

FaceMe: Robust Blind Face Restoration with Personal Identification

Siyu Liu^{1*}, Zheng-Peng Duan^{1*}, Jia OuYang², Jiayi Fu¹,
Hyunhee Park³, Zikun Liu², Chun-Le Guo^{1, 4†}, Chongyi Li^{1, 4}

¹VCIP, CS, Nankai University

²Samsung Research, China, Beijing (SRC-B)

³The Department of Camera Innovation Group, Samsung Electronics

⁴NKIARI, Shenzhen Futian

liusiyu29@mail.nankai.edu.cn, adamduan0211@gmail.com, jia.ouyang@samsung.com, fujiayi@mail.nankai.edu.cn,
{inextg.park, zikun.liu}@samsung.com, {guochunle, lichongyi}@nankai.edu.cn

Abstract

Blind face restoration is a highly ill-posed problem due to the lack of necessary context. Although existing methods produce high-quality outputs, they often fail to faithfully preserve the individual’s identity. In this paper, we propose a personalized face restoration method, FaceMe, based on a diffusion model. Given a single or a few reference images, we use an identity encoder to extract identity-related features, which serve as prompts to guide the diffusion model in restoring high-quality and identity-consistent facial images. By simply combining identity-related features, we effectively minimize the impact of identity-irrelevant features during training and support any number of reference image inputs during inference. Additionally, thanks to the robustness of the identity encoder, synthesized images can be used as reference images during training, and identity changing during inference does not require fine-tuning the model. We also propose a pipeline for constructing a reference image training pool that simulates the poses and expressions that may appear in real-world scenarios. Experimental results demonstrate that our FaceMe can restore high-quality facial images while maintaining identity consistency, achieving excellent performance and robustness.

Code — https://modyu-liu.github.io/FaceMe_Homepage

Introduction

Face restoration focuses on improving the quality of facial images by removing complex degradation and enhancing details. In the real world, facial images frequently suffer from complex degradation, *e.g.*, blurring, noise, and compression artifacts, which can significantly impact downstream tasks like face recognition and detection.

Face restoration is inherently a highly ill-posed task because a single low-quality input can correspond to many potential high-quality counterparts scattered throughout the high-quality image space. Most advanced face restoration methods cannot guarantee identity consistency, *i.e.*, the identity of the restored portrait may significantly deviate from

the actual identity. Such deviations in identity are unacceptable in practical applications, *e.g.*, restoring facial images in personal photo albums.

Recently, some studies have used high-quality reference images of the same identity to enhance identity consistency in restored facial images. These approaches improve results through feature alignment (Li et al. 2018, 2020, 2022) or by fine-tuning the diffusion model to constrain the generation prior (Varanka et al. 2024; Ding et al. 2024; Chari et al. 2023). Although these methods have achieved significant results, they have obvious limitations. The feature alignment-based methods are influenced by the degradation of the input image and the pose of the reference image. If the features are not well aligned, the quality of the restoration can significantly deteriorate. The fine-tuning-based methods typically require 5~20 reference images and need to be fine-tuned again whenever the identity changes, making it a time-consuming process. Additionally, the quality of the reference images may significantly impact image restoration quality. We illustrate the differences between our proposed personalized method and the fine-tuning-based personalized method in Fig. 1.

In this paper, we proposed FaceMe, a fine-tuning-free personalized blind face restoration method based on the diffusion model. Given a low-quality input and either a single or a few high-quality reference images of the same identity, FaceMe restores high-quality facial images and maintains identity consistency within seconds. Remarkably, changing identities does not require fine-tuning, and the reference images can have any posture, expression, or illumination. Furthermore, the quality of the reference image does not significantly impact the quality of the restored image. To our knowledge, this is the first approach that leverages diffusion prior for personalized face restoration tasks, which does not require fine-tuning when changing identity.

The fine-tuning-free personalized blind face restoration currently faces three significant challenges: 1) *Influence of identity-irrelevant features from reference images*; 2) *Balancing the dependency between low-quality input and reference images*; and 3) *Insufficient datasets for personalized blind face restoration*.

To address the **first challenge**, we propose to use an

*These authors contributed equally.

†Corresponding author.

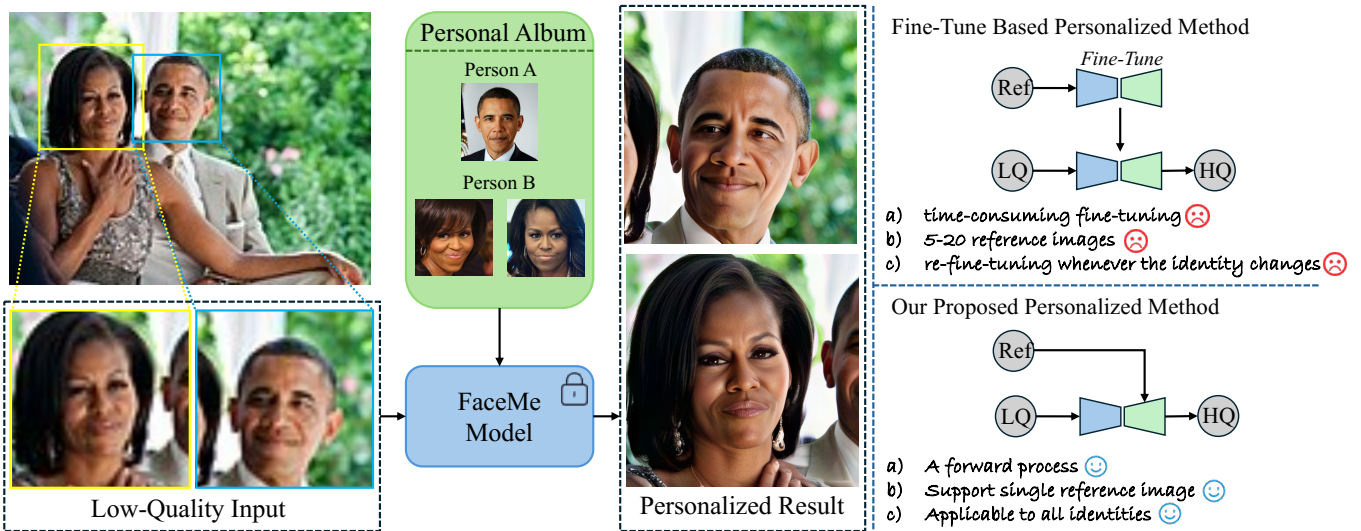


Figure 1: Make the people in the photo look like you and those you are familiar with. Using a single or a few reference images, we can restore realistic images without any fine-tuning for identity. **Zoom in for best view.**

identity encoder to extract identity-related features and a multi-reference image training approach to minimize the impact of identity-irrelevant features. In addition, we found that the **second challenge** of balancing dependencies between low-quality input and reference images arises from simultaneous personalization and restoration training. To address it, we propose a two-stage training approach that effectively balances these dependencies. In the training stage I, we focus on training the model’s personalization ability, *i.e.*, training the proposed identity encoder. In the training stage II, we fix the identity encoder and train a ControlNet, which corresponds to the image restoration capability of the model. Considering the **third challenge** of insufficient training data, we propose to use synthetic datasets as reference images. Specifically, to better align reference images with real-world application scenarios, we create a pose-reference pool categorized by poses and expressions. During the synthesis process, we sample several images from this pose-reference pool to serve as pose references and select one image from a publicly available facial image dataset to serve as identity control. Using a personalized generative model, we synthesize reference images that match the poses of the given pose references while maintaining the identity of the provided identity control image, thereby producing a set of images with the same identity but diverse poses.

Related Work

Blind Face Restoration

Recently proposed blind face restoration methods rely on generating priors to achieve high-quality image restoration. Methods based on GAN priors, such as GPEN (Yang et al. 2021) and GFP-GAN (Wang et al. 2021), enhance the quality of restored images by embedding rich generative priors from pre-trained face GANs into their network structures. Inspired by VQGAN’s (Esser, Rombach, and Om-

mer 2021) success in image generation, methods such as VQFR (Gu et al. 2022), CodeFormer (Zhou et al. 2022), and RestoreFormer (Wang et al. 2022) focus on learning high-quality codebooks and achieving superior image restoration by matching high-quality vectors. Recently, diffusion models have been applied to face restoration, with methods like DR2 (Wang et al. 2023) and DiffFace (Yue and Loy 2024) using these models for robust degradation removal. Although these blind face restoration methods achieve impressive results, restored identities often deviate from the real ones due to the lack of identity information.

Reference-based Face Restoration

To enhance the quality of restored images and ensure identity similarity, several methods leverage reference images of the same identity for face restoration. GFRNet (Li et al. 2018), ASFFNet (Li et al. 2020), and DMDNet (Li et al. 2022) require alignment modules to align the features of low-quality images and reference images. However, when low-quality images’ degradation is severe or the quality of the reference images is suboptimal, the restoration quality is compromised. Recently, some methods (Chari et al. 2023; Varanka et al. 2024; Ding et al. 2024) achieve personalized face restoration based on diffusion models. Their core idea is to utilize 5 ~ 20 reference images of the same identity to constrain the generation priors of the diffusion model through fine-tuning. Although these methods are effective in maintaining identity, they are generally time-consuming and require re-fine-tuning whenever the identity is changed.

Personalized Generation

Recently, some personalized face generation models have been proposed. Due to the constraints of the dataset, IP-Adapter (Ye et al. 2023), InstantID (Wang et al. 2024b), and Arc2Face (Paraperas Papantoniou et al. 2024) are trained

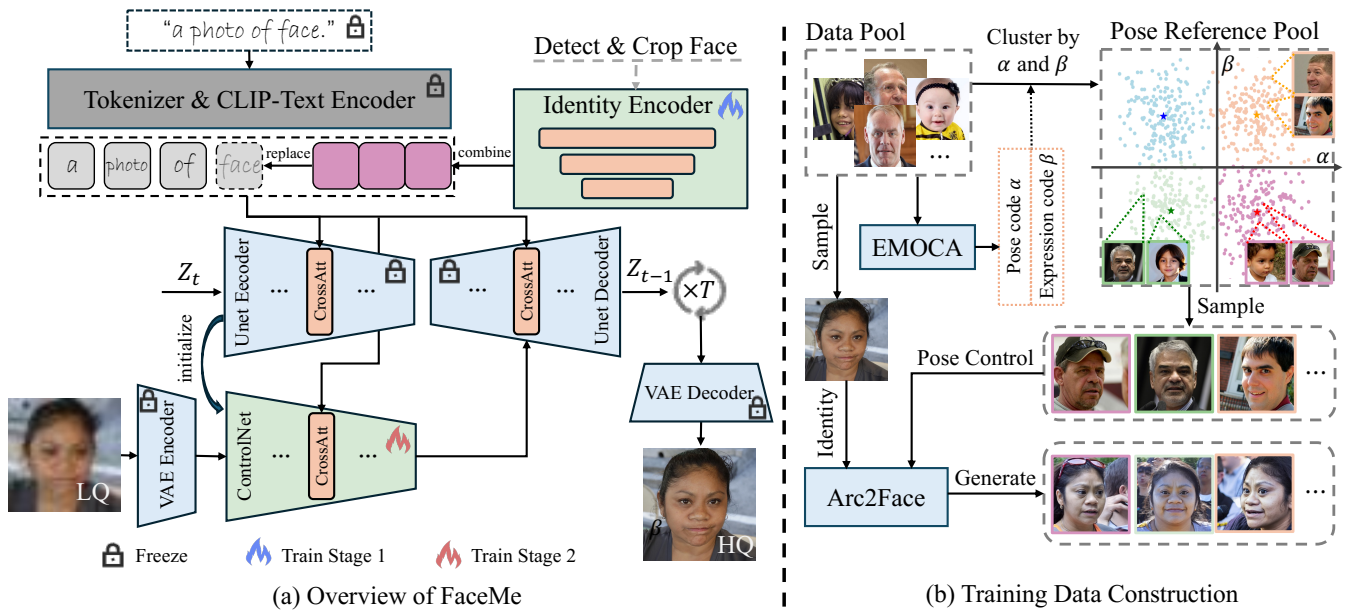


Figure 2: **Overview of proposed FaceMe (left) and training data construction pipeline (right).** For the proposed FaceMe, identity-related features from the reference image are extracted by the identity encoder, by simply combining to support multi-reference image inputs. We first use a fixed text, *i.e.*, *a photo of face*. and then apply the combined identity-related features to replace the *face* embedding. The updated embeddings are sent to the cross-attention layer of the diffusion model to guide personalized face image restoration.

in a self-supervised manner, *i.e.*, using the same facial image as reference and ground truth. Although this training strategy can generate realistic images that match the identity of a given reference facial image, it also introduces many identity-irrelevant features, such as posture and expressions. Photomaker (Li et al. 2024) collects a large-scale set of images of the same identity and supports the input of any number of reference images. The essence of personalized generation lies in blending reference images with text descriptions to create character images that both align with the identity of the reference images and adhere to the text descriptions. Personalized image restoration, on the other hand, seeks to balance the reliance on both reference images and low-quality input, aiming to restore high-fidelity facial images that maintain consistency with the identity of the reference image. Given the similarities in the core objectives of these tasks, personalized generation methods can be effectively applied to personalized restoration.

Methodology

Overview

An overview of the proposed FaceMe is shown in Fig. 2. Given a few reference images of the same individual as the target image, the primary objective of this work is to leverage prior knowledge from a diffusion model to restore realistic facial images from degraded ones, while ensuring the individual’s identity in the restored images. By utilizing the identity encoder to extract identity-related features, we can obtain a unified identity representation from a given reference image without needing to fine-tune the model for each

new identity. By simply combining identity-related features, our method enables an arbitrary number of reference images (single or a few). Furthermore, to effectively train our model, we design a dataset construction pipeline to remedy the lack of personalized facial image restoration datasets.

Problem Definition

Let X , Y , \mathcal{D} , and \mathcal{G} denotes the degraded facial image, the corresponding high-quality facial image, the degradation function, and the generation function, respectively. The objective of facial image restoration is to generate $\hat{Y} = \mathcal{G}(X)$, while ensuring the following two constraints are satisfied:

$$\text{Consistency} : \mathcal{D}(\hat{Y}) \equiv X, \text{Realness} : \hat{Y} \sim q(Y), \quad (1)$$

where $q(Y)$ denotes the distribution of high-quality facial images.

Let Ref denote the reference images of the same identity as Y . Based on reference images, personalized face restoration can be formulated as $\hat{Y} = \mathcal{G}(X|Ref)$. In addition to meeting *Consistency* and *Realness*, it also needs to meet the following constraint:

$$\text{Identity Consistency} : \hat{Y} \sim ID(Y), \quad (2)$$

where $ID(Y)$ denotes the distribution of high-quality facial images of the same identity as Y .

If we consider only *Realness* and *Identity Consistency*, it becomes a personalized face generation task. From the perspective of problem definition, we can decompose the problem of personalized face restoration into a combination of

personalized face generation and face restoration. Next, we detail how we achieve personalized face restoration based on personalized face generation.

Proposed Personal Face Restoration

Identity encoder Following recent works (Wu et al. 2024b; Guo et al. 2024), we combine the CLIP (Radford et al. 2021) image encoder ε and the facial recognition module ψ , *i.e.*, ArcFace (Deng et al. 2019) network, to extract identity features from facial images. Due to the different feature dimensions extracted by ε and ψ , we employ MLPs to map the features extracted by ψ to align with those extracted by ε . The final embeddings, which correspond to the identity information of the given reference images, are obtained by merging these two sets of features using MLPs.

Specifically, let $\{R_i\}_{i=1}^N$ denote the given N reference facial images. After face detection and cropping, r_i is obtained from R_i . We use ε and ψ to extract features from r_i , denoted as $f_i \in \mathbb{R}^d$ and $g_i \in \mathbb{R}^{512}$, respectively, where d denotes the dimension of cross-attention in the diffusion model. We then employ MLPs to align the dimension of g_i with f_i , resulting in $\hat{g}_i \in \mathbb{R}^d$. Finally, we use MLPs to fuse \hat{g}_i and f_i , resulting in $s_i \in \mathbb{R}^d$.

Combining and replacing We begin by combining the extracted e_i , obtaining:

$$s = \text{Concat}([s_1, \dots, s_N]), s \in \mathbb{R}^{N \times d}. \quad (3)$$

Since our approach does not require semantic guidance, we consistently use a simple prompt: “a photo of face.”, during both the training and inference phases. Let $c_{text} = \{e_1, e_2, e_3, e_4, e_5\}$ represent the embedding of the text obtained through the tokenizer and CLIP text encoder. We replace the embedding e_4 corresponding to the word “face” with s , resulting in $c_{id} = \{e_1, e_2, e_3, s, e_5\}$, which is used as prompt embedding to guide the diffusion model for personalized face image restoration.

Training strategy The model consists of two trainable modules, *i.e.*, ControlNet and ID encoder (Identity Encoder in Fig. 2). In this work, we propose a two-stage training strategy. In training stage I, we simultaneously train ControlNet and ID encoder, but only save ID encoder’s weights. In training stage II, we fix the ID encoder and only train ControlNet. In this training process, we randomly replace identity embedding c_{id} with non-identity embedding c_{text} with a 50% probability. The loss function for the two training stages can be expressed as:

$$\mathcal{L} = \mathbb{E}_{z_t, I_{LQ}, I_{ref}, \epsilon} \|\epsilon - \epsilon_\theta(z_t, I_{LQ}, I_{ref})\|_2, \quad (4)$$

where ϵ is target noise, ϵ_θ is the proposed model, z_t is the latent code at time t , I_{LQ} is the low-quality input, and I_{ref} are the reference images. Next, we will discuss the reasons for using such a training strategy.

Our training strategy is based on the following experimental insights: 1) When training the ID encoder independently, *i.e.*, training a personalized generative model, we encounter unsatisfactory results. The ability to generate personalized faces is significantly compromised due to insufficient training data and the absence of semantics. 2) Simultaneously

training the ControlNet and ID encoder can lead to instability. As training progresses, although the model’s generative ability improves, the control ability of low-quality image diminishes, leading to undesired artifacts in restored images. However, we found that upon completing the simultaneous training, the personalized face restoration model can seamlessly transition into a personalized generative model by simply removing the ControlNet.

The primary goal of training stage I is to train the ID encoder, *i.e.*, to train the model’s **personalized generation** ability. Training with ControlNet, as opposed to training the ID encoder only, can compensate for the limitations of insufficient training data and the absence of semantics. Since ControlNet provides the layout information of the target image, the ID encoder can focus on guiding the model to restore identity-specific detail features.

In training stage II, we focus on balancing the dependency between low-quality input and reference images. When training the ControlNet and ID encoder simultaneously, the continuous updating of the ID encoder’s weights prevents the ControlNet from providing stable low-quality image control. To enhance the ControlNet’s ability, we fix the ID encoder obtained in stage I and then train only the ControlNet, replacing c_{id} with c_{text} with a 50% probability.

Inference strategy Following (Wu et al. 2024a), we embed the low-quality input directly into the initial random Gaussian noise according to the training noise scheduler. Using Classifier-free guidance (CFG) (Ho and Salimans 2022) for personalized guidance, the CFG in the inference stage can be represented as:

$$\begin{aligned} z_{t-1}^{id} &= \phi(z_t, z_{LQ}, c_{id}), z_{t-1} = \phi(z_t, z_{LQ}), \\ \bar{z}_{t-1} &= z_{t-1} + \lambda_{cfg} \times (z_{t-1}^{id} - z_{t-1}), \end{aligned} \quad (5)$$

where $\phi(\cdot)$ denotes the proposed model, λ_{cfg} is a hyperparameter, z_{t-1} is the output of model without identity control, and z_{t-1}^{id} is the output of model with identity control. In addition, to mitigate the possibility of color shift, we apply wavelet-based color correction (Wang et al. 2024a) to the final result.

Training Data Pool Construction

To our knowledge, no publicly available facial dataset can support diffusion models for training with multiple reference images of the same identity. Additionally, following the Photomaker (Li et al. 2024) to collect a large-scale facial dataset of the same identity is both time-consuming and challenging. Thus, in this study, we employ synthetic facial images as reference facial images to construct our training data pool. Unlike previous methods (Li et al. 2018, 2020, 2022) that extract features from reference facial images and align them with degraded ones to enhance image details, our proposed method uses reference facial images as prompt to guide the diffusion model denoising process. Without a substantial shift in identity, it allows for lower-quality reference images without significantly affecting the results. Subsequently, we will detail the method of synthesizing reference facial images.

We synthesize multiple reference facial images of the same identity as the given facial image using Arc2Face (Paraperas Papantoniou et al. 2024), equipped with ControlNet. Given a pair of facial images (x_{ref}, x_{pose}) , Arc2Face can synthesize facial images that maintain the identity of x_{ref} and the pose of x_{pose} .

Pose reference data pool We begin by constructing the pose-reference pool. Let $X = \{x_i\}_{i=1}^n$ represent the images in the FFHQ dataset (Karras, Laine, and Aila 2019). Using EMOCA v2 (Daněček, Black, and Bolkart 2022; Filntisis et al. 2023), we extract the pose attribute $\theta \in \mathbb{R}^{1 \times 6}$ and the expression attribute $\psi \in \mathbb{R}^{1 \times 50}$ for each image in X . Initially, we conduct K-Means (MacQueen et al. 1967) clustering on X based on θ , setting the number of the cluster centers to c_1 . This divides X into c_1 disjointed parts, denoted as $\{X_i\}_{i=1}^{c_1}$. Subsequently, we conduct K-Means clustering on X_i based on ψ , setting the number of the cluster centers to c_2 . This results in $c_1 \times c_2$ disjoint subsets of X , denoted as $P = \{P_j\}_{j=1}^{c_1 \times c_2}$, forming the pose-reference pool.

Same identity For each x_i , we randomly sample an image p_j from P_j to serve as the pose-reference image. Using (x_i, p_j) as input for Arc2Face, we synthesize the image y_i^j . We then assess the identity similarity between x_i and y_i^j . If the identity similarity falls below δ , we re-sample p_j from P_j for regeneration. If an acceptable y_i^j is not obtained after three attempts, we stop generating it. We name the synthesized reference images for the FFHQ dataset as FFHQRef. The detailed settings can be found in the Supplementary Material.

Experiments

Experimental Setup

Training datasets Our training dataset consists of FFHQ dataset (Karras, Laine, and Aila 2019) and our synthesized FFHQRef dataset, with all images resized to 512×512 . I_h denotes the high-quality image from the FFHQ dataset. To form training pairs, 1~4 images with the same identity as I_h are randomly selected from the FFHQRef dataset as reference images. The corresponding degraded image I_l is synthesized using the following degradation model (Wang et al. 2021; Zhou et al. 2022; Yue and Loy 2024):

$$I_l = \left\{ [(I_h \otimes k_\sigma) \downarrow_r + n_\delta]_{JPEG_q} \right\} \uparrow_r, \quad (6)$$

where \otimes denotes 2D convolution, k_σ denotes a Gaussian kernel with kernel width σ , \downarrow_r and \uparrow_r denote downsampling and upsampling operators with scale r , respectively. n_δ denotes Gaussian noise with standard deviation δ , and $[\cdot]_{JPEG_q}$ denotes the JPEG compression process with quality factor q . We randomly sample σ , r , δ , and q from $[0.2, 10]$, $[1, 16]$, $[0, 15]$, and $[30, 100]$, respectively.

Implementation details We employ the SDXL model (Podell et al. 2023) *stable-diffusion-xl-base-1.0* fine-tuned by PhotoMaker as our base diffusion model. We employ the CLIP image encoder, fine-tuned by PhotoMaker, as part of

our identity encoder. We use the AdamW (Loshchilov and Hutter 2019) optimizer to optimize the network parameters with a learning rate of 5×10^{-5} for two training stages. The training process is implemented using the PyTorch framework and is conducted on eight A40 GPUs, with a batch size of 4 per GPU. The two training stages are trained 130K and 210K iterations, respectively.

Testing datasets We use one synthetic dataset CelebRef-HQ (Li et al. 2022) and three real-world datasets: LFW-Test (Huang et al. 2008), WebPhoto-Test (Wang et al. 2021), and WIDER-Test (Zhou et al. 2022) for test. CelebRef-HQ is collected by crawling images of celebrities from the internet. It contains 1,005 identities and a total of 10,555 images. LFW-Test consists of 1,711 mildly degraded face images from the LFW dataset. WebPhoto-Test consists of 407 medium degraded face images from the internet. WIDER-Test consists of 970 severely degraded face images from the WIDER Face (Yang et al. 2016) dataset.

For the synthetic dataset, we randomly select 150 identities and select one image per identity as the ground truth, using 1~4 images of the same identity as reference images. We employ the same degradation model described in Eq.(6) to synthesize the corresponding degraded images, maintaining the same hyperparameter settings used during training.

For the real-world datasets, due to the lack of reference images with the same identity, we first use face restoration method, *i.e.*, Codeformer (Zhou et al. 2022), to restore low-quality input. The restored images are then used as input to Arc2Face to generate reference images. It is worth noting that this method of synthesizing reference images does not compromise generalization capability. The synthesized images have varying poses and relatively low quality, whereas, in real-world scenarios, the reference images provided by users are often of much higher quality.


Comparisons with State-of-the-art Methods

Comparison methods We compare FaceMe with state-of-the-art methods, including GFP-GAN (Wang et al. 2021), CodeFormer (Zhou et al. 2022), VQFR (Gu et al. 2022), GPEN (Yang et al. 2021), DiffFace (Yue and Loy 2024), and DR2 (Wang et al. 2023), PGDiff (Yang et al. 2024), and DMDNet (Li et al. 2022). Since PGDiff only supports a single reference image, we use the first image in the reference dataset as the reference input during testing.


Evaluation metrics For the synthetic dataset that contains ground truth, we adopt the following metrics for quantitative comparison: full-reference metrics PSNR, SSIM, and LPIPS (Zhang et al. 2018), as well as non-reference metrics MUSIQ (Ke et al. 2021) and FID (Heusel et al. 2017). We also use LMD (landmark distance using L_2 norm) and IDS¹ (cosine similarity with ArcFace (Deng et al. 2019)) to evaluate identity preservation. For real-world datasets without ground truth, we employ the widely-used non-reference perceptual metric FID. All evaluation metrics are measured by PyIQA² except for LMD and IDS.

¹<https://github.com/deepinsight/insightface>

²<https://github.com/chaofengc/IQA-PyTorch>



(a) **Qualitative comparison on CelebRef-HQ.** In comparison to the state-of-the-art methods, our FaceMe can restore high-quality faces while maintaining identity consistency. **Zoom in for best view.**



(b) **Qualitative comparison on real-world faces.** The first row is from the LFW-Test; the second row is from the WebPhoto-Test; and the third row is from the Wider-Test. Our method can restore high-fidelity and high-quality images, while previous methods produce unrealistic artifacts. **Zoom in for best view.**

Figure 3: Comparison of our method with state-of-the-art on different datasets.

Evaluation on synthetic data Tab. 1 shows the performance of FaceMe on the synthetic dataset CelebRef-HQ. As shown, our method achieves the best performance in PSNR, FID, LMD, and IDS, and the second-best performance in LPIPS. Additionally, note that FaceMe has significantly improved in both LMD and IDS, which demonstrates its ability in personalization while minimizing the impact of ID-irrelevant features. The visualization results are shown in Fig. 3a. It can be observed that the compared methods either produce many artifacts in the restored images or restore high-quality images but fail to maintain identity consistency with the ground truth. In contrast, our FaceMe can restore high-quality images while preserving identity consistency.

Evaluation on real-world data As presented in Tab. 1, our FaceMe achieves the best FID score on the LFW-Test and Wider-Test datasets. LFW-Test and Wider-Test are mildly degraded and heavily degraded real-world dataset, respectively. The excellent performance on both datasets indicates that FaceMe is capable of adapting to complex degradation scenarios in the real world, demonstrating exceptional robustness. Fig. 3b shows the visual comparisons of

different methods. In comparison, FaceMe can handle more complex scenes and restore high-quality images without introducing unpleasant artifacts.

Ablation Studies

Due to space limitations, more ablation studies and the complete table for the ablation studies are provided in the Supplementary Material.

The number of reference images We study the effects of different numbers of reference image inputs on the restored results. As shown in Tab. 2a, as the number of reference images increases, the quality of the results improves slightly. The improvement in LMD and IDS metrics indicates that adding reference images enhances identity consistency. We attribute this improvement to the influence of the most effective reference image. When the ground truth (GT) image is used as a reference (0Ref. w GT), the model achieves its best performance. Even when the GT image is mixed with other reference images (3Ref. w GT), the model still benefits positively from GT’s inclusion.

Dataset Method	CelebRef-HQ						LFW-Test	WebPhoto-Test	WIDER-Test	Ref.	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MUSIQ \uparrow	FID \downarrow	LMD \downarrow	IDS \uparrow	FID \downarrow	FID \downarrow		FID \downarrow
GFP-GAN	23.10	0.653	0.249	75.67	59.85	3.736	0.546	51.33	91.51	40.46	
CodeFormer	23.98	0.665	0.214	76.88	<u>52.21</u>	<u>3.012</u>	0.546	53.75	86.14	40.04	
GPEN	24.08	0.686	0.317	70.01	72.96	4.197	0.546	56.02	87.04	47.90	
VQFR	22.75	0.620	0.262	<u>76.26</u>	62.36	3.521	0.541	51.83	77.44	45.02	
DifFace	<u>24.11</u>	<u>0.679</u>	0.255	71.85	53.26	3.293	0.438	<u>46.99</u>	<u>82.65</u>	<u>38.47</u>	
DR2+VQFR	22.71	0.650	0.272	75.66	64.59	4.017	0.402	64.44	117.48	56.84	
PGDiff	21.97	0.636	0.297	68.47	72.59	4.149	<u>0.615</u>	-	-	-	✓
DMDNet	23.71	0.664	0.263	74.70	58.29	5.837	0.559	53.34	89.27	41.98	✓
FaceMe (Ours)	24.37	0.678	<u>0.227</u>	75.62	51.01	2.908	0.647	43.95	92.46	33.59	✓

Table 1: Quantitative comparison. The **bold** and underlined represent the best and the second best performance, respectively.

Number	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	LMD \downarrow	IDS \uparrow
one	24.65	0.233	49.38	2.899	0.620
two	24.59	0.232	48.94	2.914	0.631
three	24.54	0.232	48.52	2.896	0.632
four	24.52	0.232	48.12	2.871	0.634
one*	24.74	0.218	48.19	2.753	0.641
four*	24.56	0.227	48.12	2.813	0.640

(a) Ablation study on the number of reference images. * denote using GT as part of reference.

Training Strategy	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	LMD \downarrow	IDS \uparrow
one-stage training	22.23	0.274	61.84	4.195	0.534
two-stage training	24.37	0.227	51.01	2.908	0.647

(b) Ablation study on the proposed training strategy.

Identity Encoder	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	LMD \downarrow	IDS \uparrow
w/o ID encoder	24.88	0.281	59.87	3.083	0.568
w/ ID encoder	24.37	0.227	51.01	2.908	0.647

(c) Ablation study on the proposed identity encoder.

Table 2: Ablation studies.

Training strategy Tab. 2b presents the comparison of our proposed two-stage training strategy with the one-stage training strategy, *i.e.*, jointly training ControlNet and ID encoder. While the one-stage training strategy can restore high-quality images, it significantly reduces the control over low-quality inputs and compromises identity protection. In Fig. 4a, we observe that the one-stage training produces some unnatural characteristics, such as excessive wrinkles and clearly unrealistic expressions, indicating that the generative capacity is overly strong. In contrast, the two-stage training effectively balances the dependency between low-quality input and reference images.

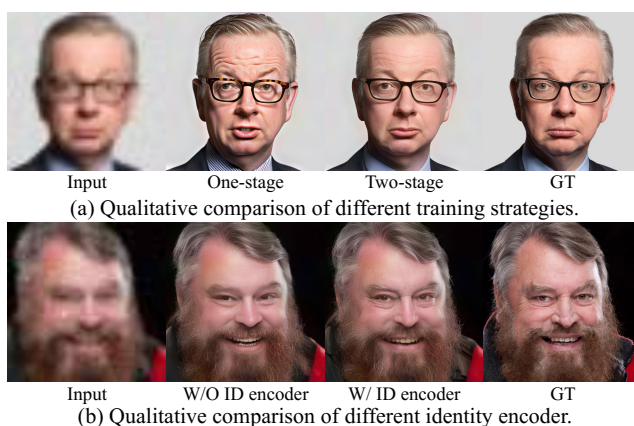


Figure 4: Qualitative comparison of ablation studies.

Identity encoder We compare the proposed identity encoder with not using an identity encoder, as shown in Tab. 2c. Using an identity encoder can significantly improve identity consistency compared to the model without an identity encoder. Additionally, as shown in Fig. 4b, using the proposed identity encoder produces clearer and higher-quality images. In contrast, the absence of an identity encoder results in blurry restored images.

Conclusion

We propose a method to address the issue of identity shift in blind facial image restoration. Based on the diffusion model, we use identity-related features extracted by the identity encoder to guide the diffusion model in recovering face images with consistent identities. Our method supports any number of reference images by combining identity-related features. In addition, the robustness of the identity encoder allows us to use synthetic images as reference images for training. Moreover, our method does not require fine-tuning the model when changing identities. The experimental results demonstrate the superiority and effectiveness of our method.

Acknowledgments

This work is funded by the National Natural Science Foundation of China (62306153) and the Fundamental Research Funds for the Central Universities (Nankai University, 070-63243143). The computational devices of this work is supported by the Supercomputing Center of Nankai University (NKSC).

References

- Chari, P.; Ma, S.; Ostashev, D.; Kadambi, A.; Krishnan, G.; Wang, J.; and Aberman, K. 2023. Personalized Restoration via Dual-Pivot Tuning. *arXiv preprint arXiv:2312.17234*.
- Daněček, R.; Black, M. J.; and Bolkart, T. 2022. Emoca: Emotion driven monocular face capture and animation. In *CVPR*, 20311–20322.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 4690–4699.
- Ding, Z.; Zhang, X.; Tu, Z.; and Xia, Z. 2024. Restoration by Generation with Constrained Priors. In *CVPR*, 2567–2577.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *CVPR*, 12873–12883.
- Filntisis, P. P.; Retsinas, G.; Paraperas-Papantoniou, F.; Katsamanis, A.; Roussos, A.; and Maragos, P. 2023. Spectre: Visual speech-informed perceptual 3d facial expression reconstruction from videos. In *CVPR*, 5745–5755.
- Gu, Y.; Wang, X.; Xie, L.; Dong, C.; Li, G.; Shan, Y.; and Cheng, M.-M. 2022. Vqfr: Blind face restoration with vector-quantized dictionary and parallel decoder. In *ECCV*, 126–143. Springer.
- Guo, Z.; Wu, Y.; Chen, Z.; Chen, L.; Zhang, P.; and He, Q. 2024. PuLID: Pure and Lightning ID Customization via Contrastive Alignment. In *NIPS*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, volume 30.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Huang, G. B.; Mattar, M.; Berg, T.; and Learned-Miller, E. 2008. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *CVPR*, 4401–4410.
- Ke, J.; Wang, Q.; Wang, Y.; Milanfar, P.; and Yang, F. 2021. Musiq: Multi-scale image quality transformer. In *ICCV*, 5148–5157.
- Li, X.; Li, W.; Ren, D.; Zhang, H.; Wang, M.; and Zuo, W. 2020. Enhanced blind face restoration with multi-exemplar images and adaptive spatial feature fusion. In *CVPR*, 2706–2715.
- Li, X.; Liu, M.; Ye, Y.; Zuo, W.; Lin, L.; and Yang, R. 2018. Learning warped guidance for blind face restoration. In *ECCV*, 272–289.
- Li, X.; Zhang, S.; Zhou, S.; Zhang, L.; and Zuo, W. 2022. Learning dual memory dictionaries for blind face restoration. *IEEE TPAMI*, 45(5): 5904–5917.
- Li, Z.; Cao, M.; Wang, X.; Qi, Z.; Cheng, M.-M.; and Shan, Y. 2024. Photomaker: Customizing realistic human photos via stacked id embedding. In *CVPR*, 8640–8650.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- MacQueen, J.; et al. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, 281–297.
- Paraperas Papantoniou, F.; Lattas, A.; Moschoglou, S.; Deng, J.; Kainz, B.; and Zafeiriou, S. 2024. Arc2Face: A Foundation Model for ID-Consistent Human Faces. In *ECCV*.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PMLR.
- Varanka, T.; Toivonen, T.; Tripathy, S.; Zhao, G.; and Acar, E. 2024. PFStorer: Personalized Face Restoration and Super-Resolution. In *CVPR*, 2372–2381.
- Wang, J.; Yue, Z.; Zhou, S.; Chan, K. C.; and Loy, C. C. 2024a. Exploiting diffusion prior for real-world image super-resolution. *IJCV*, 1–21.
- Wang, Q.; Bai, X.; Wang, H.; Qin, Z.; and Chen, A. 2024b. InstantID: Zero-shot Identity-Preserving Generation in Seconds. *arXiv preprint arXiv:2401.07519*.
- Wang, X.; Li, Y.; Zhang, H.; and Shan, Y. 2021. Towards real-world blind face restoration with generative facial prior. In *CVPR*, 9168–9178.
- Wang, Z.; Zhang, J.; Chen, R.; Wang, W.; and Luo, P. 2022. Restoreformer: High-quality blind face restoration from un-degraded key-value pairs. In *CVPR*, 17512–17521.
- Wang, Z.; Zhang, Z.; Zhang, X.; Zheng, H.; Zhou, M.; Zhang, Y.; and Wang, Y. 2023. Dr2: Diffusion-based robust degradation remover for blind face restoration. In *CVPR*, 1704–1713.
- Wu, R.; Yang, T.; Sun, L.; Zhang, Z.; Li, S.; and Zhang, L. 2024a. Seesr: Towards semantics-aware real-world image super-resolution. In *CVPR*, 25456–25467.
- Wu, Y.; Li, Z.; Zheng, H.; Wang, C.; and Li, B. 2024b. Infinite-ID: Identity-preserved Personalization via ID-semantics Decoupling Paradigm. In *ECCV*.
- Yang, P.; Zhou, S.; Tao, Q.; and Loy, C. C. 2024. PGDiff: Guiding Diffusion Models for Versatile Face Restoration via Partial Guidance. *NIPS*, 36.

- Yang, S.; Luo, P.; Loy, C.-C.; and Tang, X. 2016. Wider face: A face detection benchmark. In *CVPR*, 5525–5533.
- Yang, T.; Ren, P.; Xie, X.; and Zhang, L. 2021. Gan prior embedded network for blind face restoration in the wild. In *CVPR*, 672–681.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Yue, Z.; and Loy, C. C. 2024. DifFace: Blind Face Restoration With Diffused Error Contraction. *IEEE TPAMI*, 46(12): 9991–10004.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 586–595.
- Zhou, S.; Chan, K.; Li, C.; and Loy, C. C. 2022. Towards robust blind face restoration with codebook lookup transformer. *NIPS*, 35: 30599–30611.