Face Re-morphing: Differential Morphing Attack Detection via Feature-Space Similarity Changes

Supplementary Materials

Jie Jin[1] Masakatsu Nishigaki[1] Tetsushi Ohki[1,2]
[1]Graduate School of Science and Technology, Shizuoka University [2]RIKEN AIP
3-5-1 Johoku, Chuo-ku, Hamamatsu, Shizuoka 432-8011
1-4-1 Nihonbashi, Chuo-ku, Tokyo 103-0027, Japan
jie.jin@sec.inf.shizuoka.ac.jp, {nisigaki,ohki}@inf.shizuoka.ac.jp

A. Additional Score Distributions and Qualitative Examples

This supplementary material complements the analysis in Section 5.2 of the main paper by providing additional Δs distributions for all feature extractors. Here, Δs denotes the change in face similarity induced by re-morphing, computed as the similarity between the live image and the re-morphed image minus the original document-live similarity. While the main paper discusses IR101 as a representative example, Figures 1 and 2 show the corresponding distributions for IR101, IR50, ViT, and ViT-k. These results examine whether the feature-space similarity-change behavior discussed in the main paper is specific to a particular face recognition backbone or commonly observed across different feature extractors.

Overall, the distributions are consistent with the observations in the main paper. For MorDIFF (MD) and OpenCV (OCV), bona-fide samples are concentrated around $\Delta s = 0$, whereas morphed samples are shifted to the positive side. In particular, under the OCV setting on FRLL-Morphs, morphed samples are mainly distributed in a clearly positive range, roughly around $\Delta s = 0.25$ to 0.50, while bona-fide samples remain close to zero. This clear separation is consistent with the strong quantitative performance of OCV on FRLL-Morphs.

For MD, a similar tendency is observed, although the morphed distributions are broader than those of OCV. Morphed samples generally move to positive Δs values, while bona-fide samples remain near zero. On FEI Morph Version 2, the morphed distributions spread more widely and overlap more with the bona-fide distributions than on FRLL-Morphs. This tendency is consistent with the higher error rates observed on FEI Morph Version 2 in the main experiments.

In contrast, StyleGAN2-ADA (SG2-A) shows a qualitatively different behavior. For SG2-A, both bona-fide and morphed samples are located mainly on the negative side of Δs. For example, bona-fide samples tend to concentrate around approximately $\Delta s = -0.5$, while morphed samples are shifted closer to zero but still remain largely negative. Thus, SG2-A does not show the same positive-shift tendency for morphed samples as MD and OCV.

Figure 3 shows qualitative examples of Re-morphed images generated by OCV, MD, and SG2-A for bona-fide and morphing cases. OCV preserves the image-level structure of the input images through landmark-based warping and blending, while MD generates smoother Re-morphed images through representation-space interpolation. In contrast, the SG2-A examples show larger changes in facial appearance relative to the Document image and the Live image. This suggests that GAN inversion and synthesis may introduce identity drift during re-morphing, causing the Re-morphed feature to move less consistently toward the Live image feature. Such behavior can weaken the expected positive shift of Δs and increase the overlap between bona-fide and morphed distributions.

The qualitative examples also indicate that perceptual image quality alone does not fully explain the behavior of the proposed score. Although SG2-A can synthesize visually plausible faces, the generated Re-morphed images may deviate from the identity relationship implied by the Document image and the Live image. This observation is consistent with the discussion in the main paper that the behavior of Face Re-morphing is influenced not only by visual quality, but also by identity consistency in the face feature space.

Overall, these additional distributions and qualitative examples indicate that the main observations in the paper are not specific to IR101. Across multiple face recognition feature spaces, MD and OCV tend to produce the expected positive shift for morphed samples, whereas SG2-A shows negative-shifted and more overlapping distributions. These results support that the proposed score is generally effective across different feature extractors, while its behavior also depends on how the re-morphing method affects the identity relationship among the Document image, the Live image, and the Re-morphed image.

B. Choice of the Interpolation Coefficient

We fixed the interpolation coefficient to $\alpha = 0.5$ so that the Document and Live images contribute equally to the re-morphing process. This symmetric setting does not favor either input identity and provides a neutral basis for evaluating the proposed framework. Although other values of α may produce different similarity changes, this study focuses on the commonly used balanced setting of $\alpha = 0.5$. The same coefficient was therefore applied to all re-morphing methods and datasets without condition-specific tuning.

C. Identity Preservation across Re-morphing Methods

To further discuss the performance differences across re-morphing methods, we examine the cosine similarities between the Re-morphed feature and the two input features, namely $\cos(\mathcal{F}_{Re}, \mathcal{F}_D)$ and $\cos(\mathcal{F}_{Re}, \mathcal{F}_L)$. These similarities indicate how closely the Re-morphed image remains related to the Document and Live identities in the face feature space.

MD and OCV generally maintain high similarities between F_{Re} and both input features. On FEI Morph Version 2, MD showed mean similarities of 0.825 and 0.829 under the Accomplice condition, and 0.707 and 0.690 under the Criminal condition, for $\cos(\mathcal{F}_{Re}, \mathcal{F}_D)$ and $\cos(\mathcal{F}_{Re}, \mathcal{F}_L)$, respectively. OCV showed even higher values of 0.871 and 0.886 under the Accomplice condition, and 0.757 and 0.765 under the Criminal condition. In these morphing cases, the similarities involving F_{Re} were higher than the original Document–Live similarity $\cos(\mathcal{F}_D, \mathcal{F}_L)$, which was 0.711 for Accomplice and 0.421 for Criminal. This tendency suggests that MD and OCV preserve the relationships with both input identities more effectively than SG2-A, which is consistent with their positive Δs distributions and relatively low error rates.

In contrast, SG2-A yields substantially lower similarities between F_{Re} and both input features. On FEI Morph Version 2, the mean similarities were 0.419 and 0.435 under the Accomplice condition, and 0.341 and 0.361 under the Criminal condition. A related tendency was also observed for the StyleGAN subset of FRLL-Morphs, which was generated using a StyleGAN-based generation process. In this subset, the mean Document–Live similarity $\cos(\mathcal{F}_D, \mathcal{F}_L)$ was 0.247 under the Accomplice condition and 0.233 under the Criminal condition. These values were markedly lower than those for the FaceMorpher, OpenCV, and WebMorph subsets, whose mean similarities were approximately 0.54–0.56. This suggests that low identity similarity is a recurring tendency of the StyleGAN-based generation process rather than a phenomenon specific to FEI Morph Version 2.

Together, these observations suggest that GAN inversion and synthesis may introduce identity drift during both original morph creation and Re-morphing. Consequently, $\mathcal{F}_{Re}$ may not remain close to both input features, resulting in negative or weakened Δs and poorer SG2-A detection performance.

Table 1. Mean cosine similarities between the Re-morphed feature and the two input features on FEI Morph Version 2. Higher values indicate that the Re-morphed feature remains closer to the input identities.

Method	Condition	$\cos(\mathcal{F}_{Re}, \mathcal{F}_D)$	$\cos(\mathcal{F}_{Re}, \mathcal{F}_L)$	$\cos(\mathcal{F}_D, \mathcal{F}_L)$
MD	Accomplice	0.825	0.829	0.711
	Criminal	0.707	0.690	0.421
OCV	Accomplice	0.871	0.886	0.711
	Criminal	0.757	0.765	0.421
SG2-A	Accomplice	0.419	0.435	0.711
	Criminal	0.341	0.361	0.421

Table 2. Mean Document–Live cosine similarities $\cos(\mathcal{F}_D, \mathcal{F}_L)$ for different morphing subsets of FRLL-Morphs. Bold values indicate the notably low similarities observed for the StyleGAN subset.

Morphing subset	Condition	$\cos(\mathcal{F}_D, \mathcal{F}_L)$
AMSL	Accomplice	0.635
	Criminal	0.445
FaceMorpher	Accomplice	0.545
	Criminal	0.540
OpenCV	Accomplice	0.555
	Criminal	0.551
StyleGAN	Accomplice	0.247
	Criminal	0.233
WebMorph	Accomplice	0.547
	Criminal	0.547

Table 3. Mean runtime per image pair in seconds.

Method	Preprocessing	Re-morphing	Detection	Total
OCV	1.38	0.010	0.058	1.45
MD	1.49	11.63	0.058	13.18
SG2-A	1.61	65.34	0.058	67.00

D. Runtime Analysis

Runtime was measured using one CPU thread for OCV and one NVIDIA RTX 6000 Ada GPU for MorDIFF, SG2-A, and CVLFace. Each method was evaluated using 10 image pairs after one warm-up pair. Preprocessing covered face detection and alignment, while detection covered image loading, alignment, feature extraction, and cosine-similarity computation. Table 3 summarizes the mean runtime per pair.

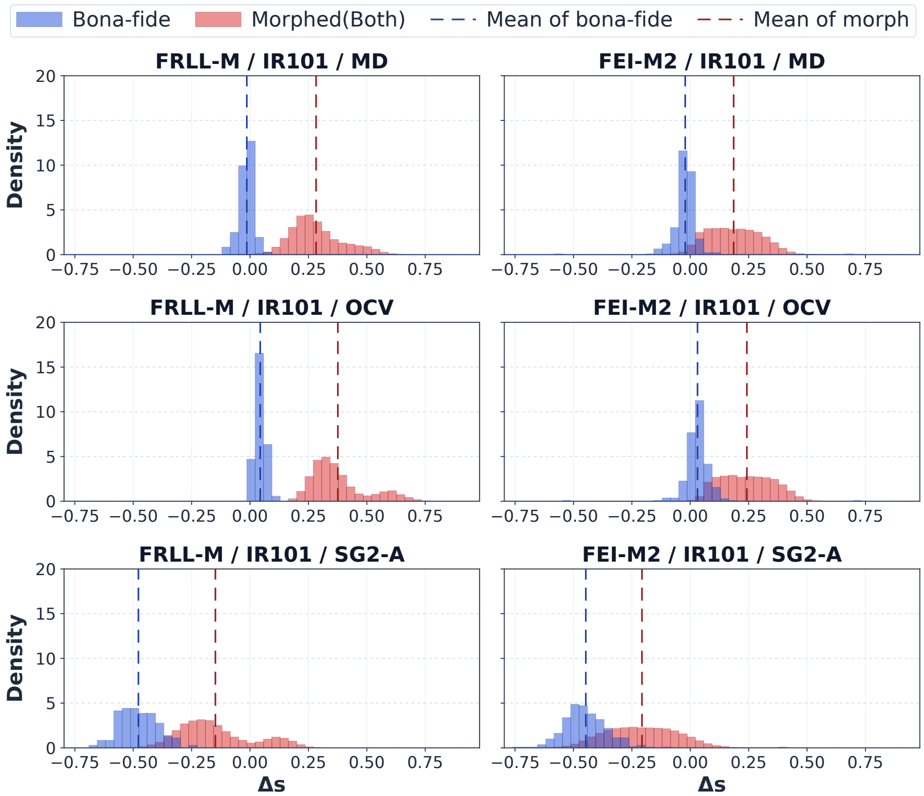

IR101

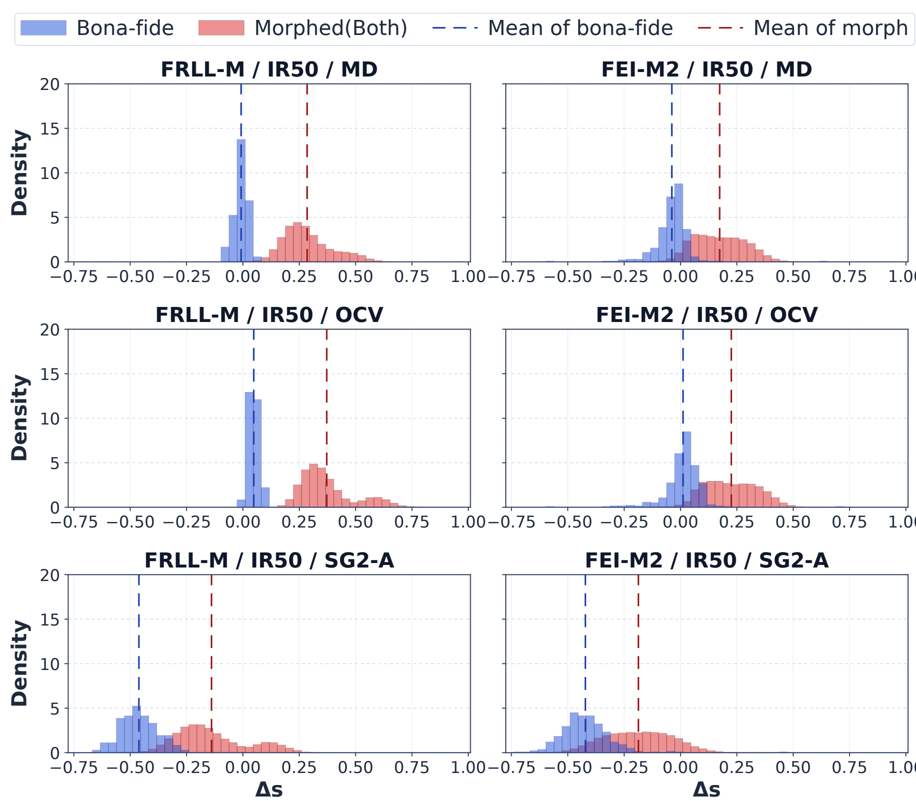

IR50

Figure 1. Distributions of Δs for bona-fide and morphed samples using IR101 and IR50 on FRLL-Morphs and FEI Morph Version 2.

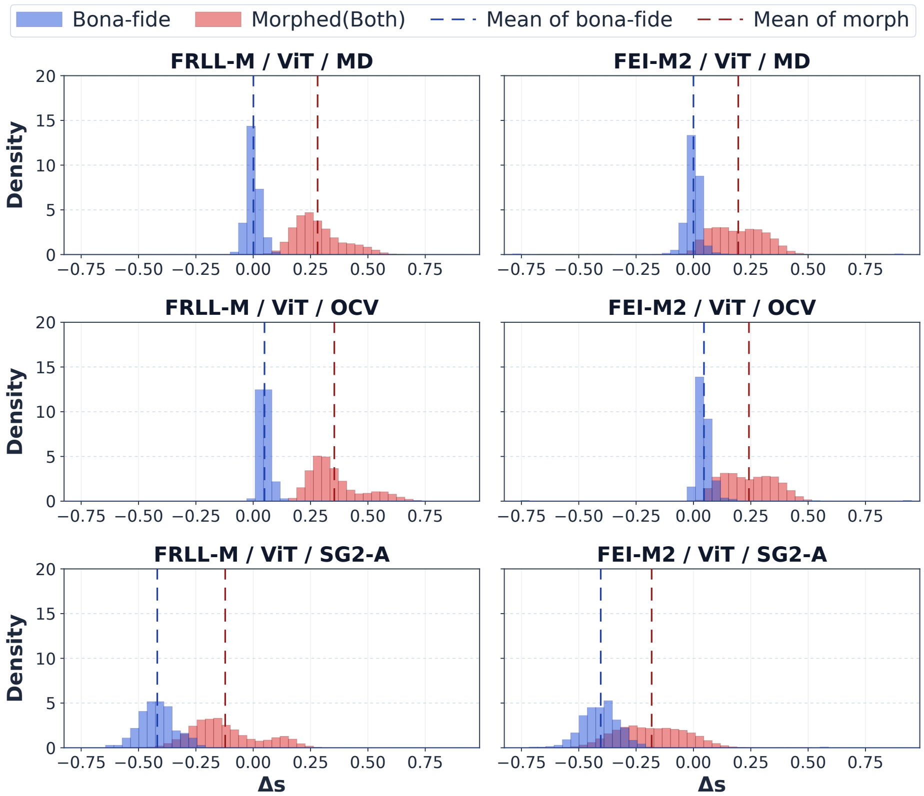

ViT

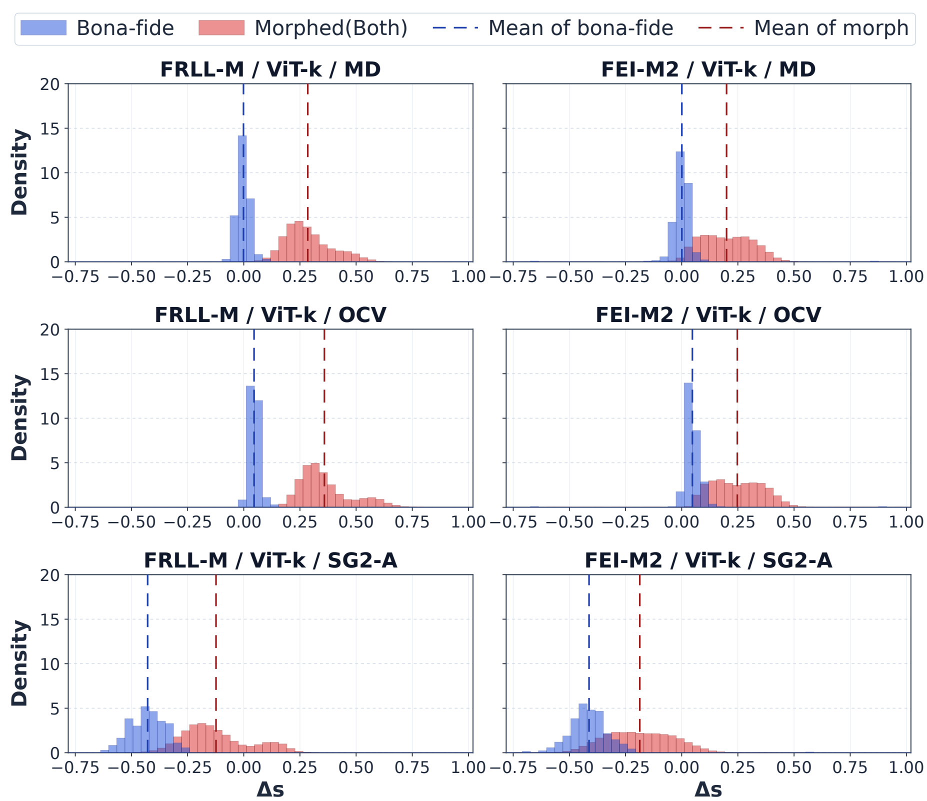

ViT-k

Figure 2. Distributions of Δs for bona-fide and morphed samples using ViT and ViT-k on FRLL-Morphs and FEI Morph Version 2.

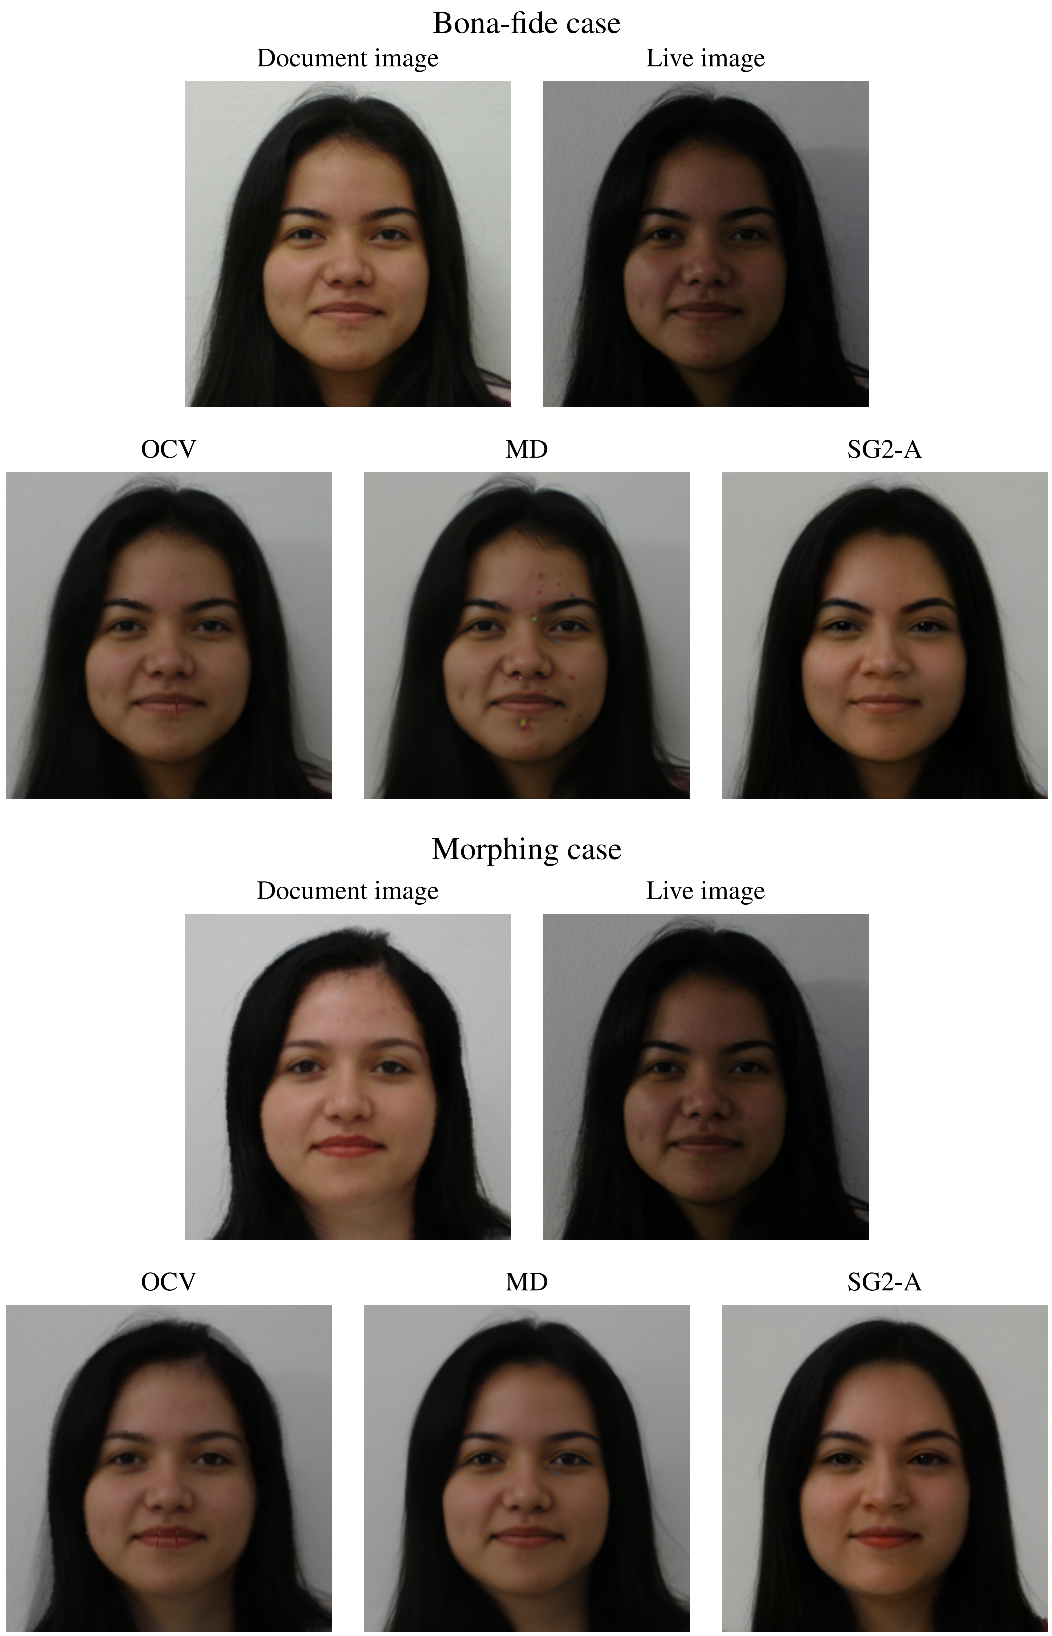

Figure 3. Qualitative examples from the FEI Morph dataset showing Re-morphed images for bona-fide and morphing cases. For each case, the Re-morphed images are generated from the Document image and the Live image using OCV, MD, and SG2-A.