

# Frequency-domain Manipulation for Face Obfuscation

Jintae Kim  
Korea University  
jtkim@mcl.korea.ac.kr

Keunsoo Ko  
The Catholic University of Korea  
ksko@catholic.ac.kr

Chang-Su Kim\*  
Korea University  
changsupkim@korea.ac.kr

Table S-1. Runtime comparison across two GPU configurations.

	RTX 2080 Ti	RTX 3090
IdentityHider [51]	86.77	68.14
Forbes [20]	1086.83	739.57
FreM	88.82	67.19

## A. Analysis

To evaluate the computational efficiency of training-based (IdentityHider [51]) and training-free (Forbes [20] and FreM) methods, we measure their runtimes on two GPUs, as shown in Table S-1. All methods are tested on the LFW dataset [16] using  $112 \times 112$  images.

Forbes requires optimizing multiple spatial-domain filters via a repetitive parameter refinement process, resulting in substantial computational overhead and the slowest runtime. In contrast, IdentityHider, which relies on a pre-trained network, performs obfuscation through a single feed-forward pass and therefore runs fast. Notably, although FreM performs iterative parameter refinement, its runtime remains close to that of IdentityHider, indicating that a few refinement steps are required in practice.

## B. Dataset

**Face recognition:** LFW [16] dataset contains 13,233 facial images with 5,749 different identities. We use 6,000 pairs of LFW datasets with half matched and half unmatched, which is a benchmark for face verification [10, 21]. CALFW [59] and CPLFW [58] datasets are based on the same identities as LFW, each consisting of 6,000 pairs of images specifically selected to have significant differences in age and pose, respectively. AgeDB [39] is another verification dataset consisting of 6,000 pairs, specifically constructed to evaluate recognition performance under large age gaps (approximately 30 years). The CFP-FP [46] dataset consists of 7,000 pairs that specifically test the robustness against significant cross-pose differences, where each pair consists of two non-centric profile faces captured from different angles.

**Age estimation:** MORPH II [44] dataset consists of 55,134 facial images with 13,617 identities spanning an age range of [16, 77]. Following MWR [47], we use a subset of 21K images, maintaining a 1:1 ratio of Caucasians to Africans and a 1:3 ratio of females to males. UTKFace [57] provides about 20K facial images spanning the age range [0, 116], of which we use 3,287 for test, as in [43, 47].

**Expression recognition:** RAF-DB [9] dataset contains 15,339 facial images with six emotion classes, whereas FERPlus [4] dataset contains 35,887 facial images with seven emotion classes. Following the conventional method (Faceptor [43]), we use 3,068 facial images from the RAF-DB dataset and 3,589 images from the FERPlus dataset.

**Binary attribute classification:** We use the CelebA [30] test dataset, which contains about 20K celebrity face images. Each image is annotated with 40 facial attributes such as smiling, wearing glasses, and mustache.

## C. Reconstruction attacks

For the reconstruction attack, we employ a U-Net encoder-decoder as the reconstruction network. The network takes an obfuscated image  $I_{\text{out}}$  as input and produces a reconstructed image  $I_{\text{rec}}$ . We optimize the network by minimizing the  $\ell_1$  distance between  $I_{\text{rec}}$  and the original image  $I_{\text{in}}$  using the AdamW optimizer (learning rate  $10^{-4}$ , weight decay  $10^{-4}$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ). We train the reconstruction network on the MS1MV2 [10] dataset. The model is trained for 100K iterations with a batch size of 64.

## D. Comparison Results

Note that Tables 1 and 2 in the main paper demonstrate that FreM outperforms conventional obfuscation methods for face recognition. To further assess its generalization, we conduct experiments using Faceptor [43]. As shown in Table S-2, FreM again achieves superior performance.

Figure S-1, S-2, S-3, S-4, and S-5 compare the qualitative results and their corresponding reconstructed results on the LFW [16], AgeDB [39], CALFW [59], CPLFW [58], and CFP-FP [46] datasets, respectively.

\*Corresponding author.

Table S-2. Quantitative comparison (XDR/ODR) of face obfuscation results on the LFW [16], AgeDB [39], CALFW [59], CPLFW [58], and CFP-FP [46] datasets. In each test, the best result is **boldfaced**, while the second best is underlined. The recognition accuracies are measured using Faceptor [42].

	Dataset				
	LFW	AgeDB	CALFW	CPLFW	CFP-FP
Original	99.45 / -	93.10 / -	94.55 / -	92.03 / -	95.86 / -
PRO-Face (Blur) [53]	97.13 / 95.15	84.57 / 74.22	87.98 / 81.88	87.95 / 80.92	88.96 / 83.86
PRO-Face (Pixelate) [53]	96.87 / 94.17	81.32 / 70.53	85.95 / 79.27	86.95 / 79.27	89.84 / 82.44
PRO-Face (FaceShifter) [53]	97.17 / 95.88	81.23 / 74.68	87.20 / 84.30	85.13 / 76.83	89.41 / 80.40
PRO-Face (SimSwap) [53]	96.28 / 95.87	82.02 / 77.12	88.55 / 85.98	84.93 / 77.28	<u>90.04</u> / <u>84.27</u>
Forbes [19]	<u>98.30</u> / <u>96.96</u>	88.27 / 82.15	88.85 / 84.26	<u>88.89</u> / <u>83.82</u>	89.63 / 83.80
IdentityHider [50]	97.00 / 96.13	<u>89.98</u> / <u>88.40</u>	<u>91.62</u> / <u>89.02</u>	85.05 / 77.80	87.54 / 81.74
Proposed Algorithm	<b>98.96</b> / <b>97.97</b>	<b>92.03</b> / <b>88.78</b>	<b>93.73</b> / <b>91.42</b>	<b>90.33</b> / <b>87.55</b>	<b>93.30</b> / <b>89.99</b>




Figure S-1. Qualitative comparison of obfuscated images and their corresponding reconstructed results. The images are selected from LFW [16] dataset.






Figure S-2. Qualitative comparison of obfuscated images and their corresponding reconstructed results. The images are selected from AgeDB [39] dataset.

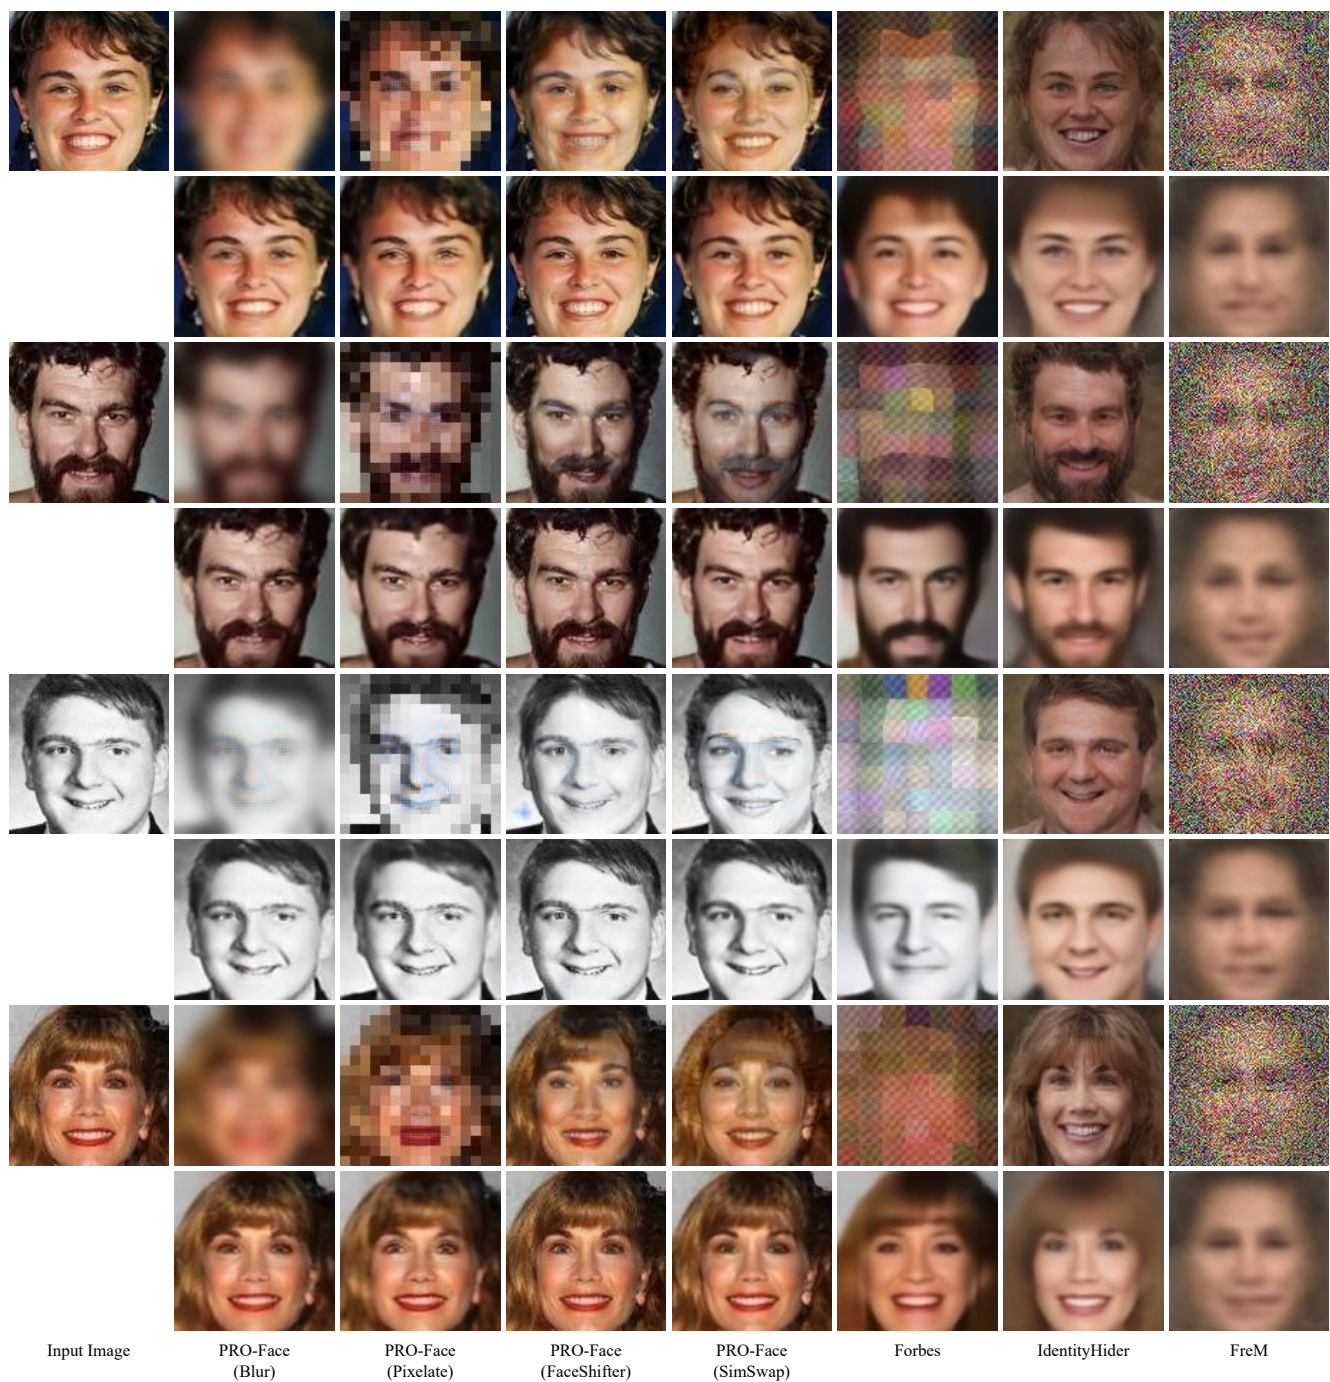


Figure S-3. Qualitative comparison of obfuscated images and their corresponding reconstructed results. The images are selected from CALFW [59] dataset.






Figure S-4. Qualitative comparison of obfuscated images and their corresponding reconstructed results. The images are selected from CPLFW [58] dataset.

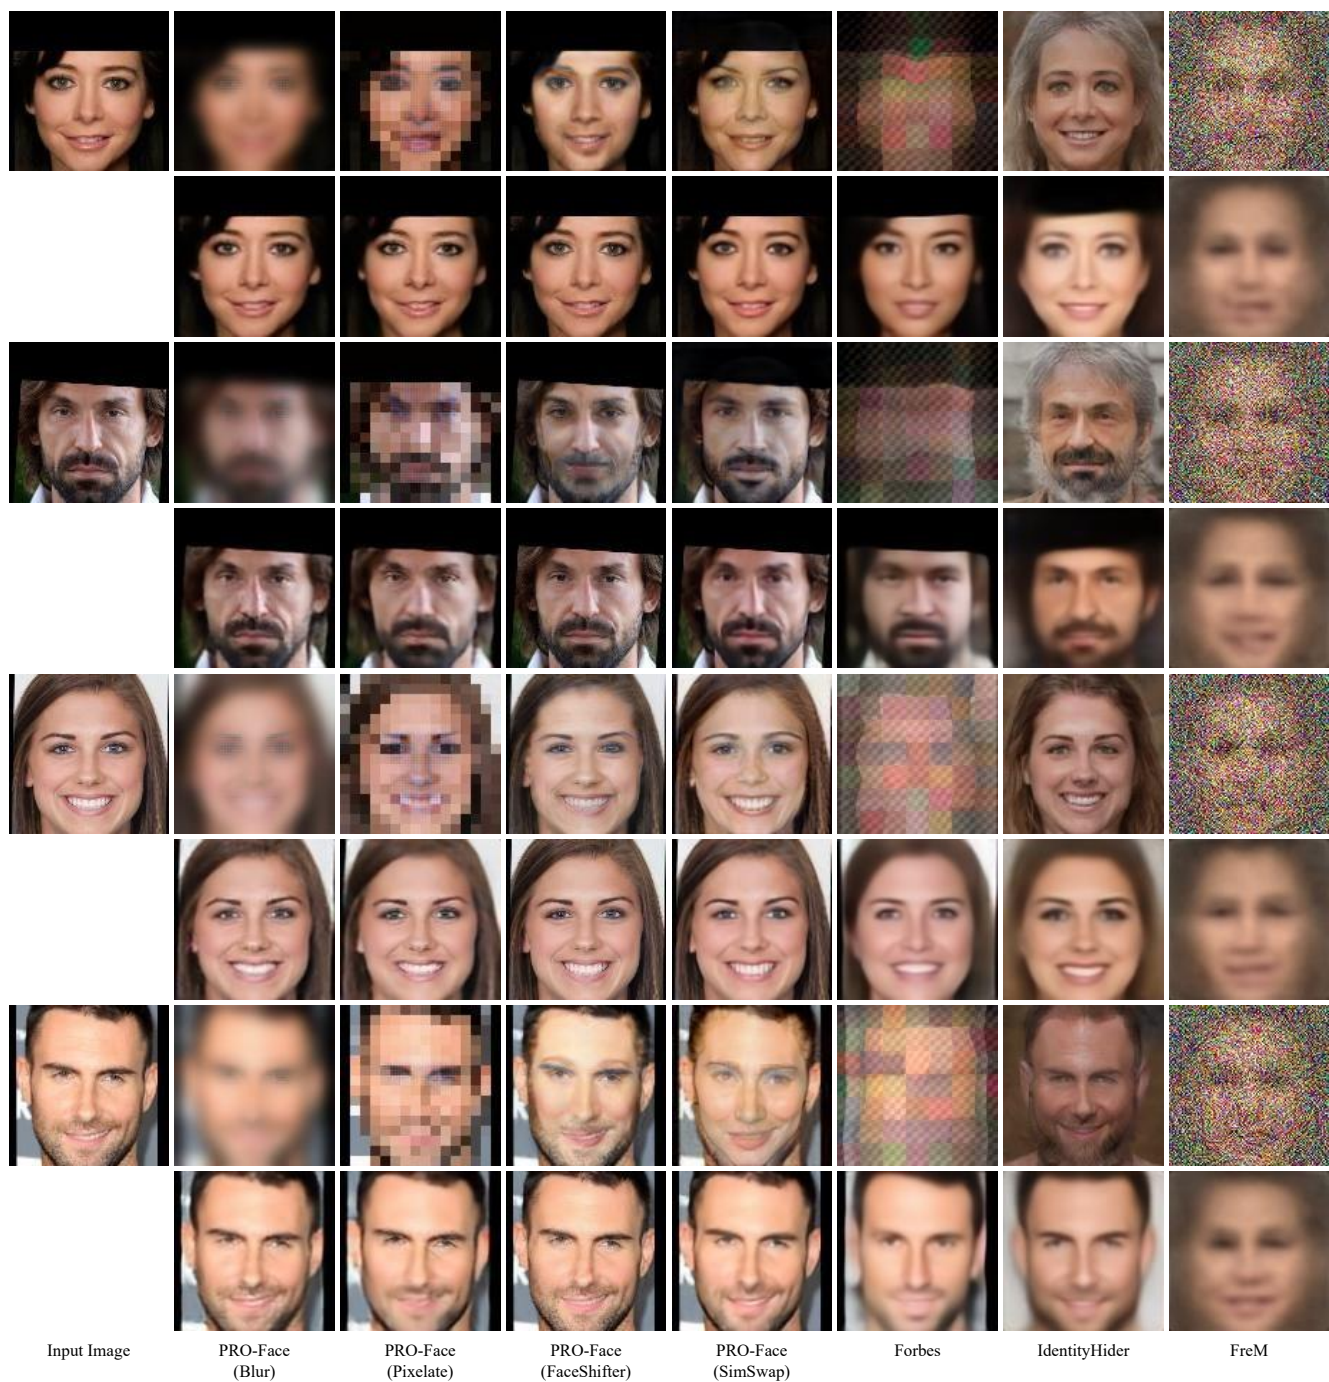


Figure S-5. Qualitative comparison of obfuscated images and their corresponding reconstructed results. The images are selected from CFP-FP [46] dataset.