


VQFR: Blind Face Restoration with Vector-Quantized Dictionary and Parallel Decoder

Yuchao Gu^{*1,2}, Xintao Wang^{†2}, Liangbin Xie^{2,5}, Chao Dong^{4,5},
Gen Li³, Ying Shan², and Ming-Ming Cheng ¹ [0000–0001–5550–8758]

¹TMCC, CS, Nankai University ²ARC Lab, Tencent PCG

³Platform Technologies, Tencent Online Video ⁴Shanghai AI Laboratory

⁵Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences


<https://github.com/TencentARC/VQFR/>

Abstract. Although generative facial prior and geometric prior have recently demonstrated high-quality results for blind face restoration, producing fine-grained facial details faithful to inputs remains a challenging problem. Motivated by the classical dictionary-based methods and the recent vector quantization (VQ) technique, we propose a VQ-based face restoration method – VQFR. VQFR takes advantage of high-quality low-level feature banks extracted from high-quality faces and can thus help recover realistic facial details. However, the simple application of the VQ codebook cannot achieve good results with faithful details and identity preservation. Therefore, we further introduce two special network designs. 1). We first investigate the compression patch size in the VQ codebook and find that the VQ codebook designed with a proper compression patch size is crucial to balance the quality and fidelity. 2). To further fuse low-level features from inputs while not “contaminating” the realistic details generated from the VQ codebook, we proposed a parallel decoder consisting of a texture decoder and a main decoder. Those two decoders then interact with a texture warping module with deformable convolution. Equipped with the VQ codebook as a facial detail dictionary and the parallel decoder design, the proposed VQFR can largely enhance the restored quality of facial details while keeping the fidelity to previous methods.

Keywords: Blind Face Restoration, Vector Quantization, Parallel Decoder

1 Introduction

Blind face restoration aims at recovering low-quality (LQ) faces with unknown degradations, such as noise [40], blur [21,29], down-sampling [9,25], *etc.* This task becomes more challenging in real-world scenarios, where there are more complicated degradations, diverse face poses and expressions. Previous works typically exploit face-specific priors, including geometric priors [6,39,5], generative priors [33,37] and reference priors [23,24,8]. Specifically, geometric priors usually consist of facial landmarks [6], face parsing maps [6,5] and facial component heatmaps [39]. They could provide global guidance for restoring accurate face shapes, but do not help generate realistic details.

[†] Project lead.  Corresponding author.

^{*} Yuchao Gu is an intern in ARC Lab, Tencent PCG.




Fig. 1: Comparisons of restoration quality between GFP-GAN [33] and VQFR. Our VQFR can restore high-quality facial details on various facial regions and keep the fidelity as well, while GFP-GAN lacks realistic fine details. (**Zoom in for best view**)

Besides, geometric priors are estimated from degraded images and thus become inaccurate for inputs with severe degradations. These properties motivate researchers to find better priors.

Recent works [33,37] begin to investigate generative priors in face restoration and achieve superior performance. They usually leverage the powerful generation ability of a pre-trained face generative adversarial network (*e.g.*, StyleGAN [18,19]) to generate realistic textures. These methods typically project the degraded images back into the GAN latent space, and then decode high-quality (HQ) faces with the pre-trained generator. Although GAN-prior-based methods achieve decent overall restoration quality at first glance, they still fail to produce fine-grained facial details, especially the fine hair and delicate facial components (see examples in Fig. 1). This can be partially attributed to the imperfect latent space of the well-trained GAN model. Reference-based methods explore the high-quality guided faces [8,24] or facial component dictionary [23] to solve face restoration problems. DFDNet [23] is a representative method, which does not need to access the faces of the same identity. It explicitly establishes a high-quality “texture bank” for several facial components and then replaces degraded facial components with the nearest HQ facial components in the dictionary. Such a discrete replacement operation directly bridges the gap between the low-quality facial components and high-quality ones, thus having the potential to provide decent facial details. However, the facial component dictionary in DFDNet still has two weaknesses. 1) It offline generates the facial component dictionary with a pre-trained VGGFace [3] network, which is optimized for the recognition task but is sub-optimal for restoration. 2) It only focuses on several facial components (*i.e.*, eyes, nose, and mouth), but does not include other important areas, such as hair and skin.

The limitations of the facial component dictionary motivate us to explore Vector Quantized (VQ) codebook, a dictionary constructed for all facial areas. The proposed face restoration method – VQFR, takes advantage of both dictionary-based methods and GAN training, yet does not require any geometric or GAN prior. Compared to the facial component dictionary [23], the VQ codebook could provide a more comprehensive low-level feature bank that is not restricted to limited facial components. It is also learned in an end-to-end manner by the face reconstruction task. Besides, the mechanism of vector quantization makes it more robust for diverse degradations. Nevertheless, it is not easy to achieve good results simply by applying the VQ codebook. We further introduce two special network designs, which allow VQFR to surpass previous methods in both detail generation and identity preserving.

First, to generate realistic details, we find that it is crucial to select a proper compression patch size f , which indicates “how large a patch is represented” by an atom of the codebook. As shown in Fig. 2, a larger f could lead to better visual quality but worse fidelity. After a comprehensive investigation, we suggest using $f = 32$ for the input image size 512×512 . However, such a selection is only a trade-off between quality and fidelity. The expression and identity could also change a lot even with a proper compression patch size (see Fig. 2). A straightforward solution is to fuse low-level features from input into different decoder layers, just like in GFP-GAN [33]. Although the input features could bring more fidelity information, they will also “contaminate” the realistic details generated from the VQ codebook. This problem leads to our second network design – a parallel decoder. Specifically, the parallel decoder structure consists of a texture decoder and a main decoder. The texture decoder only receives information from the latent representations from the VQ codebook, while the main decoder warps the features from the texture decoder to match the information from degraded input. In order to eliminate the loss of high-quality details and better match the degraded faces, we further adopt a texture warping module with deformable convolution [43] in the main decoder. Equipped with the VQ codebook as a facial dictionary and the parallel decoder design, we can achieve more high-quality facial details while reserving the fidelity for face restoration.

Our contributions are summarized as follows:

1. We propose the VQ dictionary of HQ facial details for face restoration. Our analysis of the VQ codebook shows the potential and limitations of the VQ codebook, together with the importance of the compression patch sizes in face restoration.
2. A parallel decoder is proposed to gradually fuse input features and texture features from VQ codebooks, which keeps the fidelity without sacrificing HQ facial details.
3. Extensive experiments with quantitative and qualitative comparisons show VQFR largely surpasses previous works in restoration quality while keeping high fidelity.

2 Related Work

Blind Face Restoration

Early works explore different facial priors in face restoration. Those priors can be categorized into three types: geometric priors [6,39,5], generative priors [33,37] and reference priors [23,24,8]. The geometric priors include facial landmark [6], face parsing

maps [6,5] and facial component heatmaps [39]. Those priors are estimated from degraded images and thus become inaccurate for inputs with severe degradations. Besides, the geometric structures cannot provide sufficient information to recover facial details. In this work, we do not explicitly integrate geometric priors, but we use landmark distance to estimate restoration fidelity.

Recent works investigate the generative priors to provide facial details and achieve decent performance. In the early arts, GAN inversion methods [26,14] aim to find the closest latent vector in the GAN space given an input image. Within this category, PULSE [26] iteratively optimizes the latent code of a pre-trained StyleGAN [18]. mGANprior [14] simultaneously optimize several codes to promote its reconstruction. Recent works GFP-GAN [33] and GPEN [37] extract fidelity information from inputs and then leverage the pre-trained GAN as a decoder, which achieves a good balance between visual quality and fidelity. Those methods still fail to produce fine-grained facial details. We conjecture that StyleGAN constructs a continuous latent space, and thus GAN-prior methods easily project degraded faces into a suboptimal latent code.

Reference priors [23,24,8] typically rely on reference images of the same identity. To address this issue, DFDNet [23] constructs an offline facial component dictionary [23] with VGGFace [3] for face recognition. It then conducts discrete replacement, *i.e.*, replacing the degraded facial components with high-quality ones in the dictionary by the nearest search. With the facial component dictionary, DFDNet [23] restores better facial components. However, it still has two limitations: 1) The dictionary is offline generated with a recognition model, which is sub-optimal for face restoration. 2) It only builds on several facial components (*i.e.*, eyes, nose and mouth), leaving other facial regions like skin and hair untouched. In this work, we explore the VQ codebook as a facial dictionary, which can be end-to-end trained by a reconstruction objective and provide realistic facial details. Recent RestoreFormer [34] also exploits the VQ codebook, but their work mainly discusses the diverse cross-attention mechanism for LQ latent and HQ code interaction. It does not explore deeply in VQ codebook while we show the dilemma between fidelity and realness when using the VQ codebook of different scales. Moreover, restoreformer does not exploit the input feature of larger resolution and thus results in limited restoration fidelity. Instead, we explore the parallel decoder to achieve realness and fidelity simultaneously.

Vector-Quantized Codebook

The vector-quantized codebook is first introduced in VQ-VAE [31]. With this codebook, the encoder network outputs are discrete rather than continuous, and the prior encapsulated in the codebook is learned rather than static. The following works propose different improvements to codebook learning. VQVAE2 [28] introduces a multi-scale codebook for better image generation. VQGAN [11] trains the codebook with the adversarial objective and thus the codebook can achieve high perceptual quality. To improve the codebook usage, some works [22,38] explore training techniques like L2-normalization or periodically re-initialization. Such a VQ codebook is a patch tokenizer and can be adopted in multiple tasks, like image generation [4,38,11], multi-modal generation [35] and large-scale pretraining [10,1]. Different from previous works that use the VQ codebook to get token features, we explore the potential of the VQ codebook as an HQ facial details dictionary.

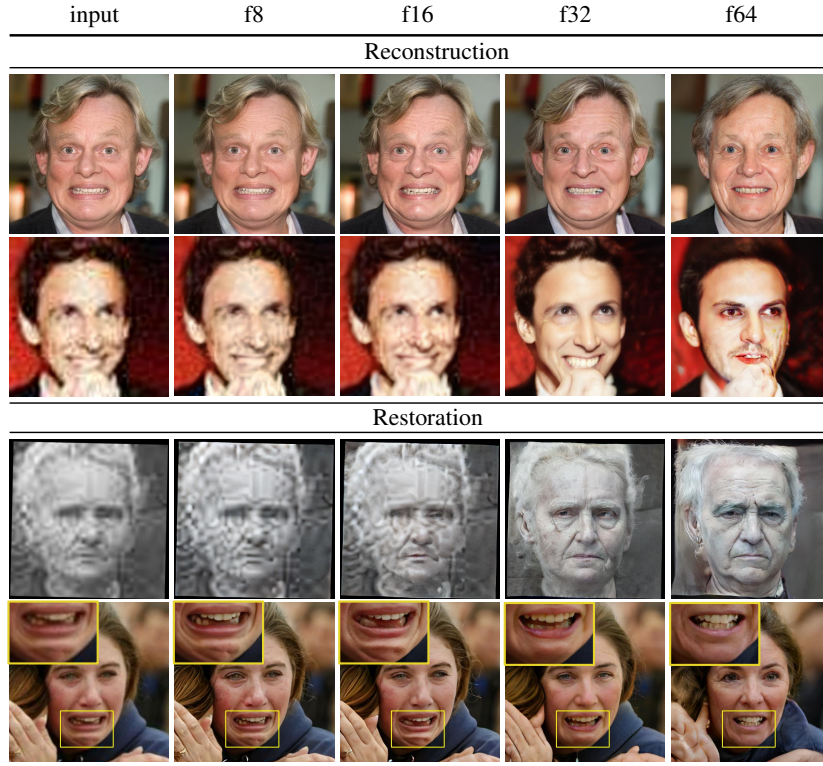


Fig. 2: Reconstruction and restoration results based on codebook with compression patch size $f = \{8, 16, 32, 64\}$. In the reconstruction experiment, we analyze the HQ reconstruction (row 1) and LQ reconstruction (row 2) based on pretrained HQ codebook. In the restoration experiment, we visualized the restoration result of faces of large degradation (row 3) and small degradation (row 4), respectively.

3 Methodology

We describe the proposed VQFR framework in this section. Our goal is to restore high-quality faces with realistic facial details while reserving the fidelity of degraded faces. To achieve this goal, VQFR explores two key ingredients: Vector-Quantized (VQ) dictionary and parallel decoder. The overview of VQFR framework is illustrated in Fig. 3. VQFR contains an encoder, a parallel decoder and a pretrained HQ codebook. We first learn a VQ codebook from only HQ faces with an encoder-decoder structure by the vector-quantization technique[31,11]. Then, a degraded face is encoded to a compact latent representation with a downsampling factor f . At each spatial location of this latent representation, we replace the latent vector with its nearest code in the HQ codebook. After that, the substituted latent representation is then decoded back to the image space, *i.e.*, restored image with high-quality facial details.

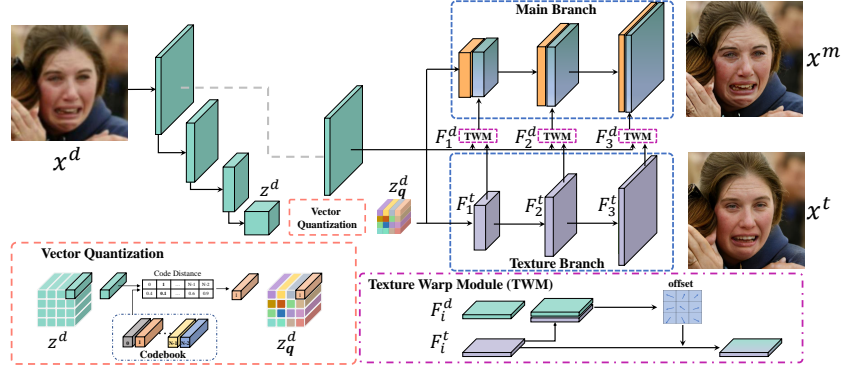


Fig. 3: **Overview of VQFR framework.** It consists of an encoder to map degraded face into latent and a parallel decoder to exploit the HQ code and input feature. The encoder and decoder are bridged by vector quantization model and a pretrained HQ codebook to replace the encoded latent to HQ code.

The remaining parts are arranged as follows: In Sec. 3.1, we first analyze the potential and limitation of vector-quantization technique as a texture dictionary in face restoration task. Then we introduce the parallel decoder to promote the fidelity while maintaining the rich high-quality facial details in Sec. 3.2. In Sec. 3.3, we describe the overall training objective of VQFR framework.

3.1 Vector-Quantized Codebook

Preliminary. The Vector-Quantized (VQ) codebook is first introduced in VQVAE [31], which aims to learn discrete priors to encode images. The following work VQGAN [11] proposes a perceptual codebook by further using perceptual loss [17] and adversarial training objectives [16]. We briefly describe the VQGAN model with its codebook in this section, and more details can be found in [11]. VQGAN is comprised of an encoder E , a decoder G and a codebook $\mathcal{Z} = \{z_k\}_{k=1}^K$ with K discrete codes. For an input image $x \in \mathbb{R}^{H \times W \times 3}$, the encoder E maps the image x to its spatial latent representation $\hat{z} = E(x) \in \mathbb{R}^{h \times w \times n_z}$, where n_z is the dimension of latent vectors. The vector-quantized representation z_q is then obtained by applying element-wise quantization $\mathbf{q}(\cdot)$ of each spatial code $\hat{z}_{ij} \in \mathbb{R}^{n_z}$ onto its closest codebook entry z_k :

$$z_q = \mathbf{q}(\hat{z}) := \left(\arg \min_{z_k \in \mathcal{Z}} \|\hat{z}_{ij} - z_k\| \right) \in \mathbb{R}^{h \times w \times n_z}. \quad (1)$$

The decoder G maps the quantized representation z_q back to the image space, and the overall reconstruction $\hat{x} \approx x$ can be formulated as:

$$\hat{x} = G(z_q) = G(\mathbf{q}(E(x))). \quad (2)$$

The encoder E maps images of size $H \times W$ into discrete codes of size $H/f \times W/f$, where f denotes the downsampling factor. It can be regarded as compressing each $f \times f$

patch in the image x into one code. In other words, for each code in z_q , f also denotes the corresponding spatial size in the original image x . We name this downsampling factor f as the *compression patch size* in our paper.

Since the quantization operation in Equ. 2 is discrete and non-differentiable, VQGAN adopts straight-through gradient estimator [2], which simply copies the gradients from the decoder to the encoder. Thus, the model and codebook can be trained end-to-end via the loss function \mathcal{L}_{vq} . VQGAN also employs perceptual loss and adversarial loss to encourage reconstructions with better perceptual quality. The full training objective of VQGAN and its codebook is:

$$\mathcal{L}(E, G, \mathcal{Z}) = \underbrace{\|x - \hat{x}\|_1 + \|\text{sg}[E(x)] - z_q\|_2^2 + \beta \|\text{sg}[z_q] - E(x)\|_2^2}_{\mathcal{L}_{vq}} + \mathcal{L}_{per} + \mathcal{L}_{adv}, \quad (3)$$

where $\|x - \hat{x}\|_1$ is the reconstruction loss and $\text{sg}[\cdot]$ denotes stop gradient operation. The codebook is updated by $\|\text{sg}[E(x)] - z_q\|_2^2$. $\beta \|\text{sg}[z_q] - E(x)\|_2^2$ is the commitment loss [31] to reduce the discrepancy between the encoded latent vectors and codes, where β is the commitment weight and set to 0.25 in all experiments.

Analysis In order to better understand the potential and limitations of VQ codebooks for face restoration, we conduct several preliminary experiments and draw the following observations.

Observation 1: *Degradations in LQ faces can be removed by VQ codebooks trained only with HQ faces, when we adopt a proper compression patch size f .*

Following [11] and Equ. 3, we first train VQGAN with perceptual codebooks on HQ faces using different compression patch sizes $f = \{8, 16, 32, 64\}$. All our experiments are conducted on 512×512 input faces. The illustration of training architectures is shown in Fig. 4. We then examine the reconstruction quality for different compression patch sizes f . Note that the reconstruction output is expected to be the same as the input. As shown in row 1 of Fig. 2, the reconstruction quality of HQ faces is as expected. The reconstruction quality of HQ faces exhibits a reasonable trend: a smaller compression patch size f will lead to more faithful outputs. After that, we are curious whether the codebook trained only on HQ faces can also reconstruct LQ faces. We use LQ faces as inputs and examine the outputs. Interestingly, we can observe that under the compression patch sizes $f = 32$, the degradation in LQ faces can be removed (see row 2 in Fig. 2). This is because the codebook trained only on HQ faces has almost no degradation-related codes. During reconstruction, the vector quantization operation can replace the “degraded” vectors of inputs with the “clean” codes in codebooks.

Such a phenomenon shows the potential of the VQ codebook. However, it only happens with a large compression patch size f , as the codebook with a small compression patch size cannot well distinguish the degradations and detailed textures. In the extreme, for $f = 1$ with each pixel having a quantized code, both the degradation and detailed textures can be well recovered by a combination of codebook entries. On the other hand, a too large compression patch size (e.g., $f = 64$) will result in a significant change in identity, even it could also reconstruct “clean” face images.

Observation 2: *When training for the restoration task, there is also a trade-off between improved detailed textures and fidelity changes.*

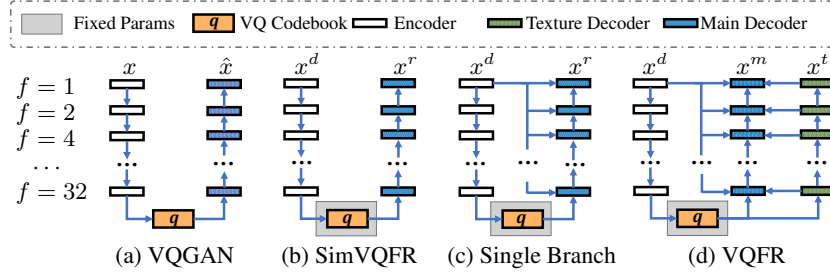


Fig. 4: Illustration of the architecture variants. (a) The VQGAN structure is used in codebook training. (b) The SimVQFR structure. (c) Single branch decoder. (d) The proposed parallel decoder.

We then investigate the VQ codebooks in face restoration, *i.e.*, training with LQ-HQ pairs as most face restoration works do [33]. Based on the trained model for reconstruction, we then fix the VQ codebooks, and finetune the encoder and decoder with LQ-HQ pairs (the training details are the same as that in Sec. 4.1). We denote this simple VQ model for face restoration as SimVQFR. We can observe from Fig. 2 that there is still a trade-off between improved detailed textures and fidelity changes. The key influential factor is the compression patch size f . With a small f (*i.e.*, $\{8, 16\}$), the SimVQFR model fails to remove degradations and cannot recover sufficient detailed textures. While with a large f (*i.e.*, $\{32, 64\}$), the textures are largely improved but the fidelity (*i.e.*, expression and identity) also changes a lot by the codebook.

Our choice. Based on the above analysis, we can conclude that the VQ codebook, as a texture dictionary, has its value in generating high-quality facial textures and removing input degradations. However, there is a trade-off between the improved detailed textures and fidelity changes. The key influential factor is the compression patch size f . In order to better leverage the strength of generating high-quality facial textures, we choose the compression patch size $f = 32$ for 512×512 input faces. The left problem is how to preserve the fidelity with the VQ codebook of $f = 32$. We will present our parallel decoder solution to address this problem in Sec. 3.2.

3.2 Parallel Decoder

The VQ codebook with the compression patch size of $f = 32$ can be used as a texture bank to provide realistic facial details, but it also brings the problem of fidelity changes. From the results in Fig. 2, we can find that the position of facial components and the facial lines are changed, making the expression and identity largely deviate from the inputs. A straightforward solution is to integrate the feature information from degraded inputs to help improve the fidelity. However, simply fusing input features into the decoder with a single branch (as shown in Fig. 4) will lead to inferior details (see Fig. 8). In other words, such a single branch fusion strategy tends to corrupt the generated high-quality details. Though input features can bring more fidelity information, these features also contain input degradations. During feature propagation, *i.e.*, the upsampling process from small spatial size to large spatial size, the intermediate features will largely

be influenced by input features, and gradually contain fewer high-quality facial details from the VQ codebook.

To overcome the dilemma of keeping the realistic facial details and promoting fidelity, we propose a parallel decoder structure with a texture warping module. The core idea of the parallel decoder (Fig. 4) is to decouple the two goals of face restoration, *i.e.*, generating high-quality facial details and keeping the fidelity. As shown in Fig. 3, given a degraded face $x^d \in \mathbb{R}^{H \times W \times 3}$, we first encode it to the latent vector z^d by $z^d = E(x^d)$. The encoder consists of several residual blocks and downsampling operations. Then we replace the code z^d with the HQ codebooks by Equ. 1 to get the quantized code z_q^d . Since the quantized code z_q^d is from the HQ codebook, it contains HQ facial details without degradations. In order to keep its realistic textures, we use a texture decoder G_t to decode it back to image space $x^t = G_t(z_q^d)$. We denote the multi-level feature of the texture branch as $F^t = \{F_i^t\}$. Since the texture branch only decodes from the HQ code, F^t can keep realistic facial details.

The main branch decoder G_m aims to generate faces x^m with high fidelity while having the realistic facial details from the texture decoder G_t . As illustrated in Fig. 3, the main branch decoder G_m warps the texture feature F^t based on the input feature extracted from degraded inputs at multiple spatial levels. We use the input feature with the largest spatial resolution as it retains the richest fidelity information of degraded faces. We then directly downsample the input feature to different resolution levels to get the multi-level features of degraded faces $F^d = \{F_i^d\}$.

For the i -th resolution level, we first warp F_i^t with high-quality facial details towards F_i^d by a texture warping module, which will be described later. After that, we fuse the warped feature F_i^w and the upsampled feature from F_{i-1} , to get the F_i feature in the main decoder. The process can be formulated as:

$$F_i^w = TWM(F_i^d, F_i^t), \quad F_i = \text{Conv}(\text{Concat}(\text{Upsample}(F_{i-1}), F_i^w)), \quad (4)$$

where TWM is the texture warping module. Our parallel decoder shares the same spir- its as reference-based restoration [24]. In our case, the features from the texture decoder serve as the reference features containing high-quality details. Unlike reference-based restoration, our method does not require extra high-quality images with rich textures. Instead, the parallel decoder learns the main feature and “reference” feature jointly.

Texture Warping Module (TWM) The texture warping module aims to warp realistic facial details to match the fidelity of degraded inputs, especially the position of facial components and expressions. There are two inputs of TWM, one is the input feature F^d and the other is the texture feature F^t . As analyzed above, the texture features are decoded from the HQ codebook and have high-quality facial details. But their fidelity probably deviates from inputs. Therefore, we adopt a deformable convolution [43] to better warp the realistic facial details F^t towards the input feature F^d . Specifically, we first concatenate those two features to generate offsets. Then the offset is used in the deformable convolution to warp the texture features to match the fidelity of input, which can be formulated as:

$$\text{offset} = \text{Conv}(\text{Concat}(F_i^d, F_i^t)), \quad F_i^w = \text{Dconv}(F_i^t, \text{offset}), \quad (5)$$

where $Dconv$ denotes the deformable convolution. We also adopt a separable convolution with a large kernel size to model large position offsets between texture features and input features.

3.3 Model Objective

The training objective of VQFR consists of 1) pixel reconstruction loss that constraints the restored outputs close to the corresponding HQ faces; 2) code alignment loss that forces the codes of LQ inputs to match the codes of the corresponding HQ inputs; 3) perceptual loss to improve the perceptual quality in feature space; and 4) adversarial loss for restoring realistic textures. We denote the degraded face as x^d , decoder restored results as $x^r = \{x^t, x^m\}$ and the ground truth HQ image as x^h . The loss definitions are as follows.

Pixel Reconstruction Loss. We use the widely-used L1 loss in the pixel space as the reconstruction loss, which is denoted as: $\mathcal{L}_{pix} = \|x^r - x^h\|_1$. Empirically, we find that with the input feature of degraded faces, the pixel reconstruction loss has a negative impact on facial details. Therefore, we discard pixel reconstruction loss in the main decoder.

Code Alignment Loss. The code alignment loss aims to improve the performance of matching the codes of LQ images with codes of HQ images. We adopt the L2 loss to measure the distance, which can be formulated as: $\mathcal{L}_{code} = \|z^d - z_q^h\|_2^2$, where z_q^h is the ground truth code obtained from encoding HQ face to the quantized code by pretrained VQGAN.

Perceptual Loss. We use the widely used perceptual loss [17,41] for both two decoder outputs: $\mathcal{L}_{per} = \|\phi(x^r) - \phi(x^h)\|_1 + \lambda_{style} \|Gram(\phi(x^r)) - Gram(\phi(x^h))\|_1$, where ϕ is the pretrained VGG-16 [30] network and $Gram$ means the Gram matrix [13]. The former term measures the content distance and the latter term measures the style difference between the restoration results and the corresponding HQ images.

Adversarial Loss. We employ the global discriminator D_g in SWAGAN [12] to encourage VQFR to favor the solutions in the natural image manifold and generate realistic textures. In order to increase the local quality of facial details, we further adopt the local discriminator D_l of PatchGAN [16] in our training. The objectives of the global and local discriminators are defined as

$$\mathcal{L}_{adv}^{global} = -\mathbb{E}_{x^r}[\text{softplus}(D_g(x^r))], \mathcal{L}_{adv}^{local} = \mathbb{E}_{x^r}[\log D_l(x^h) + \log(1 - D_l(x^r))]. \quad (6)$$

Total objective. The total training objective is the combination of above losses:

$$\mathcal{L}_{total} = \lambda_{pix}\mathcal{L}_{pix} + \lambda_{code}\mathcal{L}_{code} + \lambda_{per}\mathcal{L}_{per} + \lambda_{global}\mathcal{L}_{adv}^{global} + \lambda_{local}\mathcal{L}_{adv}^{local}, \quad (7)$$

where the λ_{pix} , λ_{code} , λ_{per} , λ_{global} and λ_{local} are scale factors of corresponding loss.

4 Experiments

4.1 Implementation and Evaluation Settings

Implementation. We implement the VQGAN and VQFR with six resolution levels, *i.e.*, $\{1, 2, 4, 8, 16, 32\}$. For the VQ codebook, we use 1024 codebook entries with 256

Table 1: Quantitative comparison on the **CelebA-Test** dataset for blind face restoration. **Red** and **blue** indicates the best and second best performance.

Methods	LPIPS↓	FID↓	NIQE↓	Deg.↓	LMD.↓	PSNR↑	SSIM↑
Input	0.4866	143.98	13.440	47.94	3.76	25.35	0.6848
Wan <i>et al.</i> [32]	0.4826	67.58	5.356	43.00	2.92	24.71	0.6320
HiFaceGAN [36]	0.4770	66.09	4.916	42.18	3.16	24.92	0.6195
DFDNet [23]	0.4341	59.08	4.341	40.31	3.31	23.68	0.6622
PSFRGAN [5]	0.4240	47.59	5.123	39.69	3.41	24.71	0.6557
mGANprior [14]	0.4584	82.27	6.422	55.45	411.72	24.30	0.6758
PULSE [26]	0.4851	67.56	5.305	69.55	7.35	21.61	0.6200
GFP-GAN [33]	0.3646	42.62	4.077	34.60	2.41	25.08	0.6777
VQFR (ours)	0.3515	41.28	3.693	35.75	2.43	24.14	0.6360




Fig. 5: Qualitative comparisons on the **CelebA-Test**. VQFR is able to restore high-quality facial details in various facial components, *e.g.*, eyes and mouth.

channels in all experiments. In order to increase the codebook usage, we follow [22] and periodically re-initialize the codebook with k-means clustering. More implementation details are provided in the supplementary.

In the first stage of codebook training, we adopt ADAM [20] optimizer with a learning rate of $1e - 4$. The training iteration is set to 800K with a batch size of 16. For the second stage training for restoration, the loss weights are set to $\lambda_{pix} = \lambda_{code} = \lambda_{per} = 1$, $\lambda_{style} = 2000$ and the adversarial loss weights are set to $\lambda_{local} = \lambda_{global} = 0.5$. We train VQFR for 200K iterations by ADAM optimizer with the learning rate of $1e - 5$.

Training Datasets. The VQFR is trained on the FFHQ dataset [18] including 70,000 high-quality faces. The images are resized to 512^2 during training. Following the common practice in [23, 24, 33], we use the following degradation model to synthesize training pairs: $x = [(y \otimes k_\sigma) \downarrow_r + n_\delta]_{JPEG_q}$, where σ , r , δ and q are randomly sampled from $\{0.2 : 10\}$, $\{1 : 8\}$, $\{0 : 15\}$ and $\{60 : 100\}$, respectively.

Testing Datasets. Following the practice in GFP-GAN [33], we conduct experiments on the synthetic dataset *CelebA-Test* and three real-world datasets - *LFW-Test*, *CelebChild-Test* and *WebPhoto-Test*. These datasets have diverse and complicated degradations. All these datasets have no overlap with the training dataset.

Evaluation Metrics. Our evaluation metrics contain two widely-used non-reference perceptual metrics: FID [15] and NIQE [27]. We also measure the pixel-wise metrics (PSNR and SSIM) and perceptual metric (LPIPS [42]) for benchmarking CelebA-Test

Table 2: Quantitative comparison on the *real-world* **LFW**, **CelebChild**, **WebPhoto**. **Red** and **blue** indicates the best and second best performance.

Dataset Methods	LFW-Test		CelebChild		WebPhoto	
	FID↓	NIQE ↓	FID↓	NIQE ↓	FID↓	NIQE ↓
Input	137.56	11.214	144.42	9.170	170.11	12.755
Wan <i>et al.</i> [32]	73.19	5.034	115.70	4.849	100.40	5.705
HiFaceGAN [36]	64.50	4.510	113.00	4.855	116.12	4.885
DFDNet [23]	62.57	4.026	111.55	4.414	100.68	5.293
PSFRGAN [5]	51.89	5.096	107.40	4.804	88.45	5.582
mGANprior [14]	73.00	6.051	126.54	6.841	120.75	7.226
PULSE [26]	64.86	5.097	102.74	5.225	86.45	5.146
GFP-GAN [33]	49.96	3.882	111.78	4.349	87.35	4.144
VQFR (ours)	50.64	3.589	105.18	3.936	75.38	3.607

with Ground-Truth (GT). We follow previous work [33] to use the embedding angle of ArcFace [7] as the identity metric, which is denoted by ‘Deg.’. In order to better measure the fidelity with accurate facial positions and expressions, we further adopt landmark distance (LMD) as the fidelity metric. More details are provided *in the supplementary*.

4.2 Comparisons with State-of-the-art Methods

We compare our VQFR with several state-of-the-art face restoration methods: Wan *et al.* [32], HiFaceGAN [36], DFDNet [23], PSFRGAN [5], mGANprior [14], PULSE [26] and GFP-GAN [33].

Synthetic CelebA-Test. From the quantitative results in Table. 1, VQFR achieves the lowest LPIPS, implying that VQFR generates the restored faces with the closest perceptual quality to ground-truth. VQFR also achieves the best FID and NIQE, with a large improvement over GFP-GAN, demonstrating that results of VQFR are closer to real faces and have more realistic details. For fidelity, VQFR can achieve comparable landmark distance and identity degree to GFP-GAN, showing that it can recover accurate facial expressions and detail positions.

Qualitative results are presented in Fig. 5. Thanks to the VQ codebook design, VQFR generates high-quality facial components like eyes and mouth as well as other facial regions. VQFR also maintains the fidelity with the help of the parallel decoder.

Real-World LFW, CelebChild, and WebPhoto-Test. We evaluate VQFR on three real-world test datasets to test the generalization ability. Table. 2 shows the quantitative results. It is observed that VQFR largely improves the realness and perceptual quality of all three real-world datasets. PULSE [26] achieves a higher perceptual quality on CelebChild, but its fidelity is severely affected. From the qualitative results in Fig. 6, the face restored by VQFR is of the most high quality on different facial regions.

4.3 Ablation Study

Importance of input features from degraded faces. The SimVQFR model directly uses the VQ codebook without exploiting input features. As shown in Fig. 7(a), though




Fig. 6: Qualitative comparisons on three real-world datasets. **Zoom in for best view.**


it can achieve high perceptual quality (low FID and NIQE), its landmark distance is large, indicating the low fidelity. After incorporating the input features of degraded faces, the fidelity improves significantly. From the visualization of Fig. 8(a), we can observe that the facial lines of SimVQFR are deviated from LQ faces, resulting in an expression change. Such a phenomenon can also be observed in the examples in Fig. 3, where the texture decoder output changes the woman’s expression apparently.

Importance of parallel decoder. When comparing Variant 1 and Variant 2, the NIQE (which favors high-quality details) of Variants 1 is higher and clearly inferior to Variant 2 (Fig. 7 (a)). From the visualization in Fig. 8(b), we observe that, in Variant 1 without the parallel decoder, the eyes and hair lose high-frequency details and are biased to degraded ones. While with the parallel decoder, the details are clearer and more realistic.

Texture Warping Module. As shown in Fig. 7 (a), Variant 2 does not have a TWM module and directly utilizes concatenation fusion of texture features and input features of degraded features. It lacks the ability to dynamically adjust the features with warping. Therefore, Variant 2 without TWM keeps the high-quality textures but changes the fidelity, as it cannot well adjust the fine details and expression. From the visualization

Models	Configurations			CelebA-Test		
	Inp Feat.	Para Dec.	TWM	FID↓	NIQE↓	LMD↓
SimVQFR				40.51	3.844	2.96
Variant 1	✓			39.64	4.019	2.36
Variant 2	✓	✓		42.61	3.714	2.48
VQFR	✓	✓	✓	41.28	3.693	2.43

(a) Ablation results on CelebA-Test. *Inp Feat.*: input feature of degraded faces; *Para Dec.*: parallel decoder; *TWM*: texture warping module.




(b) Visualization of the balance between realness of fidelity of different configurations.


Fig. 7: Quantitative ablation studies of key designs in VQFR.



(a) Ablation of input feature for fidelity.



(b) Ablation of parallel decoder for realness.



(c) Ablation of TWM for fidelity and realness.

Fig. 8: Qualitative ablation studies of key designs on VQFR.

of Fig. 8(c), the TWM module helps accurately warp the high-quality facial details to better match the degraded input.

5 Conclusion

In this paper, we propose the vector-quantized (VQ) dictionary of high-quality facial details for face restoration. Our analysis of the VQ codebook shows the potential and limitations of the VQ codebook. In order to keep the fidelity without the loss of high-quality facial details, a parallel decoder is further proposed to gradually fuse input features and texture features from VQ codebooks. Equipped with the VQ codebook as a dictionary and the parallel decoder, our proposed vector-quantized face restoration (VQFR) can produce high-quality facial details while preserving fidelity. Extensive experiments show that our methods surpass previous works on both synthetic and real-world datasets.

Acknowledgement

This work is funded by the National Key Research and Development Program of China Grant No.2018AAA0100400 and NSFC (NO. 62176130). And this work is partially supported by National Natural Science Foundation of China (61906184, U1913210), and the Shanghai Committee of Science and Technology, China (Grant No. 21DZ1100100).

References

1. Bao, H., Dong, L., Wei, F.: Beit: Bert pre-training of image transformers. arXiv preprint arXiv:2106.08254 (2021) [4](#)
2. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432 (2013) [7](#)
3. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: 2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018). pp. 67–74. IEEE (2018) [2](#), [4](#)
4. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. arXiv preprint arXiv:2202.04200 (2022) [4](#)
5. Chen, C., Li, X., Yang, L., Lin, X., Zhang, L., Wong, K.Y.K.: Progressive semantic-aware style transformation for blind face restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11896–11905 (2021) [1](#), [3](#), [4](#), [11](#), [12](#)
6. Chen, Y., Tai, Y., Liu, X., Shen, C., Yang, J.: Fsrnet: End-to-end learning face super-resolution with facial priors. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2492–2501 (2018) [1](#), [3](#), [4](#)
7. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4690–4699 (2019) [12](#)
8. Dogan, B., Gu, S., Timofte, R.: Exemplar guided face image super-resolution without facial landmarks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019) [1](#), [2](#), [3](#), [4](#)
9. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: European conference on computer vision. pp. 184–199. Springer (2014) [1](#)
10. Dong, X., Bao, J., Zhang, T., Chen, D., Zhang, W., Yuan, L., Chen, D., Wen, F., Yu, N.: Peco: Perceptual codebook for bert pre-training of vision transformers. arXiv preprint arXiv:2111.12710 (2021) [4](#)
11. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12873–12883 (2021) [4](#), [5](#), [6](#), [7](#)
12. Gal, R., Hochberg, D.C., Bermano, A., Cohen-Or, D.: Swagan: A style-based wavelet-driven generative model. ACM Transactions on Graphics (TOG) **40**(4), 1–11 (2021) [10](#)
13. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2414–2423 (2016) [10](#)
14. Gu, J., Shen, Y., Zhou, B.: Image processing using multi-code gan prior. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3012–3021 (2020) [4](#), [11](#), [12](#)
15. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017) [11](#)
16. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017) [6](#), [10](#)
17. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016) [6](#), [10](#)

18. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4401–4410 (2019) 2, 4, 11
19. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8110–8119 (2020) 2
20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014) 11
21. Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J.: Deblurgan: Blind motion deblurring using conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8183–8192 (2018) 1
22. Łańcucki, A., Chorowski, J., Sanchez, G., Marxer, R., Chen, N., Dolfing, H.J., Khurana, S., Alumäe, T., Laurent, A.: Robust training of vector quantized bottleneck models. In: *2020 International Joint Conference on Neural Networks (IJCNN)*. pp. 1–7. IEEE (2020) 4, 11
23. Li, X., Chen, C., Zhou, S., Lin, X., Zuo, W., Zhang, L.: Blind face restoration via deep multi-scale component dictionaries. In: *European Conference on Computer Vision*. pp. 399–415. Springer (2020) 1, 2, 3, 4, 11, 12
24. Li, X., Liu, M., Ye, Y., Zuo, W., Lin, L., Yang, R.: Learning warped guidance for blind face restoration. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 272–289 (2018) 1, 2, 3, 4, 9, 11
25. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 136–144 (2017) 1
26. Menon, S., Damian, A., Hu, S., Ravi, N., Rudin, C.: 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*. pp. 2437–2445 (2020) 4, 11, 12
27. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. *IEEE Signal processing letters* **20**(3), 209–212 (2012) 11
28. Razavi, A., Van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems* **32** (2019) 4
29. Shen, Z., Lai, W.S., Xu, T., Kautz, J., Yang, M.H.: Deep semantic face deblurring. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8260–8269 (2018) 1
30. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014) 10
31. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017) 4, 5, 6, 7
32. Wan, Z., Zhang, B., Chen, D., Zhang, P., Chen, D., Liao, J., Wen, F.: Bringing old photos back to life. In: *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2747–2757 (2020) 11, 12
33. Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9168–9178 (2021) 1, 2, 3, 4, 8, 11, 12
34. Wang, Z., Zhang, J., Chen, R., Wang, W., Luo, P.: Restoreformer: High-quality blind face restoration from undegraded key-value pairs. *arXiv preprint arXiv:2201.06374* (2022) 4
35. Wu, C., Liang, J., Ji, L., Yang, F., Fang, Y., Jiang, D., Duan, N.: N\” uwa: Visual synthesis pre-training for neural visual world creation. *arXiv preprint arXiv:2111.12417* (2021) 4
36. Yang, L., Wang, S., Ma, S., Gao, W., Liu, C., Wang, P., Ren, P.: Hifacegan: Face renovation via collaborative suppression and replenishment. In: *Proceedings of the 28th ACM International Conference on Multimedia*. pp. 1551–1560 (2020) 11, 12

37. Yang, T., Ren, P., Xie, X., Zhang, L.: Gan prior embedded network for blind face restoration in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 672–681 (2021) [1](#), [2](#), [3](#), [4](#)
38. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldrige, J., Wu, Y.: Vector-quantized image modeling with improved vqgan. arXiv preprint arXiv:2110.04627 (2021) [4](#)
39. Yu, X., Fernando, B., Ghanem, B., Porikli, F., Hartley, R.: Face super-resolution guided by facial component heatmaps. In: Proceedings of the European conference on computer vision (ECCV). pp. 217–233 (2018) [1](#), [3](#), [4](#)
40. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. IEEE transactions on image processing **26**(7), 3142–3155 (2017) [1](#)
41. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) [10](#)
42. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018) [11](#)
43. Zhu, X., Hu, H., Lin, S., Dai, J.: Deformable convnets v2: More deformable, better results. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9308–9316 (2019) [3](#), [9](#)