, pages 1510–1518, 2024.
- [40] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer, and C. Busch. Mipgan—generating strong and high quality morphing attacks using identity prior driven gan. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(3):365–383, 2021.
- [41] 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*, pages 7912–7921, 2024.