

Progressive Semantic-Aware Style Transformation for Blind Face Restoration

Chaofeng Chen^{1, 3*} Xiaoming Li^{2, 3, 5} Lingbo Yang^{4, 3} Xianhui Lin³
Lei Zhang^{3, 5} Kwan-Yee K. Wong¹

¹ Department of Computer Science, The University of Hong Kong ² Faculty of Computing, Harbin Institute of Technology

³ DAMO Academy, Alibaba Group ⁴ Institute of Digital Media, Peking University

⁵ Department of Computing, The Hong Kong Polytechnic University

<https://github.com/chaofengc/PSFRGAN>



Figure 1: Result comparison with state-of-the-art methods. (Please zoom in to see the details)

Abstract

Face restoration is important in face image processing, and has been widely studied in recent years. However, previous works often fail to generate plausible high quality (HQ) results for real-world low quality (LQ) face images. In this paper, we propose a new progressive semantic-aware style transformation framework, named PSFR-GAN, for face restoration. Specifically, instead of using an encoder-decoder framework as previous methods, we formulate the restoration of LQ face images as a multi-scale progressive restoration procedure through semantic-aware style transformation. Given a pair of LQ face image and its corresponding parsing map, we first generate a multi-scale pyramid of the inputs, and then progressively modulate different scale features from coarse-to-fine in a semantic-aware style transfer way. Compared with previous networks, the proposed PSFR-GAN makes full use of the semantic (parsing maps) and pixel (LQ images) space information from different scales of input pairs. In addition, we further introduce a semantic aware style loss which calculates the feature style loss for each semantic region individually to improve the details of face textures. Finally, we pretrain a face parsing network which can generate decent parsing maps from

real-world LQ face images. Experiment results show that our model trained with synthetic data can not only produce more realistic high-resolution results for synthetic LQ inputs but also generalize better to natural LQ face images compared with state-of-the-art methods.

1. Introduction

Blind face restoration refers to recovering the HQ images from the LQ inputs which suffer from unknown degradation such as low resolution, noise, blur and lossy compression. It has drawn more and more interest due to its wide applications. However, most current restoration methods still focus on a specific type of restoration, especially super resolution, and few of them can generalize well to real LQ images.

Unlike general image restoration, face restoration can exploit strong prior knowledge of the face to recover details of the face components even when the images are severely degraded. Therefore, many recent works about face super resolution incorporate face prior knowledge to improve the performance, such as parsing maps [4], face landmarks [1, 5, 37] and identity prior [42]. Most of these works are based on an encoder-decoder like structure, which follow the practice of general image restoration and aim to learn a direct black-box mapping from LQ to HQ images. AI-

* This work was done when he was an intern at Alibaba.

though they managed to get better results with extra face prior knowledge as inputs or supervision, few of them reports satisfactory results for real LQ images. Other methods [22, 21] try to utilize high quality references to facilitate the restoration of LQ images. However, their practical applications are limited when there are no high quality references.

In this work, we propose a new progressive framework, named PSFR-GAN, which formulates face restoration as multi-scale semantic-aware style transformation procedure. Inspired by recent success of style based GAN, *i.e.*, StyleGAN [15], we use a semantic-aware style transfer approach to modulate the features of different scales progressively. To be specific, the proposed PSFR-GAN starts with a learned constant latent code and then generates features of different scales through several upsample layers. We modulate the “styles” of different scale features by generating the corresponding style transformation parameters from different scale inputs. The LQ input provides color information and the parsing map provides shape and semantic information. In this way, more details are added to the final features in a coarse-to-fine manner. In addition, we proposed a *semantic aware style loss* which calculates the gram matrix loss for each semantic region separately. Gram matrix loss is usually applied in neural style transfer [7], recent works [8] found it also effective in recovering textures. In this work, we show that the *semantic aware style loss* can help to improve the restoration of textures and alleviate the occurrence of artifacts in different face regions.

Finally, to make our framework more practical, we pre-train a face parsing network (FPN) for LQ face images. Intuitively, predicting face parsing maps is easier than face restoration because we do not need to care the texture details. Experiments demonstrate that the FPN is pretty robust on parsing real-world LQ face images. During test time, we first generate parsing maps for LQ inputs with FPN, and then produce the HQ outputs with PSFR-GAN.

Our contributions are summarized as follows:

1. We propose a novel multi-scale progressive framework for practical blind face restoration, *i.e.* PSFR-GAN. Our model can recover high quality face details progressively through semantic aware style transformation with multi-scale LQ images and parsing maps as inputs. Compared with previous works, the proposed PSFR-GAN can make better use of multi-scale inputs in both pixel domain and semantic domain.
2. We introduce the *semantic aware style loss* which helps to improve the texture restoration of different semantic regions and reduce the occurrence of artifacts.
3. Extensive experiments demonstrate that our model trained with synthetic dataset generalizes better to natural LQ images than current state-of-the-arts.

4. By introducing a pretrained LQ face parsing network, our model can generate HQ images given only LQ inputs, making it highly practical and applicable.

2. Related Works

In this section, we briefly review the existing methods for face super-resolution, blind face restoration and HQ face generation based on generative adversarial networks.

Face Super-Resolution. Face super resolution is often studied as a basic task for face restoration. Zhu *et al.* [45] proposed a cascaded two-branch network to optimize face hallucination and dense correspondence field estimation in a unified framework. Yu *et al.* [39] exploited generative adversarial networks (GAN) to directly super-resolve LR inputs. They further improved their model to handle unaligned faces [40], noisy faces [41] and faces with different attributes [38]. Instead of directly inferring HR face images, Huang *et al.* [12] proposed to predict wavelet coefficients from LR images to reconstruct HR images. More recent works exploited extra face prior information to improve the SR performance. Chen *et al.* [4] used face landmark heatmaps and parsing maps to super-resolve unaligned LR faces. They first predicted landmark heatmaps and parsing maps from LR faces, and then concatenated them with feature maps to fine tune the SR results. Adrian *et al.* [1] jointly learned face SR and landmark prediction. Yu *et al.* [37] employed face component heatmaps to preserve face structure while super-resolving LR faces. They used one CNN branch to predict face component heatmaps and then concatenated the heatmaps to the SR CNN branch. Kim *et al.* [5] proposed a multi-scale facial attention loss by multiplying the heatmap values to the pixel differences of different scales to better restore pixels around landmarks. However, most of these works followed a general encoder-decoder framework and did not fully utilize the face prior, making them unsuitable to handle real-world LQ images.

Blind Face Restoration To deal with blind face restoration, recent works either tried to improve the framework or adopted reference based approach. Adrian *et al.* [2] proposed a two-stage GAN framework to learn real degradation. Yang *et al.* [36] introduced the HiFaceGAN which progressively replenish face details. As for reference based methods, Song *et al.* [31] constructed a high-resolution dictionary to help enhance face components. Li *et al.* [22] proposed GRFNet which learns to warp a guidance image for blind face restoration. They further improve their work by replacing single reference with multiple reference images [21] and feature dictionaries [20]. The aforementioned model based methods fail to make full use of face prior, and the reference based approaches require HQ references which limits their practical application.

High Quality Face Generation with GAN. Among all GAN based image generation methods, the style based

GANs (StyleGAN) proposed by Karras *et al.* [15, 16] exceed other works by a large margin and generate HQ faces which are almost indistinguishable from real photos. Instead of directly generating the faces from random latent vector, they first generate style transform parameters from latent vector, and modulated the features with adaptive instance normalization (AdaIN) [13]. SEAN [44] and GauGAN [27] extend StyleGAN to generate realistic faces from parsing maps. PULSE [24] describes a self-supervised face restoration approach by exploring the latent space of a pre-trained StyleGAN. mGANprior [10] extended PULSE with multiple latent codes. Inspired by the above works, we introduce the style transformation approach to face restoration in this paper. Compared with PULSE, our model does not need a time consuming optimization process and preserves the shape and identity better.

3. Proposed Method

In this section, we first describe our formulation and framework in detail, then introduce the *semantic-aware style loss* and the other objectives used to train our networks.

3.1. Progressive Semantic-Aware Style Transformation

As illustrated in Fig. 2, the architecture of our PSFR-GAN is inspired by GauGAN [27] and StyleGANs [15, 16]. It starts with a learned constant of size $C \times 16 \times 16$, denoted as F_0 , where C is the channel size. Then, F_0 goes through several upsample residual blocks and generates the final features with the same size as HQ images. Let's define the output features of i -th residual block as F_i , then the features are progressively upsampled in the following way

$$F_i = \begin{cases} \Phi_{ST}(\Phi_{RES}(F_{i-1})), & i = 1 \\ \Phi_{ST}(\Phi_{UP}(F_{i-1})), & 1 < i \leq 6 \end{cases} \quad (1)$$

where $\Phi_{RES}(\cdot)$ denotes the residual convolution block, $\Phi_{UP}(\cdot)$ denotes the upsample residual convolution block and $\Phi_{ST}(\cdot)$ denotes the style transformation block. The last feature F_6 goes through a single TORGB convolution layer and predicts the final output \hat{I}_H .

The $\Phi_{ST}(\cdot)$ blocks are the key parts of our framework. Each of them learns the style transformation parameters $\mathbf{y}_i = (\mathbf{y}_{s,i}, \mathbf{y}_{b,i})$ for F_i from corresponding scale of input pairs, *i.e.*, LQ images and parsing maps denoted as (I_L^i, I_P^i) . (I_L^i, I_P^i) are resized to the same size as F_i through bicubic interpolation. Then, $\Phi_{ST}(\cdot)$ can be formulated as follows

$$\begin{aligned} \mathbf{y}_i &= \Psi(I_L^i, I_P^i), \quad (2) \\ F_i &= \mathbf{y}_{i,s} \frac{\Phi_{UP}(F_{i-1}) - \mu(\Phi_{UP}(F_{i-1}))}{\sigma(\Phi_{UP}(F_{i-1}))} + \mathbf{y}_{i,b}, \quad 1 < i \leq 6 \end{aligned} \quad (3)$$

where $\Psi(\cdot)$ is a lightweight network composed of several convolution layers. $\mu(\cdot)$ and $\sigma(\cdot)$ are the mean and standard variation of features. Compared with StyleGAN which adpots spatially invariant styles, we follow [9, 33] and compute the spatially adaptive style parameters \mathbf{y}_i with the same size as F_i . This helps to make full use of the spatial-wise color and texture information from I_L as well as shape and semantic guidance from I_P . We use the pretrained face parsing network (FPN) to generate I_P from LQ inputs I_L .

3.2. Semantic-Aware Style Loss

Recent super-resolution work [8] has demonstrated that gram matrix loss which is usually used in style transfer helps a lot in recovering textures. To better synthesize texture details, we introduce the *semantic-aware style loss* \mathcal{L}_{ss} which calculates the gram matrix loss for each semantic region separately. We use VGG19 features of layer `relu3_1`, `relu4_1` and `relu5_1` to calculate \mathcal{L}_{ss} . Denote ϕ_i as the i -th layer feature in VGG19 and the parsing mask with label j as M_j (the background is denoted as M_0), the semantic aware style loss is formulated as

$$\mathcal{L}_{ss} = \sum_{i=3}^5 \sum_{j=0}^{18} \|\mathcal{G}(\phi_i(\hat{I}_H), M_j) - \mathcal{G}(\phi_i(I_H), M_j)\|_2, \quad (4)$$

$\mathcal{G}(\cdot)$ computes the gram matrix of feature $\phi_i(\cdot)$ with semantic label mask M_j as below

$$\mathcal{G}(\phi_i, M_j) = \frac{(\phi_i \odot M_j)^T (\phi_i \odot M_j)}{\sum M_j + \epsilon}, \quad (5)$$

where \odot is the element-wise product, and $\epsilon = 1e-8$ is used to avoid zero division.

3.3. Model Objectives

Following previous image restoration works [4, 34, 32, 21, 20], we apply reconstruction loss and adversarial loss apart from L_{ss} .

Reconstruction Loss. It is the combination of pixel and feature space mean square error (MSE) which aims to constrain the network output \hat{I}_H close to ground truth I_H . We formulate the reconstruction loss as

$$\mathcal{L}_{rec} = \|\hat{I}_H - I_H\|_2 + \sum_s \sum_{k=1}^4 \|D_s^k(\hat{I}_H^s) - D_s^k(I_H^s)\|_2 \quad (6)$$

The second term in \mathcal{L}_{rec} is the multi-scale feature matching loss [32] which matches the discriminator features of \hat{I}_H and I_H . $s \in \{1, \frac{1}{2}, \frac{1}{4}\}$ is the downscale factors, and $D_s^k(\cdot)$ denotes the k -th layer features in D_s .

Adversarial Loss. It has been proved to be effective and critical in generating realistic textures in image restoration

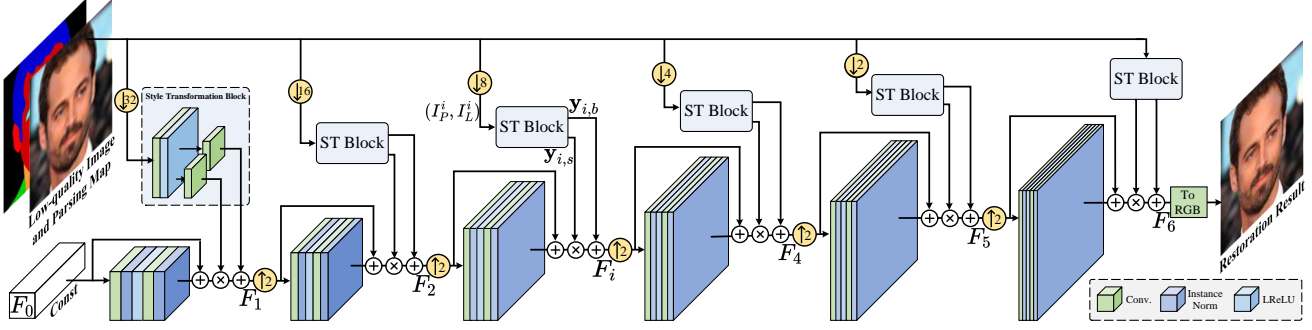


Figure 2: Visualization of the proposed progressive semantic-aware style transformation network for face restoration.

tasks. In this work, we use multi-scale discriminators and hinge loss as the objective function, defined as

$$\mathcal{L}_{GAN_G} = \sum_s -\mathbb{E}(\hat{I}_H^s), \quad (7)$$

$$\mathcal{L}_{GAN_D} = \sum_s \left\{ \mathbb{E}[\max(0, 1 - D_s(I_H^s))] + \mathbb{E}[\max(0, 1 + D_s(\hat{I}_H^s))] \right\}, \quad (8)$$

For training stability, we incorporate spectral normalization [26] in both generator and discriminators. In summary, the final loss function for our generator network is defined as

$$\mathcal{L}_G = \lambda_{ss}\mathcal{L}_{ss} + \lambda_{rec}\mathcal{L}_{rec} + \lambda_{adv}\mathcal{L}_{GAN_G}, \quad (9)$$

where λ are the weights for different terms. PSFR-GAN is trained by minimizing \mathcal{L}_G and \mathcal{L}_{GAN_D} alternatively.

4. Implementation Details

4.1. Degradation Model

According to previous works [22, 35] and common practice in SISR framework, we generate the LR image I_L^r with the following degradation model:

$$I_L^r = ((I_H \otimes \mathbf{k}_\varrho) \downarrow_r + \mathbf{n}_\delta)_{JPEG_q}, \quad (10)$$

where \otimes represents the convolution operation between the HQ image I_H and a blur kernel \mathbf{k}_ϱ with parameter ϱ . \downarrow_r is the downsampling operation with a scale factor r . \mathbf{n}_δ denotes the additive white Gaussian noise (AWGN) with a noise level δ . $(\cdot)_{JPEG_q}$ indicates the JPEG compression operation with quality factor q . The hyper parameters ϱ, r, δ, q are randomly selected for each HR image I_H , and I_L^r is generated online. More details are described in the appendix. After we get I_L^r , $I_L = (I_L^r) \uparrow_r$ is used for the parsing map prediction and face restoration.

4.2. Datasets

Training Data. We adopt the FFHQ [15] as the training dataset. This dataset consists of 70,000 high-quality images at a size of 1024×1024 . All images are automatically

cropped and aligned. We resize the images to 512×512 with bilinear downsampling as the ground-truth HR images, and synthesize the LQ inputs online with Eq. 12 where the parameters are randomly selected.

Testing Data. We construct two testing datasets, a synthetic one and a real one. For the synthetic test dataset, we randomly choose 2,800 HQ images from CelebAHQ [14] which has no identity intersection with FFHQ, and then generate the corresponding LQ images in the same way as training dataset. We denote this synthetic test dataset as **CelebAHQ-Test**. For the real LQ test dataset, we collect 1,020 faces smaller than 48×48 from CelebA [23] and 106 images provided by GFRNet [22]. The GFRNet-Test contains LR images from VGGFace2 [3] and IMDB-WIKI [30]. We also collect some old photos from the internet for testing. All images are cropped and aligned in the same manner as FFHQ, and then resized to 512×512 using bicubic upsampling. We merge all these images together and create a new dataset containing 1,157 real LQ faces, denoted as **PSFR-RealTest**.

4.3. Training Details

We use Adam optimizer [17] to train our networks. We choose $\beta_1 = 0.5, \beta_2 = 0.999$, and set the learning rate of the generator and discriminator to 0.0001 and 0.0004 respectively. The trade-off parameters of different losses are set as $\lambda_{rec} = 10, \lambda_{ss} = 100$ and $\lambda_{adv} = 1$. The training batch size is set to 4. All models were implemented by PyTorch [28] and trained on a Tesla V100 GPU.

5. Experiments

In this section, we conduct experiments to compare our framework with other methods on both synthetic and real LQ test datasets, and carry out extensive ablation studies to evaluate the effectiveness of the multi-scale parsing map guidance and *semantic aware style loss*.



Figure 3: Visual comparisons on CelebAHQ-Test dataset. The proposed PSFR-GAN shows the best results for LQ inputs with light degradation, severe degradation, large pose and different skin colors. The predicted parsing maps of *LQ inputs* through pretrained FPN are overlapped at the corner of GT images. For better visualization experience, we only show results with top-5 FID scores in Table 1. Please zoom in to see the details.

Table 1: Quantitative comparison on different restoration tasks with state-of-the-art methods. The test datasets are generated with FFHQ-Test using random parameters of each specific degradation type.

Task	Methods	CelebAHQ-Test					PSFR-RealTest
		PSNR \uparrow	SSIM \uparrow	MSSIM \uparrow	LPIPS \downarrow	FID \downarrow	FID \downarrow
JPEG artifacts removal	ARCNN	22.78	0.6538	0.7462	0.5862	133.38	124.46
Deblur	DeblurGANv2	22.66	0.6587	0.7493	0.5546	113.85	97.42
Super-Resolution	ESRGAN	21.95	0.6096	0.7293	0.5515	97.02	57.51
	Super-FAN	22.71	0.6527	0.7459	0.4908	94.95	65.45
	WaveletSRNet	<u>23.50</u>	0.6595	0.7542	0.5409	111.60	108.21
Blind-Restoration	HiFaceGAN	21.50	0.5495	0.6900	0.4569	57.81	56.48
	DFDNet	22.28	<u>0.6589</u>	<u>0.7650</u>	<u>0.3791</u>	<u>37.34</u>	<u>37.63</u>
	PSFR-GAN (ours)	23.64	0.6557	0.7740	0.3042	23.20	30.39

5.1. Evaluation Metrics

For CelebAHQ-Test with ground truth, we take the widely used PSNR, SSIM and MSSIM metrics. However, these pixel space metrics prefer smooth results and are not consistent with human perception, therefore we also adopt LPIPS score [43] to evaluate the perceptual realism of generated faces. For natural LQ images without ground truth,

we use FID score [11] to measure the statistic distance between the restoration results and a reference HQ face dataset. Compared with other no reference metrics such as NIQE [25] and IS (Inception Score) [29] which focus on natural images, FID utilizes a reference HQ face dataset and gives better measurement. We use the ground truth images from CelebAHQ-Test as the reference dataset to evaluate results of PSFR-RealTest. We also provide FID scores for




Figure 4: Visual comparisons on PSFR-RealTest dataset. Results of the proposed PSFR-GAN are clearer and more realistic. We show results with top-5 FID scores in Table 1. More results are provided in appendix.

results of CelebAHQ-Test for reference.

5.2. Comparison on Synthetic Datasets

We first evaluate the performance of different methods on the synthetic CelebAHQ-Test dataset. Following HiFaceGAN [36], we compare the proposed PSFR-GAN with architectures designed for different restoration tasks: AR-CNN [6] for JPEG artifacts removal; DeblurGANv2 [18] for image deblurring; ESRGAN [34] for natural image super-resolution; Super-FAN [1] and WaveletSRNet [12] for face super-resolution; and recent methods HiFaceGAN [36] and DFDNet [20] for blind face restoration. Table 1 shows the performance of both statistical metrics (PSNR, SSIM, MSSIM) and perceptual metrics (LPIPS, FID). We can observe that methods designed for specific tasks generally show low perceptual scores than blind face restoration methods. Both DFDNet and PSFR-GAN utilize extra face prior, and their performance is much better than HiFaceGAN. Compared with DFDNet, the proposed PSFR-GAN outperforms by a large margin in terms of most evaluation metrics, especially in FID score (50% improvement), which indicates the superiority of our proposed PSFR-GAN.

The visualization examples in Fig. 3 help us understand the quantitative results better. We show three kinds of LQ inputs in each row of Fig. 3: LQ with light degradation, LQ with severe degradation, LQ with large pose and different skin color. It can be observed that results of both ESR-

GAN and Super-FAN are over smoothed and fail to recover clear face components and textures compared with blind face restoration methods, see 2-nd and 3-rd rows. This indicates that methods designed for specific task cannot handle LQ inputs with complicated blind degradation. The results of HiFaceGAN are clearer and sharper but contain too much artifacts in the mouth, eyes and backgrounds. This is most likely because of the unstable training of GAN without face prior guidance. Different from PSFR-GAN which utilizes parsing map, DFDNet needs to first detect the LQ face components and then matches them to a HQ dictionary. It may fail to find the correct reference when the LQ are too blurry or with large pose. For example, the results of DFDNet in the 3-rd row seems to have two eyeballs which do not exist in the other methods. Meanwhile, PSFR-GAN incorporates the semantic information through parsing map which is more generic than explicit HQ reference. Therefore, the results of PSFR-GAN are more realistic and robust with all kinds of LQ inputs. Moreover, our results are generated using the predicted parsing maps of the pretrained FPN, see last column in Fig. 3. This means PSFR-GAN does not need extra information such as HQ dictionaries during the test time, making PSFR-GAN available in most situations.

5.3. Comparison on Real World LQ Images

The final target of all methods is to restore real world LQ face images. To evaluate the generalization ability of





Figure 5: Visual comparison between PULSE and PSFR-GAN. Complete results are provided in the appendix. Please zoom in to see the details.


different methods, we also compare the performance of PSFR-GAN on PSFR-RealTest dataset with methods in Table 1. The same as previous results on CelebAHQ-Test, methods designed for blind face restoration still outperform the others. FID score of PSFR-GAN surpasses other methods by a large margin, and is 20% higher than the second best result of DFDNet. Fig. 4 gives some examples from PSFR-RealTest dataset. We can observe that there are many artifacts in the results of HiFaceGAN. DFDNet seems to have difficulties finding the correct component references for real LQ inputs, for example, teeth in the first row and right eye in the second row. In contrast, the results of our method is much more natural with few conspicuous artifacts. This is because PSFR-GAN generates the results in a progressive way, and the multi-scale parsing maps provide multi-scale style guidance which makes PSFR-GAN able to synthesis realistic textures for each semantic regions, *e.g.*, teeth textures and eyes. The parsing map results in the last column illustrate that the pretrained FPN works well for real LQ inputs. We show more results of PSFR-GAN on PSFR-RealTest datasets in appendix, which demonstrate that PSFR-GAN is practical for real world scenarios.

5.4. Comparison with PULSE

PULSE [24] is a recent popular method for face restoration. Different from other methods, PULSE is an optimization based method which needs carefully finetune for each LQ input. Therefore, it is unfair to compare the quantitative result of PULSE with others on PSFR-RealTest. Instead, we follow HiFaceGAN and use the historic group photograph of famous physicists taken at the 5th Solvay Conference 1927 for visual comparison. We carefully finetune PULSE on these photos and get the best results as much as we can. Even so, we still observe several typical fail-



(a) Multi-scale LQ and parsing map input pairs (I_L^i, I_P^i) .



(b) Results of adding multi-scale inputs (I_L^i, I_P^i) progressively.



(c) Keep all scales of LQ inputs I_L^i , and add different scales of parsing maps I_P^i progressively.



(d) Keep all scales of parsing map inputs I_P^i , and add different scales of LQ inputs I_L^i progressively.

Figure 6: Analysis of PSFR-GAN. Zoom in to see details.

ure cases of PULSE shown in Fig. 5: (1) age mismatch and shape deformation in the first column; (2) large pose in the second column; (3) background interference in the third column. The final column shows the best result of the test photos, but there are many slight shape changes in the eyes, nose and mouth which make it look like another person. To conclude, although PULSE can generate better details, for example hair textures, the results of our PSFR-GAN are still much better than PULSE in face restoration. Besides, PSFR-GAN is 40 times faster than PULSE (0.1s vs. 4s) to process the same image on GPU. Complete results of Solvay conference are in the appendix.

6. Discussions and Ablation Study

6.1. Analysis of PSFR-GAN

Why PSFR-GAN works ? There are two key designs in PSFR-GAN: (1) the features are modulated progressively from coarse-to-fine; (2) the multi-scale input pairs I_L and I_P work together to provide sufficient color, shape and semantic information. In Fig. 6(b), we show how details are progressively added to the final outputs by first zeroing all inputs and then adding them back from coarse to fine. We can observe that the restoration process is consistent with




Figure 7: Results for different upscale factors while fixing other degradation parameters.

our hypothesis, and happens in the following order: high level semantic information comes first, then the mid-level shape and edges, and finally low-level color and details. Similarly, we analyze the effect of I_P^i and I_L^i separately in Fig. 6(c) and Fig. 6(d). The first column of Fig. 6(c) and Fig. 6(d) show the result of using I_L^i and I_P^i as inputs separately. We can see that in the first image of Fig. 6(c), the nose and mouth borders are not clear and there are many artifacts in the cheek region. This indicates that network without parsing map as inputs only makes the bicubic results (first row and last column in Fig. 6(b)) sharper and has difficulties in understanding the semantic meaning of each region. When we add I_P^i progressively, the artifacts are gradually removed and the edges are clearer. As for Fig. 6(d), we can observe that the semantic regions are clear in all stages, for example the nose and teeth part. Then, color and texture details are added with the entry of I_L^i . In summary, PSFR-GAN restores the LQ images in a progressive way by modulating the features with multi-scale inputs, where I_L^i provides the low level color and texture information and I_P^i contributes the semantic and shape information.

Robustness to different degrees of degradation As an example, we verify the effectiveness of the PSFR-GAN for different upscale factors in Eq. 12 and fix other parameters. We can observe from Fig. 7 that our network works well for upscale factors ≤ 12 and can produce reasonable result for $\times 16$, which demonstrates the robustness of PSFR-GAN.

6.2. Ablation Study

To explore the effectiveness of parsing map guidance and *semantic-aware style loss* \mathcal{L}_{ss} , we evaluate four variants of our framework in Table 2: **A**, baseline model with only I_L as inputs; **B**, baseline model with (I_L, I_P) as inputs but without \mathcal{L}_{ss} ; **C**, baseline model with I_L as inputs and \mathcal{L}_{ss} ; **D**, the proposed PSFR-GAN. We can observe that parsing map plays an important role in face image restoration and achieves the most improvements, and the \mathcal{L}_{ss} can also benefit the restoration results. With the combination of them, our PSFR-GAN can achieve the best performance. Figure 8 shows some example results on PSFR-RealTest. We can ob-




Figure 8: Visual comparison between different variations of our model. Please zoom in to see the details.

Table 2: Ablation study of the proposed method.

ID	Model Variations	FID↓
A	Baseline with I_L	47.48
B	+ I_P	32.37
C	+ \mathcal{L}_{ss}	34.22
D	+ $I_P + \mathcal{L}_{ss}$	30.39

serve from the bottom row that models without parsing map (*i.e.*, model **A** and **C**) fail to generate clear shapes when the face border in LQ image is not clear. On the other hand, models without \mathcal{L}_{ss} (*i.e.*, model **A** and **B**) produce apparent artifacts especially in the eyes. In contrast, model **D** with parsing map and \mathcal{L}_{ss} has none of the above flaws. It can be inferred from the above observation that 1) parsing map helps to regularize face structure, and 2) \mathcal{L}_{ss} helps to synthesize realistic textures for each semantic region.

7. Conclusion

This paper proposes a multi-scale progressive face restoration network, named PSFR-GAN, which restores LQ face inputs in a coarse to fine manner through semantic-aware style transformation. We also proposed the *semantic-aware style loss* based on original gram matrix loss. Experiments on both synthetic and real LQ test datasets demonstrate the superiority and robustness of our PSFR-GAN. By pretraining the face parsing network (FPN) for LQ inputs, our framework can generate high-resolution and realistic HQ outputs without requiring extra inputs. In summary, PSFR-GAN provides a robust and easy-to-use solution for face restoration in real-world scenarios.

8. Acknowledgements

This work was partially supported by Alibaba DAMO Academy, Hong Kong RGC RIF grant (R5001-18), and Hong Kong RGC GRF grant (project# 17203119).

References

- [1] Adrian Bulat and Georgios Tzimiropoulos. Super-fan: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [2] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [3] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face and Gesture Recognition*, 2018.
- [4] 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 (CVPR)*, 2018.
- [5] Kim Deokyun, Kim Minseon, Kwon Gihyun, and Kim Dae-Shik. Progressive face super-resolution via attention to facial landmark. In *British Machine Vision Conference (BMVC)*, 2019.
- [6] Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 576–584, 2015.
- [7] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, 2016.
- [8] Muhammad Waleed Gondal, Bernhard Schölkopf, and Michael Hirsch. The unreasonable effectiveness of texture transfer for single image super-resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 80–97. Springer, 2018.
- [9] Jinjin Gu, Hannan Lu, Wangmeng Zuo, and Chao Dong. Blind super-resolution with iterative kernel correction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1604–1613, 2019.
- [10] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3012–3021, 2020.
- [11] 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. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 6626–6637, 2017.
- [12] Huaibo Huang, Ran He, Zhenan Sun, and Tieniu Tan. Wavelet-srnet: A wavelet-based cnn for multi-scale face super resolution. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1689–1697, 2017.
- [13] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [14] 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.
- [15] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410, 2019.
- [16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. *arXiv preprint arXiv:1912.04958*, 2019.
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8878–8887, 2019.
- [19] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. *arXiv preprint arXiv:1907.11922*, 2019.
- [20] Xiaoming Li, Chaofeng Chen, Shangchen Zhou, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. Blind face restoration via deep multi-scale component dictionaries. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [21] Xiaoming Li, Wenyu Li, Dongwei Ren, Hongzhi Zhang, Meng Wang, and Wangmeng Zuo. Enhanced blind face restoration with multi-exemplar images and adaptive spatial feature fusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [22] Xiaoming Li, Ming Liu, Yuting Ye, Wangmeng Zuo, Liang Lin, and Ruigang Yang. Learning warped guidance for blind face restoration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [23] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [24] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2437–2445, 2020.
- [25] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2012.
- [26] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [27] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, volume 32, pages 8026–8037, 2019.
- [29] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 2234–2242, 2016.
- [30] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [31] Yibing Song, Jiawei Zhang, Shengfeng He, Linchao Bao, and Qingxiong Yang. Learning to hallucinate face images via component generation and enhancement. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4537–4543, 2017.
- [32] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [33] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [34] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)*, September 2018.
- [35] Xiangyu Xu, Deqing Sun, Jinshan Pan, Yujin Zhang, Hanspeter Pfister, and Ming-Hsuan Yang. Learning to super-resolve blurry face and text images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 251–260, 2017.
- [36] Lingbo Yang, Chang Liu, Pan Wang, Shanshe Wang, Peiran Ren, Siwei Ma, and Wen Gao. Hifacegan: Face renovation via collaborative suppression and replenishment. 2020.
- [37] Xin Yu, Basura Fernando, Bernard Ghanem, Fatih Porikli, and RICHARD HARTLEY. Face super-resolution guided by facial component heatmaps. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 217–233, 2018.
- [38] Xin Yu, Basura Fernando, Richard Hartley, and Fatih Porikli. Super-resolving very low-resolution face images with supplementary attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [39] Xin Yu and Fatih Porikli. Ultra-resolving face images by discriminative generative networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–333. Springer, 2016.
- [40] Xin Yu and Fatih Porikli. Face hallucination with tiny unaligned images by transformative discriminative neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 2, page 3, 2017.
- [41] Xin Yu and Fatih Porikli. Hallucinating very low-resolution unaligned and noisy face images by transformative discriminative autoencoders. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3760–3768, 2017.
- [42] Kaipeng Zhang, ZHANPENG ZHANG, Chia-Wen Cheng, Winston Hsu, Yu Qiao, Wei Liu, and Tong Zhang. Super-identity convolutional neural network for face hallucination. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 183–198, 2018.
- [43] 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 (CVPR)*, 2018.
- [44] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [45] Shizhan Zhu, Sifei Liu, Chen Change Loy, and Xiaoou Tang. Deep cascaded bi-network for face hallucination. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.

Appendix

A. Pretrain of LQ Face Parsing Network

A.1. Network Architecture

Given a LQ face image of any size, we first upsample it to 512×512 and treat it as the input I_L . FPN is trained to produce a parsing map \hat{I}_P and a HR face \hat{I}_H that approximate the ground-truth parsing map I_P and ground-truth HR face I_H respectively, *i.e.*,

$$\theta_p = \arg \min_{\theta_p} \mathcal{L}_{parse}(\hat{I}_P, I_P) + \mathcal{L}_{pix}(\hat{I}_H, I_H), (11)$$

where θ_p denotes the parameters of FPN, \mathcal{L}_{parse} is the parsing loss, and \mathcal{L}_{pix} is the pixel space L2 loss. As shown in Fig. 9, FPN adopts an encoder-resnet-decoder architecture. It begins with 4 downsample blocks, followed by 10 resnet blocks and 4 upsample blocks. Finally, two output convolution layers are used to generate \hat{I}_P and \hat{I}_H . We adopt BatchNorm-LeakyRelu after every convolution layer.

We use multi-task learning for FPN because we found that \mathcal{L}_{pix} is quite helpful for the prediction of \hat{I}_P . Since I_L is degraded, both the pixel values and pattern of the face components are not clear and stable. The network is not able to understand the meaning of each label without the extra supervision of I_H . Fig. 10 shows that the parsing results with multi-task learning are much better than that without it, especially in the eyes and eyebrows.

A.2. Datasets and Implementation

We use CelebA-Mask-HQ [19] to train FPN. The CelebA-Mask-HQ contains 30,000 HR faces with a size of 1024×1024 selected from the CelebA [23] dataset. Each image has a segmentation mask of facial attributes corresponding to CelebA. The masks of CelebA-Mask-HQ are manually-annotated with a size of 512×512 and 19 classes including background, skin, nose, eyes (left and right), eyebrows (left and right), ears (left and right), mouth, lips (up and bottom), hair, hat, eyeglass, earring, necklace, neck, and cloth. The whole dataset is split into a training set (24,183 images), a validation set (2,993 images), and a test set (2,824 images). We use the training set as ground-truth HQ faces and parsing maps, and the LQ faces are generated online with Eq. 12.

We use Adam optimizer [17] to train the FPN. We set $\beta_1 = 0.9$, $\beta_2 = 0.999$ and learning rate to 0.0002. The training batch size is set to 8.

B. Degradation Model

As described in the paper, our degradation model used the following equation:

$$I_L^r = ((I_H \otimes \mathbf{k}_\varrho) \downarrow_r + \mathbf{n}_\delta) JPEG_q, (12)$$

where

- \mathbf{k}_ϱ is the blur kernel. We randomly choose one of the following four kernels: Gaussian Blur ($3 \leq \varrho \leq 15$), Average Blur ($3 \leq \varrho \leq 15$), Median Blur ($3 \leq \varrho \leq 15$), Motion Blur ($5 \leq \varrho \leq 25$);
- \downarrow_r is the downsample operation. The scale factor r is randomly selected in $[\frac{32}{512}, \frac{256}{512}]$;
- \mathbf{n}_δ is the additive white gaussian noise (AWGN) with $0 \leq \delta \leq 0.1 \times 255$;
- $JPEG_q$ is the JPEG operation. The compression level is randomly chosen from $[10, 65]$, in which higher means stronger compression and lower image quality.

We implement the degradation model using `imgaug`¹ library with code snippets in Fig. 11.

C. More Results

In this section, we show more results on PSFR-RealTest and Solvay conference test. We mainly compare our model with DFDNet because they provide public codes and test models, and their results are current state-of-the-art. We also provide carefully finetuned results of PULSE on Solvay test.

C.1. Results of PSFR-RealTest

Fig. 12, Fig. 13 and Fig. 14 show more examples from PSFR-RealTest dataset.

C.2. Results of Solvay Conference Test

We give the overall results of the 5-th Solvay conference test images in Fig. 15. All faces are cropped out and aligned first, then enhanced by our model and finally paste back to the original photo. Complete results and detailed comparison with other methods are presented in Fig. 16, Fig. 17, Fig. 18, Fig. 19 and Fig. 20.

¹<https://github.com/aleju/imgaug>

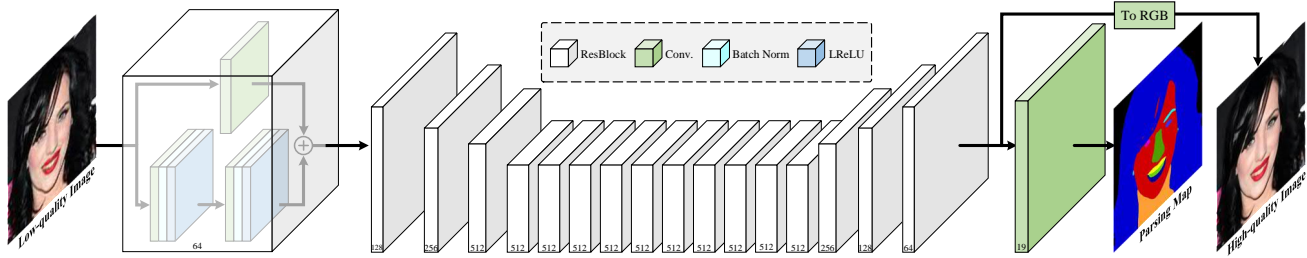


Figure 9: Architecture details of face parsing network (FPN).




Figure 10: Comparison of parsing results of natural LR faces with and without supervision of I_H .


```

import imgaug as ia
import imgaug.augmenters as iaa
scale_size = random(32, 256)
org_size = 512
aug_seq = iaa.Sequential([
    iaa.Sometimes(0.5, iaa.OneOf([
        iaa.GaussianBlur((3, 15)), iaa.AverageBlur(k=(3, 15)),
        iaa.MedianBlur(k=(3, 15)), iaa.MotionBlur((5, 25))
    ])),
    iaa.Resize(scale_size, interpolation=ia.ALL),
    iaa.Sometimes(0.2, iaa.AdditiveGaussianNoise(loc=0, scale=(0.0, 0.1*255), per_channel=0.5))

    iaa.Sometimes(0.7, iaa.JpegCompression(compression=(10, 65))),
    iaa.Resize(org_size),
])

```

Figure 11: Code snippets for degradation model.




LQ

DFDNet

PSFR-GAN (ours)

Parsing Map (ours)

Figure 12: More results from PSFR-RealTest Dataset.



LQ

DFDNet

PSFR-GAN (ours)

Parsing Map (ours)

Figure 13: More results from PSFR-RealTest Dataset.



LQ

DFDNet

PSFR-GAN (ours)

Parsing Map (ours)

Figure 14: More results from PSFR-RealTest Dataset.



Figure 15: Overall result of the 5-th Solvay conference taken in 1927. Please zoom in to see the details.



Figure 16: Results of 5-th Solvay conference test.



Figure 17: Results of 5-th Solvay conference test.



Figure 18: Results of 5-th Solvay conference test.



Figure 19: Results of 5-th Solvay conference test.




Figure 20: Results of 5-th Solvay conference test.