

# CLR-Face: Conditional Latent Refinement for Blind Face Restoration Using Score-Based Diffusion Models

Maitreya Suin , Rama Chellappa

Johns Hopkins University  
maitreyasuin21@gmail.com, rchella4@jhu.edu

## Abstract

Recent generative methods have shown promising blind face restoration performance. They usually project the degraded images to the latent space and then decode high-quality faces either by single-stage latent optimization or directly from the encoding. Generating fine-grained facial details faithful to inputs remains challenging. Most existing methods produce either overly smooth outputs or alter the identity. This could be attributed to the typical trade-off between quality and resolution in the latent space. If the latent is highly compressed, the decoded output is more robust to degradations but shows worse fidelity. On the other hand, a more flexible latent space can capture intricate details better, but is extremely difficult to optimize for highly degraded faces. We introduce a diffusion-based-prior inside a VQGAN architecture that focuses on learning the distribution over uncorrupted latent embeddings. We iteratively recover the clean embedding conditioning on the degraded counterpart. Furthermore, to ensure the reverse diffusion trajectory does not deviate from the underlying identity, we train a separate Identity Recovery Network and use its output to constrain the reverse diffusion. Specifically, using a learnable latent mask, we add gradients from a face-recognition network to a subset of latent features that correlates with the finer identity-related details in the pixel space, leaving the other features untouched. Disentanglement between perception and fidelity in the latent space allows us to achieve the best of both worlds. We perform extensive evaluations on multiple real and synthetic datasets to validate our approach.

## 1 Introduction

Blind face restoration (BFR) is a challenging problem that aims to recover a high-quality facial image from a low-quality input. With the increasing availability of degraded images captured in various real-world scenarios, the demand for effective BFR techniques has increased significantly in recent years. The primary goal of BFR is to recover undegraded facial features while preserving the person’s identity. A

complex combination of many factors, such as low resolution, blur, noise, compressions, etc., corrupts a facial image. Typically, traditional approaches [Baker and Kanade, 2000; Bourlai *et al.*, 2011] rely on a degradation model and manually designed priors, which often lead to sub-optimal results and restricted ability to handle a variety of real-life images. Recently, generative priors [Wang *et al.*, 2021; Yang *et al.*, 2021] have shown remarkable improvement in restoration performance. Generative priors encapsulated in a well-trained high-quality face generator (e.g., StyleGAN [Karras *et al.*, 2019]) are typically exploited in such approaches. The degraded image is first projected to the latent space to get a ‘cleaner’ latent vector either directly or after certain refinement operations. The decoder projects the latent vector back to the pixel space, producing a restored image. Instead of dealing with a continuous latent space, [Gu *et al.*, 2022; Zhou *et al.*, 2022] utilizes pretrained vector-quantized codebooks. The mechanism of vector quantization in a compressed latent space reduces the uncertainty of the task, making these methods robust to various degradations. Intuitively, these approaches exploit information from two sources when restoring an image. The degraded input image contains crucial information about the person’s identity. On the other hand, the pretrained decoder is utilized for a high-quality generation. Thus, ultimately, the task boils down to predicting the corresponding clean latent given a degraded image/latent. Using a highly compressed latent space usually reduces the complexity of the prediction task and prioritizes high-quality generation but fails to maintain intricate facial details essential for identity preservation. It may also lead to repeated and less diverse texture patterns in the output. On the other hand, using a less compressed latent space can be potentially more expressive and flexible but is often extremely difficult to optimize when directly projecting the degraded image to the latent space or trying to predict the correct latent code as a classification problem. A lower compression often diminishes the advantages of latent-space-based generative models, leading to suboptimal degradation removal from the input. Similar behavior was recently discussed in [Gu *et al.*, 2022], where the trade-off between fidelity and quality depending on the compression factor was analyzed. We argue that the performance of such face restoration approaches mainly depends on the nature and size of the latent space. It is non-trivial to balance the demands of face generation and restoration us-

ing existing latent refinement techniques. [Gu *et al.*, 2022; Zhou *et al.*, 2022] deploy a highly spatially-compressed latent space (downsampled by 32) and instead use skip-connections to improve the fidelity, which is often counter-productive for the output quality, as the information from the encoder is usually corrupted.

In this work, we design an alternative approach, where we first select a highly expressive latent space with lower compression and to handle its additional complexity, we introduce diffusion-based-prior to model it. We learn the clean latent space distribution by directly modeling the gradient of the log density of the data, known as the score function [Song and Ermon, 2019], using a neural network. Such gradient information can be utilized in reverse by stochastic sampling to generate diverse samples from the underlying distribution. Unlike unconditional generation as in [Song *et al.*, 2020], for the current task, we model it as a conditional generation task, where given degraded latent, we perform multi-step refinement to produce the cleaner counterpart.

Our approach has several key advantages. 1. Similar to prior state-of-the-art (SOTA) blind face restoration works [Zhou *et al.*, 2022], we also utilize a pre-trained VQVAE framework to produce high-quality face images given correct latent codes. 2. The significantly lower compression ration allows us to maintain higher fidelity, suitable for BFR task. 3. Compared to the single-step refinement strategy of prior arts, multi-step conditional refinement of the corrupted latent using score-based diffusion prior is far more effective in producing the clean embedding, enabling us to use relatively mild compression rates, achieving more faithful and sharper reconstructions. 4. Compared to pixel-space-based diffusion models used in [Saharia *et al.*, 2022] (DDPM), our score-based latent-diffusion model has much better time complexity and requires significantly fewer refinement steps ( $\approx 20\times$ ).

The diffusion prior is extremely powerful for producing a clean latent. However, during multiple refinements, deviating from the trajectory of correct identity is plausible. To further constraint the refinement process, we deploy a separate Identity Recovery Network (IRN), which prioritizes recovering identity-specific intricacies given the input image using a face recognition loss. We utilize the output of the IRN to initialize the iterative refinement process. Next, at each step, we utilize a gradient-based score-update strategy using a pre-trained face recognition network to steer the refinement process along the identity recovered by the IRN. While this approach considerably diminishes the risk of identity distortion, the guidance signal from a face-recognition network may not align well with the other objective of producing visually pleasing results. Such a dilemma is observed in existing works as well [Wang *et al.*, 2021], where usually a small amount of identity loss is introduced while prioritizing standard perceptual losses, which ultimately fails to recover the identity information adequately. Instead of blindly enforcing the identity constraint across the entire latent space, we introduce a learnable latent mask, where the gradient from the recognition network is harnessed to selectively update specific latent locations that contribute most significantly to identity information. This enables us to balance the restored identity and the perceptual quality derived from the diffusion prior.

To summarize, our main contributions are

- (1) We propose a diffusion-prior-based conditional latent refinement strategy for recovering the clean latent inside a pre-trained VQVAE framework.
- (2) The strong modeling capability of such iterative refinement allows us to design a latent space with a lower spatial-compression ratio ( $f=4$ ), improving the overall fidelity and sharpness of the output.
- (3) We further deploy an IRN to recover identity-specific features from the degraded input. The output of IRN, along with adaptive guidance using a learnable latent mask, steers the refinement process for better identity recovery without compromising the perceptual quality.
- (4) Extensive evaluation on multiple real and synthetic datasets demonstrates our proposed framework’s superior perceptual and identity-preserving properties.

## 2 Related Works

Various priors have been proposed, including geometric priors such as facial landmarks [Kim *et al.*, 2019; Chen *et al.*, 2018; Zhu *et al.*, 2016] and parsing maps [Chen *et al.*, 2021; Shen *et al.*, 2018], as well as reference-based approaches [Li *et al.*, 2020; Dogan *et al.*, 2019]. However, these priors often struggle to estimate necessary information from corrupted images. Recent generative priors optimize the latent vector for GAN inversion techniques [Menon *et al.*, 2020; Gu *et al.*, 2020] or direct projection of the input image to the latent space [Richardson *et al.*, 2021]. [Yang *et al.*, 2021], and [Wang *et al.*, 2021] exploited the generative prior inside an encoder-decoder, with structural details from the degraded input through skip connections. For severe degradations, such links lead to unwanted artifacts.

Diffusion and score-based models have shown improvements recently [Saharia *et al.*, 2022; Whang *et al.*, 2022; Choi *et al.*, 2021; Yue and Loy, 2022; Chung *et al.*, 2022]. Iterative refinement strategies has been adopted by [Saharia *et al.*, 2022], [Whang *et al.*, 2022] for super-resolution and motion deblurring tasks. [Choi *et al.*, 2021] guided the reverse process with low-frequency information from a conditional image. However, such a conditioning strategy does not translate well for the BFR problem with significant degradation and may alter the identity. [Yue and Loy, 2022] uses an unconditional diffusion model and starts from an intermediate stage of the reverse diffusion process. But, as the underlying diffusion model is unconditional, the identity of the restored face changes considerably. If it is used for a smaller timespan, the visual quality and sharpness of the output suffer considerably. [Chung *et al.*, 2022] addressed only non-blind super-resolution tasks, and its identity-preserving capability is yet to be tested for more difficult BFR scenarios. Our work mainly aims to balance identity preservation and perceptual quality for heavily degraded faces.

## 3 Method

Our framework has two main parts, a Vector-Quantized (VQ) autoencoder is our backbone that maps a degraded image to the latent space with a downsampling factor  $f = 4$  and reverses this operation after refining the latent. A conditional

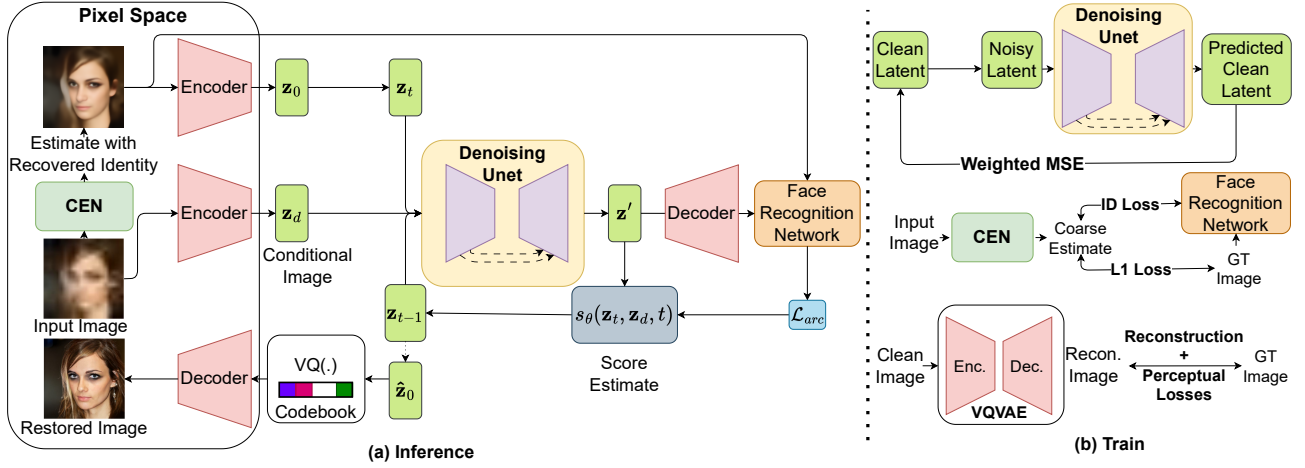


Figure 1: An overview of our inference (left) and training framework (right).

diffusion model (denoising UNet) is used for the latent refinement task. We pre-train the quantized autoencoder through self-reconstruction to obtain a discrete codebook and the corresponding decoder for projecting the restored latent back to pixel space. In addition, we also deploy IRN to produce an initial estimate of the restored face, focusing on recovering identity-specific features, which are used to steer the iterative latent refinement process. An overview of our framework is shown in Fig. 1. In the following sections, we describe the details of different parts of our framework.

### 3.1 Learning Latent Space of Images via Vector-Quantized Codebook

We use  $x$  and  $\mathbf{z}$  to represent image in pixel space and the corresponding latent representation, respectively. Our perceptual compression model utilizes a variant of the Vector-Quantize codebook called VQGAN [Esser *et al.*, 2021]. It learns a perceptual codebook through a combination of perceptual loss [Johnson *et al.*, 2016] and adversarial training [Isola *et al.*, 2017], preventing blurriness that can occur when relying solely on pixel-space  $L_2$  or  $L_1$  losses. The model consists of an encoder  $E$ , a decoder  $G$ , and a codebook  $\mathcal{Z} = \{z_k\}_{k=1}^K \in \mathbb{R}^{K \times d}$  containing a finite set of embedding vectors. Here,  $K$  represents the codebook size, and  $d$  denotes the dimension of the codes. Given an input image  $x \in \mathbb{R}^{3 \times H \times W}$ , we project it to the latent space as  $\mathbf{z} = E(x) \in \mathbb{R}^{d \times \frac{H}{f} \times \frac{W}{f}}$ . Next, a spatial collection of image tokens  $\mathbf{z}_q$  is obtained using a spatial-wise quantizer  $Q(\cdot)$ , which maps each spatial feature  $z_{ij}$  into its closest codebook entry  $z_k$

$$\mathbf{z}_q = Q(\mathbf{z}) = \left( \underset{z_k \in \mathcal{Z}}{\operatorname{argmin}} \|z_{ij} - z_k\|_2^2 \right) \in \mathbb{R}^{d \times \frac{H}{f} \times \frac{W}{f}} \quad (1)$$

The decoder reconstructs the image from the latent, giving  $\tilde{x} = G(\mathbf{z}_q)$ . The encoder  $E$ , the decoder  $G$  and the codebook  $\mathcal{Z}$  are trained end-to-end via the following loss function

$$\mathcal{L} = \|x - \tilde{x}\|_1 + \|\phi(x) - \phi(\tilde{x})\|_1 + [\log D(x) + \log(1 - D(\tilde{x}))] + \|\operatorname{sg}(E(x)) - \mathbf{z}_q\|_2^2 + \|\operatorname{sg}(\mathbf{z}_q) - E(x)\|_2^2 \quad (2)$$

where  $\operatorname{sg}(\cdot)$  stands for the stop-gradient operator,  $\phi(\cdot)$  represents the VGG-based [Simonyan and Zisserman, 2014] feature extractor. The first three terms help reconstruct the input image, whereas the last two terms reduce the distance between codebook  $\mathcal{Z}$  and input feature embeddings  $\mathbf{z}$ . As we build upon the foundation of continuous space diffusion models, we absorb the quantization layer in the decoder. Given  $\mathbf{z}$ , we use the conditional diffusion model to recover the correct latent embedding through iterative refinement. Next, quantization and the pre-trained high-quality codebook can handle any inaccuracies in the latent prediction for producing a clean, restored face. Compared to existing BFR works that use VQVAE [Gu *et al.*, 2022; Zhou *et al.*, 2022], our work has two critical differences: we use a milder compression (downsample  $f = 4$  instead 32) and a larger codebook size ( $K = 8192$  instead of 1024). Such an expressive yet highly complex latent space can be optimally modeled by iterative diffusion-prior-based refinement. As a result, the restored image is of higher fidelity and realism than prior art.

### 3.2 Latent Refinement Using Diffusion Prior

Given the high dimensional latent with a lower compression ratio, existing methods encounter challenges in generating a cleaner latent through a single pass in the encoder or a classification-based refinement module, as elaborated in [Gu *et al.*, 2022]. Instead, we aim to learn the probability distribution of the clean latent space (through gradients of log probability density functions) and map the degraded latent to the clean space in multiple refinement steps. Our approach is built upon the framework of [Song and Ermon, 2019], where stochastic differential equations drive the underlying diffusion process. We can represent this process using continuous time variables  $\mathbf{z}_t, t=0,1$ , where  $\mathbf{z}_0$  represents the initial variable (i.e., clean latent representation) and  $\mathbf{z}_t$  represents its noisy version at time  $t$ . The diffusion process is defined by an Itô SDE [Rogers and Williams, 2000]

$$d\mathbf{z}_t = f(\mathbf{z}_t, t)dt + g(t)d\mathbf{w}_t \quad (3)$$

where  $f(\cdot, \cdot)$  and  $g(\cdot)$  are the drift and diffusion coefficients, respectively, and  $\mathbf{w}_t$  denotes the standard Brownian motion.

We denote the distribution of  $\mathbf{z}_t$  as  $p(\mathbf{z}_t)$ . The SDE in Eq. 3 can be converted to a generative model by first sampling from  $\mathbf{z}_1 \sim \mathcal{N}(z_1, 0, I)$  and then running the reverse-time SDE as

$$d\mathbf{z}_t = [f(\mathbf{z}_t, t) - \frac{1}{2}g^2(t)\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t)]dt + g(t)d\bar{\mathbf{w}}_t \quad (4)$$

where  $\bar{\mathbf{w}}_t$  is a reverse-time standard Wiener process and  $dt$  is an infinitesimal negative time step. As shown in [Song *et al.*, 2020], there exists a deterministic ordinary differential equation (ODE), whose trajectories share the same marginal probability densities as the SDE (Eqs. 3, 4)

$$d\mathbf{z}_t = [f(\mathbf{z}_t, t) - g^2(t)\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t)]dt \quad (5)$$

In practice, blackbox ODE solvers can be used on Eq. 5 for faster and high quality sampling, assuming we know the score function. The score function ( $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t)$ ) is typically approximated using a neural network  $s_\theta(\mathbf{z}_t, t)$  [Vincent, 2011; Song and Ermon, 2019; Ho *et al.*, 2020], trained to minimize the weighted Fisher’s divergence with a positive weighting function  $\lambda$  as

$$\min_{\theta} \mathbb{E}_{t \sim \mathbb{U}[0,1]} (\lambda(t) \|\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t) - s_\theta(\mathbf{z}_t, t)\|_2^2) \quad (6)$$

### Conditional Diffusion for BFR

Instead of unconditionally sampling from the underlying clean latent distribution using Eq. 5, for the task of BFR, our goal is to predict the clean counterpart, given a degraded latent  $\mathbf{z}^d$  of the input image  $\mathbf{x}^d$ . We are interested in  $p(\mathbf{z}_0|\mathbf{z}^d)$ , where  $\mathbf{z}_0$  is the target clean latent distribution corresponding to the image  $\mathbf{x}^d$  and  $\mathbf{z}^d$  is the condition signal. Similar to Eq. 5, we can obtain a conditional reverse-time ODE as

$$d\mathbf{z}_t = [f(\mathbf{z}_t, t) - g^2(t)\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t|\mathbf{z}^d)]dt \quad (7)$$

We need to learn the conditional score function  $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t|\mathbf{z}^d)$  in order to be able to sample from Eq. 7 using reverse-time diffusion. In [Song *et al.*, 2020],  $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t|\mathbf{z}^d)$  was estimated using

$$\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t|\mathbf{z}^d) = \nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t) + \nabla_{\mathbf{z}_t} \log p(\mathbf{z}^d|\mathbf{z}_t) \quad (8)$$

where  $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t)$  is learned using an unconditional model and  $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}^d|\mathbf{z}_t)$  is learned using a separate auxiliary network. But, modeling such a forward process is non-trivial for BFR, as there could be many possible  $\mathbf{z}^d$  given a  $\mathbf{z}_t$ . Instead, following the intuition of [Saharia *et al.*, 2022; Tashiro *et al.*, 2021], we can extend the formulation of Eq. 6 to a conditional case as

$$\min_{\theta} \mathbb{E}_{t \sim \mathbb{U}[0,1]} (\lambda(t) \|\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t) - s_\theta(\mathbf{z}_t, \mathbf{z}^d, t)\|_2^2) \quad (9)$$

where we pass the condition  $\mathbf{z}^d$  directly to the denoising UNet, making it learn the conditional score. We empirically verify that it works well for BFR task. In [Song *et al.*, 2020; Kadkhodaie and Simoncelli, 2020], the noise was added to both the target and the conditional observations, which we found to be harmful. We follow the simplified configuration of [Karras *et al.*, 2022], where we set  $f(\mathbf{z}_t, t) = 0$  and  $g(t) = \sqrt{2t}$ . Then,  $p(\mathbf{z}_t)$  can be expressed as  $p(\mathbf{z}_0) * \mathcal{N}(0, t^2I)$ ,

where  $*$  represents convolution operation and  $p(\mathbf{z}_0)$  is the underlying distribution of the clean or non-degraded images’ latent space, which we aim to model. Once we know the score function (using a neural network), Eq. 7 can be re-written as

$$d\mathbf{z}_t = -ts_\theta(\mathbf{z}_t, \mathbf{z}^d, t) \quad (10)$$

During the reverse process, we can proceed backward in time with a black box ODE solver, such as Heun solvers [Karras *et al.*, 2022], to obtain the solution trajectory. The resulting  $\hat{\mathbf{z}}_0$  can be considered an approximate embedding of the clean image corresponding to  $\mathbf{x}^d$ .  $\hat{\mathbf{z}}_0$  is passed through the pre-trained decoder to produce the restored image. The decoder has the quantization layer, which handles minor errors in the continuous space prediction.

### 3.3 Identity Recovery Network (IRN)

Existing works typically introduce an identity-preserving loss to the same restoration network in addition to standard perceptual losses. But, often, minimizing both losses together hampers the perceptual quality and may introduce unwanted artifacts, especially if the coefficient for identity loss is high. Instead, to ease the learning process, we aim to disentangle these two objectives and use a separate IRN ( $f_\phi$ ) to focus on recovering the identity. We train  $f_\phi$ , emphasizing the identity-preserving loss, with a small amount of  $L_1$  loss for better stability. The loss function for training  $f_\phi$  is as follows:

$$\mathcal{L}_{IRN} = \alpha L_1(f_\phi(\mathbf{x}^d), \mathbf{x}) + D_{cos}(f_{arc}(f_\phi(\mathbf{x}^d)), f_{arc}(\mathbf{x})) \quad (11)$$

where  $D_{cos}$  denotes the cosine distance between two feature vectors and  $f_{arc}$  denotes a pre-trained ArcFace model. We denote the second term in Eq. 11 as  $\mathcal{L}_{arc}$ . The IRN aims to generate a stable approximation  $\mathbf{x}^{id}$  that primarily recovers the identity information rather than focusing on visually pleasing sharp results. Such crucial identity features are also transferred to the latent embedding  $\mathbf{z}^{id}$ . To ensure that during reverse diffusion process, the output does not deviate from the already recovered identity in  $\mathbf{x}^{id}$ , we add a gradient term  $\nabla_{\hat{\mathbf{z}}_{0,t}} \mathcal{L}_{arc}$  to the estimated score at each step, where the  $\mathcal{L}_{arc}$  is calculated between the image obtained after projecting  $\hat{\mathbf{z}}_{0,t}$  to pixel-space using the pretrained decoder and  $\mathbf{x}^{id}$ ,  $\hat{\mathbf{z}}_{0,t}$  is the denoised output at step  $t$ . It follows the intuition of classifier guidance in [Dhariwal and Nichol, 2021]. But, we have observed that simply adding the gradient for the whole latent again deteriorates the perceptual quality. We hypothesize that the latent embedding should have specific features that correlate to identity-specific details in the pixel space. We should ideally regularize the score function of only those features while keeping the rest untouched.

To verify this, we introduce a learnable latent mask  $M_l \in \mathbb{R}^{d \times \frac{H}{F} \times \frac{W}{F}}$ , denoting the identity-specific latent features. We use a stack of a few convolutional layers followed by a sigmoid layer as  $f_M$ , that takes the  $\hat{\mathbf{z}}_t^i$  as input and produces  $M_l$ . To train  $f_M$ , we first update the score to obtain the corresponding  $\hat{\mathbf{z}}_{0,t}$  and project it back to pixel space. We use LPIPS, cosine-distance-based ID loss on the output image, and a sparsity constraint on  $M_l$  to ensure  $M_l$  selects only the subset of identity-specific features without harming the LPIPS score. As discussed in our ablation study, the observed



Figure 2: Qualitative comparisons on CelebA-Test set for BFR.


 Figure 3: Qualitative comparisons on CelebA-Test set for  $\times 32$  upsampling. Although the input is severely degraded, our approach works better than existing works in restoring the face faithfully.

Methods	FID ↓	LPIPS ↓	IDS ↑	LDM ↓	PSNR ↑	SSIM ↑
GPEN	101.12	0.3362	0.4022	10.76	22.43	0.6009
GFPGAN	99.03	0.2812	0.4633	10.49	<u>22.50</u>	<u>0.6060</u>
PSFRGAN	64.81	0.2513	0.3983	6.51	21.75	0.5450
CodeFormer	54.41	0.2288	<u>0.5009</u>	5.85	22.35	0.5736
DifFace	<u>52.18</u>	<u>0.2061</u>	0.4833	<u>5.40</u>	<b>22.74</b>	<b>0.6116</b>
RestoreFormer	59.51	0.2899	0.4041	7.65	21.73	0.5256
VQFR	60.75	0.3147	0.37	8.32	21.02	0.4931
Ours	<b>46.15</b>	<b>0.1868</b>	<b>0.5667</b>	<b>5.00</b>	21.66	0.5652

 Table 1: Quantitative evaluation on 3000 images of size  $512 \times 512$  from the CelebA-Test (BFR). Bold and underline indicate the best and the second best performance.

improvement after introducing the latent mask validates our design choice. For all the cases, straight-through gradient estimator [Esser *et al.*, 2021] is utilized to copy the gradients from the decoder to the continuous latent embedding to handle the non-differentiable quantization operation.

## 4 Experimental Results

### 4.1 Training Dataset

The FFHQ dataset [Karras *et al.*, 2019] contains 70,000 high-quality face images at  $1024 \times 1024$  resolution. We resized the images to  $512 \times 512$  for training. We synthesized degraded images on the FFHQ dataset using the degradation model proposed in [Yang *et al.*, 2021; Wang *et al.*, 2021; Wang *et al.*, 2022; Chen *et al.*, 2021]:  $x^d = ((x \otimes k) \downarrow_s + n_\sigma)_q$ . Here,  $x$ ,  $x^d$ ,  $k$ ,  $n_\sigma$ ,  $s$ , and  $q$  are the clean face image, the corresponding degraded image, the blur kernel, the Gaussian noise with standard deviation  $\sigma$ , downscaling factor, and the JPEG-compression quality factor, respectively. We randomly and uniformly sampled  $\sigma$ ,  $s$ , and  $q$  from  $[0, 20]$ ,  $[1, 32]$ , and  $[30, 90]$ , respectively. We evaluated our approach

Methods	LPIPS ↓	FID ↓	IDS ↑	LDM ↓	PSNR ↑	SSIM ↑
GPEN	0.4350	148.39	0.1843	21.80	19.91	<u>0.5346</u>
GFPGAN	0.4028	160.29	0.2243	23.10	<u>19.95</u>	<b>0.5366</b>
CodeFormer	0.3565	73.45	0.2546	11.24	19.14	0.4639
DifFace	<u>0.3001</u>	<u>53.93</u>	<u>0.2892</u>	8.85	<b>20.12</b>	0.5314
RestoreFormer	0.4193	103.13	0.1438	14.10	19.26	0.4581
VQFR	0.4090	109.97	0.1583	15.30	18.88	0.4220
Ours	<b>0.2574</b>	<b>39.71</b>	<b>0.3512</b>	<b>7.67</b>	18.77	0.4751

 Table 2: Quantitative evaluation on 3000 images from the CelebA-Test for extreme upsampling ( $16 \rightarrow 512$ ). Bold and underline indicate the best and the second best performance.

on 3,000 synthetic images from CelebA-Test dataset [Karras *et al.*, 2017]. The degraded images in this dataset were synthesized using the same degradation range as our training. Additionally, we tested our method on real-world datasets: WebPhoto-Test [Wang *et al.*, 2021] (407 images), WIDER Face [Yang *et al.*, 2016] (970 images), Celeb-Child [Wang *et al.*, 2021] (180 images) and TURB. For the TURB dataset, we randomly sampled 139 images from the BRIAR [Cornett *et al.*, 2023] and LRFID [Miller *et al.*, 2019] datasets, which provided a more challenging scenario as our models were not trained on severe turbulence-affected images.

### 4.2 Evaluation Metrics

To assess the quality of the restored images quantitatively, we primarily rely on the Fréchet Inception Distances (FID) [Heusel *et al.*, 2017] and Learned Perceptual Image Patch Similarity (LPIPS) [Zhang *et al.*, 2018] metrics. We also compute PSNR and SSIM for completeness, although they often fail to capture visual quality. To evaluate the face recognition performance of the restored images, we calculate the cosine similarity between the features of the restored image and the paired GT image (IDS). Higher cosine similarity val-






Figure 4: Qualitative comparisons on real-world datasets. The first two rows represent images from WIDER face dataset, the third row represents images from WebPhoto, respectively.



Figure 5: Qualitative comparisons on real-world images from CelebA-Child for image colorization.

Methods	WIDER Face	WebPhoto	CelebA-Child
PSFRGAN	49.85	88.45	107.40
GPEN	46.99	81.77	109.55
GFPGAN	39.76	87.35	111.78
CodeFormer	39.21	78.87	116.18
VQFR	44.54	<u>75.46</u>	<u>105.68</u>
DiffFace	37.49	85.52	110.81
Ours	<b>34.25</b>	<b>75.05</b>	<b>104.40</b>

Table 3: Quantitative comparisons of FID ( $\downarrow$ ) on real-world datasets in terms of FID. Bold and underline indicate the best and the second best performance, respectively.

ues indicate better recovery of the identity information. We use the same evaluation metrics and pre-trained models for LPIPS, FID, and ArcFace as used in the prior art [Wang *et al.*, 2022; Wang *et al.*, 2021]. We also adopt landmark distance (LMD) as the fidelity metric, following [Gu *et al.*, 2022].

### 4.3 Comparisons with State-of-the-Art Methods

We compare with the following state-of-the-art (SOTA) methods: PSFRGAN [Chen *et al.*, 2021], GPEN [Yang *et al.*, 2021], GFPGAN [Wang *et al.*, 2021], CodeFormer [Zhou *et al.*, 2022], RestoreFormer [Wang *et al.*, 2022], DiffFace [Yue and Loy, 2022] and VQFR [Gu *et al.*, 2022]. We use the official results and checkpoints provided by the authors.

**Synthetic BFR:** We reported restoration accuracy on the synthetic CelebA-Test dataset for BFR task in Table 1. Our approach achieves a much better balance between fidelity (IDS, LDM) and perceptual metrics (LPIPS, FID).

**Extreme Upscaling:** We tested the algorithms under extreme BFR conditions, where we applied a fixed downscaling factor of  $\times 32$  to  $512 \times 512$  images, resulting in degraded images of size  $16 \times 16$ . To make the task more challenging, we added noise and blur to the images. The quantitative results can be found in Table 2. Despite the limited information in the in-

Methods	LPIPS $\downarrow$	FID $\downarrow$	IDS $\uparrow$	PSNR $\uparrow$
VQVAE-Rec (f=32)	0.0559	22.76	00.9012	29.46
VQVAE-Rec (f=4)	<b>0.0068</b>	<b>9.02</b>	<b>0.9903</b>	<b>35.72</b>
VQVAE-Rest (f=4)	0.2127	81.69	0.5258	19.09
IRN	0.3344	127.53	0.5155	<b>24.26</b>
IRN <sup>ID</sup>	0.3109	118.50	<b>0.5849</b>	24.16
Diff. Prior	0.1901	47.91	0.5107	21.06
IRN <sup>ID</sup> + Diff. Prior	0.2004	53.22	0.5712	22.34
IRN <sup>ID</sup> + Diff. Prior + $M_I$	<b>0.1868</b>	<b>46.15</b>	0.5667	21.66

Table 4: Quantitative comparison of different ablations of our network on CelebA-Test set. ID represents: ArcFace [Deng *et al.*, 2019] based identity loss. The first two-row represents reconstruction of GT images using VQ-VAE/GAN.

put, our approach outperforms other algorithms in terms of perceptual quality and recognition accuracy.

**Real-World Cases:** Our approach outperforms existing GAN-based methods like GPEN and GFPGAN in terms of producing more realistic and faithful reconstructions on real-world degraded datasets. CodeForm, VQFR’s outputs show visible artifacts or repetitive skin/hair texture due to highly compressed latent space. DiffFace outputs are less sharp and may alter facial details as the underlying model is unconditional. In contrast, our approach achieves superior results with fewer artifacts, even for low to medium degradation levels. Real-world colorization examples are shown in Fig. 5, for which we finetune our network for colorization. We also report the accuracy for the downstream face-recognition task on the TURB dataset in Table 5. Our approach comfortably outperforms prior arts for recognition, as well.

## 5 Ablation Analysis

In Table 4, we analyze the effect of individual components of our approach on the perceptual quality and identity-preserving aspect. We use CelebA-Test with  $512 \times 512$  images for our ablation. We empirically found that the IRN is



Figure 6: (Left): Reconstruction performance of VQVAE. (Right): Restoration performance of VQVAE, GD, IRN and our final model using latent space refinement.

model agnostic, and any SOTA restoration network results in a comparable accuracy boost. We provided detailed comparisons of possible options for IRN in the supplementary. We finally select SwinIR [Liang *et al.*, 2021] that keeps a good balance between performance and accuracy. We design our VQ backbone with the following settings:  $f=4$ , codebook = 8192, latent =  $3 \times 128 \times 128$ , feat =  $128 (\times 1, \times 2, \times 3)$ . For the diffusion model, we have used 40 time steps with heun sampler [Karras *et al.*, 2022] and a UNet-based denoiser based on GD<sup>1</sup>. More implementation and network details, additional results are provided in the supplementary material.

**Compression Factor:** First, we compared VQVAE models with different latent sizes:  $f = 32$  and  $f = 4$ . The  $f = 32$  models, commonly used in previous works [Gu *et al.*, 2022], resulted in poor reconstruction quality due to significant spatial compression to  $16 \times 16$ . In contrast, the  $f = 4$  models achieved superior reconstruction quality, as shown in Fig. 6 (Left). We have verified that it is satisfactorily accurate in reconstructing degraded images as well (PSNR=37.9). Higher spatial compression led to the loss of fine facial details and overly smooth outputs. For example, in the first row, the eye color was changed, and in the second row, both the eye and hair patterns were altered in the  $f = 32$  models.

**Latent Refinement:** We attempted to fine-tune the VQVAE model (with  $f = 4$ ) directly for restoration, but it often failed to remove degradation effectively, as reported in [Gu *et al.*, 2022]. To address this issue, we developed a modified version called VQVAE-Rest. This variant takes both the degraded image and the coarse estimate from the IRN as inputs, enabling it to focus on recovering residual details. However, even with this modification, the output of VQVAE-Rest still exhibited lower quality and noticeable artifacts (Figure 6 (Right), Column 2). These results show the need for appropriate latent-space refinement to achieve satisfactory restoration outcomes. We also compare with Guided-Diffusion (GD) for pixel-space-based diffusion (DDPM), which is run for 1000 timesteps in Figure 6 (Right). As can be observed, it is not only computationally intensive due to large no. of time steps, but also fails to reconstruct the image properly starting from pure noise. The outputs of the IRN shown in Figure 6 (Right), Column 4, may lack sharp details. However, it recovers crucial facial locations like eye, nose, etc.

**ID Loss:** Table 4 demonstrates the significance of the ArcFace [Deng *et al.*, 2019] based identity-preserving loss for the IRN (IRN<sup>ID</sup> vs. IRN), which prioritizes the recovery of identity-specific features, which are valuable for regularizing the diffusion process.

**Adaptive Mask:** While the diffusion model alone can

Methods	Top-1 (↑)	Top-3 (↑)	Top-5 (↑)
GPEN	32/50.6	50/72	62/81
GFPGAN	26/57	58/79	60/85
CodeFormer	28/61	52/77	58/82
RestoreFormer	20/62	50/81	62/88
VQFR	34/60	60/79	64/87
Ours	34/68	68/81	72/90

Table 5: Face recognition accuracy using pre-train ArcFace [Deng *et al.*, 2019] on real-world BRIAR/LRFID dataset. Our method performs best for such downstream task as well.

achieve good perceptual quality, it may inadvertently alter certain identity-specific facial features. Adding IRN<sup>ID</sup>-based gradient improves the IDS, but is detrimental for perceptual metrics. The final output could be slightly different from the output of IRN as we regularize only a subset of the latent in the diffusion stage. Our final model, combining the strengths of the Diff. Prior, IRN<sup>ID</sup> and the learnable latent mask  $M_l$  achieves superior accuracy both in terms of perceptual metrics and fidelity. Inference time(s) for GFPGAN, CodeFormer, DiffFace, Guided-Diffusion and ours are: 0.5, 0.15, 6,30, 0.61. Our design is still much faster than existing diffusion-based BFR methods, and can potentially be further accelerated by reducing the number of steps using knowledge-distillation techniques [Song *et al.*, 2023]. We also have comparable or less parameters (16M) compared to VQFR (76.3M), DiffFace (17.5), RestoreFormer (12M).

## 6 Conclusions

We propose a iterative latent-space refinement technique using diffusion prior for restoring severely degraded face images, achieving a better balance between restoration quality and identity recovery compared to existing methods. However, the performance of our approach is limited by the IRN’s ability to preserve identity information. To improve efficiency, we plan to explore the development of a recognition-model-free reverse algorithm in future research.

## Acknowledgements

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via [2022-21102100005]. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U. S. Government. The US. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

<sup>1</sup>GD: <https://github.com/openai/guided-diffusion>

## References

- [Baker and Kanade, 2000] Simon Baker and Takeo Kanade. Hallucinating faces. In *Proceedings Fourth IEEE international conference on automatic face and gesture recognition (Cat. No. PR00580)*, pages 83–88. IEEE, 2000.
- [Bourlai *et al.*, 2011] Thirimachos Bourlai, Arun Ross, and Anil K. Jain. Restoring degraded face images: A case study in matching faxed, printed, and scanned photos. *IEEE Transactions on Information Forensics and Security*, 6(2):371–384, 2011.
- [Chen *et al.*, 2018] Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, and Jian Yang. Fsrnet: End-to-end learning face super-resolution with facial priors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2492–2501, 2018.
- [Chen *et al.*, 2021] Chaofeng Chen, Xiaoming Li, Lingbo Yang, Xianhui Lin, Lei Zhang, and Kwan-Yee K Wong. Progressive semantic-aware style transformation for blind face restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11896–11905, 2021.
- [Choi *et al.*, 2021] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021.
- [Chung *et al.*, 2022] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12413–12422, 2022.
- [Cornett *et al.*, 2023] David Cornett, Joel Brogan, Nell Barber, Deniz Aykac, Seth Baird, Nicholas Burchfield, Carl Dukes, Andrew Duncan, Regina Ferrell, Jim Goddard, et al. Expanding accurate person recognition to new altitudes and ranges: The briar dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 593–602, 2023.
- [Deng *et al.*, 2019] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019.
- [Dhariwal and Nichol, 2021] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [Dogan *et al.*, 2019] Berk Dogan, Shuhang Gu, and Radu Timofte. Exemplar guided face image super-resolution without facial landmarks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.
- [Esser *et al.*, 2021] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [Gu *et al.*, 2020] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3012–3021, 2020.
- [Gu *et al.*, 2022] Yuchao Gu, Xintao Wang, Liangbin Xie, Chao Dong, Gen Li, Ying Shan, and Ming-Ming Cheng. Vqfr: Blind face restoration with vector-quantized dictionary and parallel decoder. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVIII*, pages 126–143. Springer, 2022.
- [Heusel *et al.*, 2017] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [Isola *et al.*, 2017] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [Johnson *et al.*, 2016] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016.
- [Kadkhodaie and Simoncelli, 2020] Zahra Kadkhodaie and Eero P Simoncelli. Solving linear inverse problems using the prior implicit in a denoiser. *arXiv preprint arXiv:2007.13640*, 2020.
- [Karras *et al.*, 2017] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [Karras *et al.*, 2019] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [Karras *et al.*, 2022] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *arXiv preprint arXiv:2206.00364*, 2022.
- [Kim *et al.*, 2019] Deokyun Kim, Minseon Kim, Gihyun Kwon, and Dae-Shik Kim. Progressive face super-resolution via attention to facial landmark. *arXiv preprint arXiv:1908.08239*, 2019.



- [Li *et al.*, 2020] Xiaoming Li, Chaofeng Chen, Shangchen Zhou, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. Blind face restoration via deep multi-scale component dictionaries. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, pages 399–415. Springer, 2020.
- [Liang *et al.*, 2021] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1833–1844, 2021.
- [Menon *et al.*, 2020] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2437–2445, 2020.
- [Miller *et al.*, 2019] Kevin J Miller, Bradley Preece, Todd W Du Bosq, and Kevin R Leonard. A data-constrained algorithm for the emulation of long-range turbulence-degraded video. In *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXX*, volume 11001, pages 204–214. SPIE, 2019.
- [Richardson *et al.*, 2021] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2287–2296, 2021.
- [Rogers and Williams, 2000] L Chris G Rogers and David Williams. *Diffusions, Markov processes and martingales: Volume 2, Itô calculus*, volume 2. Cambridge university press, 2000.
- [Saharia *et al.*, 2022] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [Shen *et al.*, 2018] Ziyi Shen, Wei-Sheng Lai, Tingfa Xu, Jan Kautz, and Ming-Hsuan Yang. Deep semantic face deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8260–8269, 2018.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Song and Ermon, 2019] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [Song *et al.*, 2020] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [Song *et al.*, 2023] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.
- [Tashiro *et al.*, 2021] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csd: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems*, 34:24804–24816, 2021.
- [Vincent, 2011] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [Wang *et al.*, 2021] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9168–9178, 2021.
- [Wang *et al.*, 2022] Zhouxia Wang, Jiawei Zhang, Runjian Chen, Wenping Wang, and Ping Luo. Restoreformer: High-quality blind face restoration from undegraded key-value pairs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17512–17521, 2022.
- [Whang *et al.*, 2022] Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alexandros G Dimakis, and Peyman Milanfar. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16293–16303, 2022.
- [Yang *et al.*, 2016] Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5525–5533, 2016.
- [Yang *et al.*, 2021] Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. Gan prior embedded network for blind face restoration in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 672–681, 2021.
- [Yue and Loy, 2022] Zongsheng Yue and Chen Change Loy. Difface: Blind face restoration with diffused error contraction. *arXiv preprint arXiv:2212.06512*, 2022.
- [Zhang *et al.*, 2018] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [Zhou *et al.*, 2022] Shangchen Zhou, Kelvin CK Chan, Chongyi Li, and Chen Change Loy. Towards robust blind face restoration with codebook lookup transformer. *arXiv preprint arXiv:2206.11253*, 2022.
- [Zhu *et al.*, 2016] Shizhan Zhu, Sifei Liu, Chen Change Loy, and Xiaoou Tang. Deep cascaded bi-network for face hallucination. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 614–630. Springer, 2016.