

GCFSR: a Generative and Controllable Face Super Resolution Method Without Facial and GAN Priors

Jingwen He¹ Wu Shi^{2,3} Kai Chen¹ Lean Fu¹ Chao Dong^{2,4,*}

¹ByteDance Inc, ²ShenZhen Key Lab of Computer Vision and Pattern Recognition, SIAT-SenseTime Joint Lab, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences,

³Guangdong-Hong Kong-Macao Joint Laboratory of Human-Machine Intelligence-Synergy Systems, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China,

⁴Shanghai AI Laboratory, Shanghai, China.

hejingwenhejingwen@outlook.com, wu.shi@siat.ac.cn,

{chenkai.ck9838, fulean}@bytedance.com, chao.dong@siat.ac.cn



Figure 1. The results of our proposed GCFSR on CelebA-HQ [16] for $32\times$ SR. Left: The original 32×32 LR input image is upsampled ($\times 32$) to a resolution 1024^2 by bicubic interpolation. The middle three images are the modulated results of GCFSR under different conditional upscaling factors ($s = 24, s = 32, s = 48$). The conditional upscaling factor s controls the generative strength. The best SR result ($s = 32$) is denoted by a yellow rectangle. Right: The ground truth (GT) at 1024^2 resolution. **(Zoom in for best view)**

Abstract

Face image super resolution (face hallucination) usually relies on facial priors to restore realistic details and preserve identity information. Recent advances can achieve impressive results with the help of GAN prior. They either design complicated modules to modify the fixed GAN prior or adopt complex training strategies to finetune the generator. In this work, we propose a generative and controllable face SR framework, called GCFSR, which can reconstruct images with faithful identity information without any additional priors. Generally, GCFSR has an encoder-generator architecture. Two modules called style modulation and feature modulation are designed for the multi-factor SR task. The style modulation aims to generate realistic face details and the feature modulation dynamically fuses the multi-level encoded features and the generated ones conditioned on the upscaling factor. The simple and elegant architecture can be trained from scratch in an end-

to-end manner. For small upscaling factors (≤ 8), GCFSR can produce surprisingly good results with only adversarial loss. After adding L1 and perceptual losses, GCFSR can outperform state-of-the-art methods for large upscaling factors (16, 32, 64). During the test phase, we can modulate the generative strength via feature modulation by changing the conditional upscaling factor continuously to achieve various generative effects. Code is available at <https://github.com/hejingwenhejingwen/GCFSR>.

1. Introduction

Face image super resolution (face SR or face hallucination) algorithms have been developing rapidly in recent years, with its wide application in video restoration and AI photographing. Face SR has a close relationship with general image SR [8, 9, 31–33, 38, 41] and face generation tasks [10, 15, 17, 18, 20]. Similar as general image SR, face SR is a restoration problem, whose goal is to reconstruct correct structures and preserve identity information. Differ-

* Corresponding author

ently, face SR has to deal with very large upscaling factors (8-64) [4, 30, 35], thus requiring to generate a large amount of finer details, which is similar to face generation. As a combination of restoration and generation problem, face SR has unique solution pipelines, which always involve various additional facial priors [2, 5, 6, 21, 21, 28, 36, 43], like parsing map and attribution map.

Recent advances have found that a face GAN can take the place of all previous facial priors, and produce realistic face details. This is based on the observation that a well-trained GAN model has already contained enough face information, which is sufficient to provide strong priors. For instance, GLEAN [4] adopts the intermediate features of a StyleGAN [20] as latent banks, and achieves superior performance on large-factor SR tasks. While GFPGAN [30] and GPEN [35] introduce face GAN models to solve blind face restoration problem, and both can recover promising facial details. Their success can be attributed to the utilization of GAN priors and motivates later works to find more applications.

However, if we must rely on such a prior, face SR will face two apparent limitations. First, as face GAN is trained on specific datasets (e.g., FFHQ [19]), the corresponding face SR methods can only deal with the same kind of face images (e.g., frontal faces with a fixed size), significantly restricting its applications. Second, as face GAN is not specially designed for face SR, we have to add additional operations in the network for adaptation [4, 25, 30, 35], which is a waste of computation resources. Then we will ask: can we get rid of these priors, and design a pure data-driven framework?

Another issue unsolved in face SR is the flexibility of generation. Existing methods can only output a single restoration result with a fixed style. However, in real scenarios, users might want to adjust the generative strength to meet personalized requirements. For example, they will desire more details in old photo restoration, but less hallucination effects in surveillance video enhancement. "How to control the amount of generated details" is a practical demand. Furthermore, real-world images may have various sizes, but conventional SR models (e.g., GLEAN [4] and ESRGAN [33]) for fixed upscaling factors cannot handle such diverse cases.

To address the problems, we propose a generative and controllable face SR framework, called GCFSR, which has three appealing properties. First, it could reconstruct faithful images with promising identity information. This is also the basic requirement of face SR task. Second, it could generate realistic face details, without reliance on any additional priors, including facial priors and GAN priors. This shows that GAN prior is not an essential part in face SR task. Third, its generative strength can be interactively adjusted (Figure 1). This can also be used in handling different

and continuous upscaling factors. These three properties are guaranteed by three special designs in GCFSR, which are the encoder-generator architecture, style modulation and feature modulation modules. GCFSR enjoys a very concise architecture without extra priors or initialization. We will detail our designs in the Method section. More importantly, GCFSR has a nice training property. It is end-to-end trainable and converges fast. When the upscaling factor is small (≤ 8), it is possible to discard all pixel-wise constraints and use a single GAN loss to achieve state-of-the-art performance. This has never been revealed in previous SR methods. Extensive experiments and ablation studies have demonstrated the effectiveness of each module. Combing them together, GCFSR could achieve superior performance to GAN based methods in both small and large upscaling factors. We can also observe vivid face details and gradually modulated effects in qualitative results (see Figure 1).

2. Related Work

2.1. Face Super Resolution.

We can divide face super resolution works into two groups according to the utilization of facial priors. On the one hand, the works in [3, 13, 29, 37, 39] directly use deep neural networks for face SR without any additional facial priors. On the other hand, more recent works focus on the investigation in facial priors to preserve the identity information as well as generate faithful face details. In general, facial landmarks [2, 6, 21, 21, 43], face parsing maps [5, 28] and facial attributes [36] have been demonstrated to be effective in the face image reconstruction. Chen et al. [6] predict landmark heatmaps and parsing maps from LR faces, then use them to further finetune the SR results. Differently, the work in [2] learns face SR and landmark prediction jointly. Yu et al. [36] utilize a convolution neural network to obtain face component heatmaps in order to achieve improvement for face super resolution.

Recently, significant advances have been made by using face GAN prior [4, 30, 35] instead of the previous facial priors. Generally, these state-of-the-art methods all design additional modules to extract feature maps and latent vectors, then use them to adapt the face GAN prior to handle face super resolution/restoration tasks.

2.2. GAN Prior

The pretrained GAN priors [1, 18, 20] have been deeply exploited in GAN inversion [11, 23, 25, 26, 42]. In PULSE [23], the latent code of GAN prior is iteratively optimized with L1 constraint between the input and downsampled output. While mGANprior [11] optimizes multiple latent codes to improve the capacity for reconstruction. Instead of relying on latent code alone, DGP [25] also finetunes the pretrained GAN prior for better performance. Since the GAN

inversion methods only use the insufficient low-dimension latent code for image reconstruction, they usually generate undesirable results with low fidelity. To solve this issue, GLEAN [4] uses an additional RRDBNet [33] to extract the multi-resolution features, which will be fused with the intermediate features in GAN prior. GLEAN is designed for SR on single upscaling factor and achieves the state-of-the-art performance. GPEN [35] and GFPGAN [30] both achieve the state-of-the-art for blind face restoration (BFR) problem. The two GAN-prior-based methods also adopt additional encoders to extract multi-resolution features and combine them with intermediate features of the pretrained GAN prior. More descriptions of these state-of-the-art GAN-prior-based methods could be found in Method 3.1.

Different from them, our proposed GCFSR adopts an end-to-end training strategy without any additional priors (facial prior or GAN prior), and achieves the state-of-the-art performance on face SR. In addition, our method provides the flexibility for user adjustment on the generative strength.

3. Method

3.1. Rethinking on GAN-prior-based Methods.

Before introducing the proposed GCFSR, we give a brief review on previous state-of-the-art GAN-prior-based image restoration/super resolution methods: GLEAN [4], GFPGAN [30], and GPEN [35]. Here we provide the detailed descriptions of these methods in Table 1.

Table 1. The detailed descriptions of the state-of-the-art GAN-prior-based methods: GLEAN [4], GFPGAN [30], GPEN [35].

	GLEAN [4]	GFPGAN [30]	GPEN [35]	GCFSR
Degradation	single	multiple	multiple	multiple
Network Description	encoder GAN prior decoder	encoder CS-SFT GAN prior	encoder Concat GAN prior	/
Parameters	188.29M	90.76M	71.21M	66.69M
GAN prior	fixed	fixed	finetuned	/

First of all, GLEAN [4] is proposed for image super resolution on single upscaling factor, while GFPGAN and GPEN could deal with multiple degradations. From Table 1, we observe that GLEAN relies heavily on additional modules for SR. Specifically, GLEAN utilizes a RRDBNet [33] for feature extraction, and then combines the fixed GAN prior with an additional decoder to generate the final output. Thus, GLEAN has significantly more parameters than the other methods. Similarly, GFPGAN [30] adopts an additional UNet [27] trained with L1 loss for the degradation removal, and then transforms the features in Unet to the parameters of scaling and shifting operations, which will be used to modify the fixed GAN prior. With the pretrained GAN prior, GLEAN and GFPGAN could achieve better performance than others at the beginning of training (see Figure 2). Differently, GPEN [35] directly concatenates the

features from the encoder and GAN prior. As the concatenation operation introduces new parameters to GAN prior, GPEN gives the GAN prior a small learning rate for further finetuning. However, as shown in Figure 2, this strategy leads to slow convergence and inferior performance compared with other methods.




Figure 2. The convergence curves of GCF SR (ours), GFPGAN [30], GPEN [35], and GLEAN [4]. The x and y axes denote the training iterations (k) and FID scores on CelebA-HQ for 64× SR.

In conclusion, previous GAN-prior-based methods either design complicated modules to modify the fixed GAN prior, or further finetune the GAN prior for adaptation. These observations indicate that utilizing GAN prior in face restoration/super resolution is not a trivial task. If this is the case, can we design a new generative model without reliance on the pretrained GAN prior? Thus, in this work, we propose a very concise architecture – GCFSR without extra priors. As we can see from Figure 2, the end-to-end trainable GCFSR converges fast and outperforms the state-of-the-art GAN-prior-based methods.

3.2. Overview of GCFSR

Given an input LR face image $x = \downarrow_s(\mathbf{y})$ and the up-scaling factor s , GCFSR aims to estimate an SR face image $\hat{\mathbf{y}}$, which is as close as possible to its ground truth \mathbf{y} . To achieve this goal, GCFSR learns a mapping function $G(x, s) \rightarrow \mathbf{y}$, where s comes from a set of upscaling factors (e.g., $s \in \{4, 8, 16, 32, 64\}$), and the size of \mathbf{y} is 2^u (e.g., 512, 1024). On the other hand, the target upscaling factor s corresponds to the strength to generate the missing details in the down-scaling process. During testing, the generative strength could be smoothly adjusted by changing the conditional upscaling factor s continuously.

The overall framework of GCFSR is depicted in Figure 3. In general, it consists of an encoder network and a generator network. The encoder network takes the LR face image x as input and extracts the face structure roughly. It also estimates the latent codes w in $\mathcal{W}+$ space [20] for the generative process. The generator network takes the topmost encoded feature maps as well as the latent codes w to generate realistic face details by a sequence of style-modulated convolutions [20]. To handle a wide range of

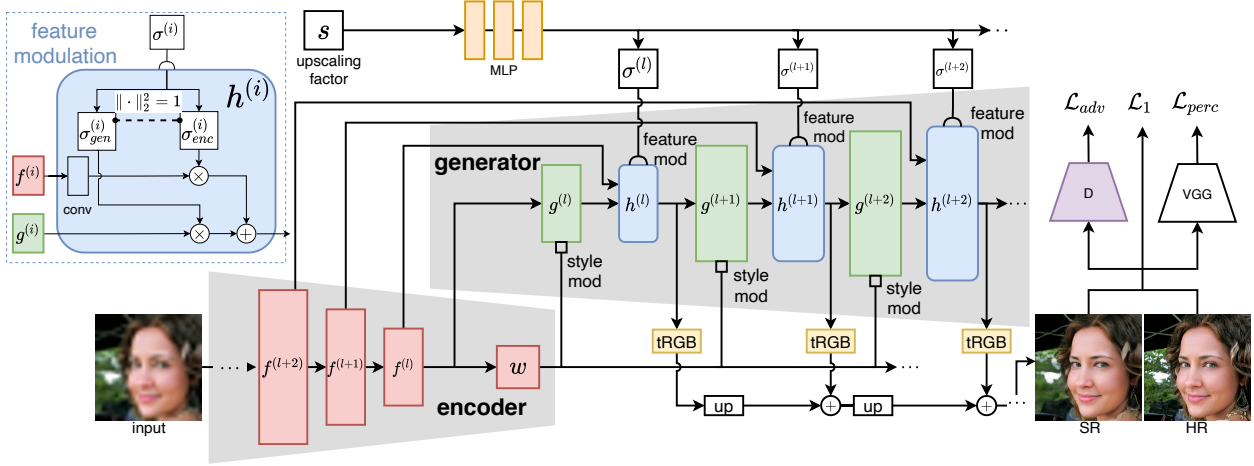


Figure 3. The architecture of GCFSR. It contains an encoder (red) and a generator (green & blue). The encoder network uses several strided convolutional layers to extract the multi-level features and latent codes w . The generator takes the topmost encoded feature maps and latent codes w to generate realistic face details by a sequence of style-modulated convolutions [20], namely as *style modulation* (green) here. While the *feature modulation* (blue) module controls how much the encoded and generated features are expressed under the conditional upscaling factor s . We train the whole network in an end-to-end manner. (The colored blocks are trained from scratch, while the other blocks are fixed or contain no trainable parameters.)

upscaling factors, we add skip connections from the extracted structural features to the corresponding generated ones. Specifically, our proposed feature modulation module at each level controls how much the encoded and generated features are expressed under the conditional upscaling factor s . The two networks collaborate closely with each other, yielding realistic results with high fidelity and providing the flexibility for user adjustment.

The framework can be trained in an elegant end-to-end manner without any GAN prior pretraining or finetuning using complex learning objectives. Our method can obtain visually pleasing results using only adversarial loss for small upscaling factors (4, 8). After adding L1 and perceptual losses, our method can achieve state-of-the-art performance for large upscaling factors (16, 32, 64).

3.3. Encoder Network

The encoder network is a simple convolution neural network (CNN) with stride of 2. The intermediate features are denoted by $\{\dots, f^{(l+1)}, f^{(l)}\}$, where the superscript indicates the base 2 logarithm of feature size. The topmost feature map, $f^{(l)}$, has a size of 2^l . The input LR image is resized by bicubic interpolation to the size of 2^u , i.e., $\uparrow_*(\mathbf{x}) \in \mathbb{R}^{2^u \times 2^u \times 3}$. Formally, we define

$$f^{(i)} = \begin{cases} \text{Conv}(\uparrow_*(\mathbf{x})), & i = u, \\ \text{Conv}(f^{(i+1)}), & l \leq i < u, \end{cases} \quad (1)$$

where Conv denotes a convolution layer with bias and activation. The encoded features are used to carry multi-level structural information of the input image. Besides, we add several convolution layers and a fully-connected (FC) layer on the topmost feature to estimate the latent codes,

$w = [w^{(l)}, w_{1,2}^{(l+1)}, \dots]$, in $\mathcal{W}+$ space. The latent codes are further used by a style-based generator to generate realistic facial details.

$$w = \text{Reshape}(\text{FC}(\text{Conv} \dots (f^{(l)}))). \quad (2)$$

3.4. Generator Network

The generator takes the topmost encoded feature maps and latent codes, w , to generate realistic facial details by *style modulation*. While the *feature modulation* controls how much the encoded and generated features are expressed at each level given the conditional upscaling factor s . We will elaborate on the details of the two modules as follows.

Style modulation. Style-modulated convolution is proposed in StyleGAN2 [20]. It uses a latent vector to modulate the convolution kernel on the input channel dimension. To approximately preserve the variances between input and output neurons, the kernel is channel-wisely normalized before conducting the convolution. Bias, noise and activation are added to the output. (Please refer to the original paper for more details.) We denote the whole module by Conv_{sm} . The generation starts from a 4-by-4 constant feature map, c . The feature map, $g^{(i)}$, progressively grows in size via up-sampling, \uparrow_2 , and Conv_{sm} .

In our method, we make several modifications for adaptation to SR tasks. First, we do not re-generate from the small constant feature map. Instead, we replace it by the topmost encoded feature, i.e., $c = f^{(l)}$ and let l be equal to the minimum input size. Then, our proposed feature modulation will join the multi-level encoded features $(\{\dots, f^{(l+1)}, f^{(l)}\})$ and the generated features $(\{\dots, g^{(l+1)}, g^{(l)}\})$ under the condition of upscaling factor

s . The fused result is denoted by $h^{(i)}$ at level i . Thus, the generator can benefit from both encoded and generated features and does not need to generate everything from scratch. Formally, we define the style modulation as

$$g^{(i)} = \begin{cases} \text{Conv}_{\text{sm}}(c, w^{(l)}), & i = l, \\ \text{Conv}_{\text{sm}}(\text{Conv}_{\text{sm}}(\uparrow_2(h^{(i-1)}), w_1^{(i)}), w_2^{(i)}), & i > l. \end{cases} \quad (3)$$

Feature modulation. To handle the multi-factor SR in a single model, the amounts of input information to preserve and output details to generate may vary depending on the upscaling factor. It is difficult for the add/concat skip connection to fulfill the various requirements of consistency and generative capacity for different SR tasks. We thus propose the feature modulation to flexibly adjust the generative strength. The upscaling factor, s , is first transformed by an MLP to a set of scaling vectors, $\sigma = \{\sigma_{1,2}^{(l)}, \sigma_{1,2}^{(l+1)}, \dots\}$. At level i , $\sigma_1^{(i)}$ and $\sigma_2^{(i)} \in \mathbb{R}^{\text{chan}(i)}$ are used to channel-wisely adjust the contributions of $f^{(i)}$ and $g^{(i)}$, where $\text{chan}(i)$ denotes the dimension of feature channel. To satisfy the assumption of unit-variance activations in the style modulation [20], we add a Conv layer after each $f^{(i)}$ to initially adjust the activations, and normalize the scaling vectors to be positive and to have channel-wise unit L2 norm. Formally, we define the feature modulation as follows

$$\begin{aligned} \sigma &= \text{Reshape}(\text{MLP}(s)), \\ \sigma_{enc/gen}^{(i)} &= \frac{|\sigma_{1/2}^{(i)}|}{\sqrt{\sigma_1^{(i)2} + \sigma_2^{(i)2} + \epsilon}}, \quad l \leq i \leq u, \quad (4) \\ h^{(i)} &= \sigma_{enc}^{(i)} \cdot \text{Conv}(f^{(i)}) + \sigma_{gen}^{(i)} \cdot g^{(i)}, \end{aligned}$$

where $\epsilon = 1e-8$ and the last two equations are both channel-wise operations.

Output. The output image is progressively computed from the fused feature, $h^{(i)}$, via tRGB layer [20]. We up-sample and sum all the intermediate RGB outputs to derive the final output, $\hat{y} = \hat{y}^{(u)}$:

$$\hat{y}^{(i)} = \begin{cases} \text{tRGB}(h^{(i)}), & i = l, \\ \uparrow_2(\hat{y}^{(i-1)}) + \text{tRGB}(h^{(i)}), & l < i \leq u. \end{cases} \quad (5)$$

3.5. Training Details

Recall that GCFSR takes the LR image x and the conditional upscaling factor s as input and estimates an SR face image, $\hat{y} = G(x, s)$. We create the LR images by down-sampling the ground-truth images and then up-sample them to the original size by bicubic interpolation. During training, the upscaling factor s is randomly chosen from $\{4, 8, 16, 32, 64\}$ and normalized to the range of $[0, 1]$.

GCFSR could be well trained with adversarial loss, where y and $\hat{y} = G(x, s)$ are treated as real and fake images respectively. We use the non-saturating logistic loss:

$$\begin{aligned} \mathcal{L}_{adv,D} &= \mathbb{E}_{y,s} [\log(1 + \exp(-D(y))) \\ &\quad + \log(1 + \exp(D(G(\downarrow_s(y), s)))] \quad (6) \\ \mathcal{L}_{adv,G} &= \mathbb{E}_{y,s} [\log(1 + \exp(-D(G(\downarrow_s(y), s)))] \end{aligned}$$

As shown in Table 4, our GCFSR_{adv} trained with only adversarial loss outperforms other blind face restoration methods for 4× and 8× SR tasks.

To further boost the SR performance, we use the conventional combinations: L1, perceptual [14] and adversarial losses. The overall training objectives are as follows:

$$\begin{aligned} \mathcal{L}_D &= \lambda_{adv} \mathcal{L}_{adv,D}, \\ \mathcal{L}_G &= \lambda_{l1} \cdot \mathcal{L}_1 + \lambda_{perc} \mathcal{L}_{perc} + \lambda_{adv} \mathcal{L}_{adv,G}. \end{aligned} \quad (7)$$

The hyper-parameters are set as: $\lambda_{l1} = 1$, $\lambda_{perc} = 0.01$, and $\lambda_{adv} = 0.01$. D and G are trained to minimize \mathcal{L}_D and \mathcal{L}_G respectively. Although GCFSR is trained from scratch, it converges fast (see Figure 2) and achieves the best performance among the GAN-prior-based methods.

3.6. Testing

During testing, the upscaling factor of the given LR image could not be strictly in the set of $\{4, 8, 16, 32, 64\}$. A by-product is that GCFSR can achieve continuous SR effect for upscaling factors in the range of $[4, 64]$. The users are encouraged to modulate the factor s to obtain various super-resolved images with different generative strength and the best one could be found during the modulation process. For example, given a LR image downsampled with factor 48, the users may find the satisfactory results between $s = 32$ and $s = 64$. As shown in Figure 1, 5, the modulation process yields smooth transitions without artifacts.

4. Experiments

Datasets and Implementation We train our GCFSR on the FFHQ dataset [19], consisting of 70k high-quality 1024×1024 face images. For testing dataset, we follow GLEAN [4] to extract 100 images from CelebA-HQ [16] dataset. We use bicubic interpolation to perform downscaling/upscaling. For evaluation, we employ the widely used non-reference perceptual metrics: FID [12] and NIQE [24]. We also adopt pixel-wise metrics (PSNR and SSIM) and the perceptual metric (LPIPS [40]). In addition, we measure the cosine similarity on the ArcFace [7] embedding space. The training mini-batch size is set to 24. We augment the training data with horizontal flip. We train our model with Adam optimizer for a total of 300k iterations. The learning rates for the generator and discriminator are both set to 2×10^{-3} . We implement our models with the PyTorch framework and train them using a NVIDIA Tesla V100 GPU.



Figure 4. Qualitative comparisons on CelebA-HQ for 16 \times (first row), 32 \times (second row), 64 \times (third row) SR. The GT image (Right) has a resolution of 1024². **Zoom in for best view.**

Table 2. Quantitative comparisons with state-of-the-art methods on CelebA-HQ for 16 \times , 32 \times , 64 \times SR. GLEAN uses three models, while the others use a single model for three SR tasks. **Red** and **blue** indicate the best and the second best performance. Similarity represents cosine similarity of ArcFace Embeddings.

	16 \times (64 ² \rightarrow 1024 ²)					32 \times (32 ² \rightarrow 1024 ²)					64 \times (16 ² \rightarrow 1024 ²)				
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	similarity \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	similarity \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	similarity \uparrow
PULSE [23]	20.08	0.6032	0.4651	77.89	0.4947	19.63	0.5848	0.4789	78.26	0.5177	18.52	0.5604	0.5119	91.55	0.4680
mGANprior [11]	23.80	0.6674	0.4668	100.87	0.6794	21.26	0.6117	0.5099	105.62	0.5230	18.69	0.5721	0.5530	108.17	0.4397
GLEAN [4]	26.88	0.6953	0.2693	29.99	0.9682	24.34	0.6534	0.3257	46.57	0.7750	21.38	0.6016	0.4109	62.93	0.6118
GPEN [35]	26.51	0.6988	0.2827	37.94	0.9473	24.65	0.6717	0.3340	50.40	0.7641	22.20	0.6291	0.3906	67.50	0.5978
GFPGAN [30]	27.07	0.7101	0.2716	34.49	0.9623	24.81	0.6751	0.3128	46.00	0.7881	22.26	0.6285	0.3675	59.33	0.6558
GCFSR (Ours)	27.17	0.7100	0.2604	30.48	0.9631	24.95	0.6748	0.3061	43.34	0.7911	22.39	0.6315	0.3663	57.15	0.6620

Comparison with state-of-the-art methods. We compare our GCFSR with several state-of-the-art methods: GAN inversion methods including PULSE [23] and mGANprior [11], and GAN-prior-based methods including GLEAN [4], GFPGAN [30] and GPEN [35]. We provide the quantitative and qualitative results for 16 \times , 32 \times , 64 \times SR tasks. Note that GLEAN is designed for SR task on single upscaling factor, thus we train three GLEAN models for different tasks. For fair comparison, we train GFPGAN, GPEN and our proposed GCFSR on the same training dataset towards the same learning objectives. In GFPGAN, the parameters of GAN prior are fixed during training, while the parameters of GAN prior are given a smaller learning rate (2×10^{-4}) in GPEN as suggested by [35].

The quantitative results are presented in Table 2. As can be seen, our GCFSR achieves the best performance in terms of PSNR and LPIPS for all three upscaling factors, indicating the superiority of GCFSR’s ability in image reconstruction with high perceptual quality. Although GLEAN is trained on a single upscaling factor, it only performs well on 16 \times SR. As GLEAN adopts a simple progressive upsampling strategy in the network design without any skip connections, the reconstruction quality cannot be guaranteed when the upscaling factor is large. On the other hand, the GAN inversion methods, PULSE and mGANprior, achieve

significantly worse results compared with GAN-prior-based methods since they can hardly preserve the identity by the simple latent code exploration strategy.

We also show qualitative results in Figure 4. It is observed that GAN inversion methods fail to maintain a good fidelity, while the GAN prior based methods achieve overall satisfactory results in term of the identity preservation. However, as shown in the second row in Figure 4, GLEAN and GFPGAN are unable to maintain the original color of the eyes since their GAN prior is fixed during training. GPEN performs better than them but shows some distortions. In general, our proposed GCFSR could achieve both fidelity and naturalness with the highest quality among all those methods. More visual comparisons could be found in the supplementary file.

Evaluation of modulation on generative strength. In this section, we provide qualitative results in Figure 1, 5 to illustrate that our GCFSR could modulate the generative strength smoothly across different levels. In Figure 1, the target upscaling factor 32 is in the predefined set {4, 8, 16, 32, 64}, so we could directly specify the conditional upscaling factor to $s = 32$ and obtain an SR result with high perceptual quality. Furthermore, we could modulate s around $s = 32$ and obtain a blurry effect ($s = 24$) or a




Figure 5. The results obtained by modulation on the generative strength. We change the conditional upscaling factor s from $s = 4$ to $s = 64$ continuously, and find satisfactory results (e.g., results denoted by yellow rectangles) between two ends. **Zoom in for best view.**

strong generative effect ($s = 48$). For target upscaling factor that is not predefined, we could still obtain satisfactory results through modulation. The LR image in the first row of Figure 5 is downsampled with factor 24. As can be seen, the modulated results between $s = 24$ and $s = 32$ (denoted by yellow rectangles) are all satisfactory. Similarly, for LR image downsampled with factor 48, we could modulate the conditional upscaling factor s continuously within the range of $[48, 64]$ and then arrive at a point with vivid and natural texture details (e.g., $s = 48$). The modulation process yields smooth transition without any noticeable artifacts. More results are in the supplementary file.

The Effectiveness on Blind Face Restoration In this section, we investigate the effectiveness of our method on blind face restoration (BFR) task. To create the blind version of our GCFSR, we fix the upscaling factor s to a constant value (e.g., 1). As suggested in [30], we create CelebA-Test with 3,000 CelebA-HQ images from its testing partition [16]. All the images of FFHQ and CelebA-Test datasets are resized to 512^2 . Then, we adopt the degradation model in GFPGAN [30] to synthesize the training and test input images. Note that we directly use the officially released models of the state-of-the-art blind face restoration methods: DFDNET [22], PSFRGAN [5], GPEN [35], and GFPGAN [30]. The quantitative results are presented in Table 3. It is observed that our blind model could achieve the best performance in PSNR, SSIM, LPIPS, and the cosine similarity of ArcFace Embeddings. Besides, we could obtain comparable results in terms of FID. This indicates the effectiveness of our method for blind face restoration.

Analysis and visualization on feature modulation. In this section, we investigate the mechanism of how the scaling vectors work to achieve effective feature modulation for different conditional upscaling factors. As we have mentioned before, the scaling vectors σ_{enc} and σ_{gen} channel-

Table 3. Quantitative comparison on CelebA-Test for blind face restoration. **Red** and **blue** indicate the best and the second best performance. Similarity represents cosine similarity of ArcFace Embeddings.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	similarity \uparrow
DFDNET	23.51	0.6674	0.4342	58.72	0.5980
PSFRGAN	24.66	0.6439	0.4199	43.33	0.6464
GPEN	24.63	0.6477	0.4004	41.99	0.6993
GFPGAN	24.65	0.6725	0.3646	42.61	0.7156
Ours	26.49	0.7120	0.3356	42.23	0.7257

wisely adjust the features from the encoder and generator, respectively. Here we provide the histograms of scaling vectors that correspond to level 64: σ_{enc}^{64} and σ_{gen}^{64} , which are illustrated in Figure 6. For σ_{enc}^{64} , its values are approaching 0 as the conditional upscaling factor s increases. Conversely, the values of σ_{gen}^{64} are approaching 1. This indicates that higher conditional upscaling factor corresponds to stronger generative effect, since the features from the encoder are weakened while the features from the decoder are strengthened. Similar trend could be found at other levels, presented in the supplementary file.

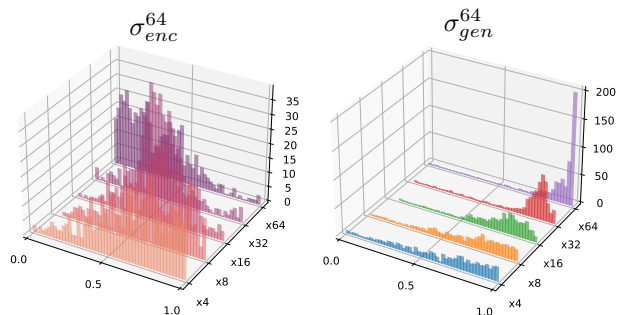


Figure 6. The visualization on feature modulation. The histograms of scaling vectors σ_{enc}^{64} and σ_{gen}^{64} for different conditional upscaling factors are presented.

Table 4. Quantitative comparison of GCFSR_{adv} (trained with only one adversarial loss) and state-of-the-art blind face restoration methods on CelebA-HQ for 4×, 8× SR. **Red** and **blue** indicate the best and the second best performance.

	4× (128 ² → 512 ²)				
	PSNR ↑	SSIM ↑	LPIPS ↓	NIQE ↓	FID ↓
GFPGAN [30]	27.32	0.7686	0.1421	4.42	36.76
GPEN [35]	27.10	0.7593	0.1534	4.18	43.81
HiFaceGAN [34]	27.39	0.7397	0.1409	4.06	30.28
DFDNet [22]	26.47	0.7802	0.1838	4.41	45.80
PSFRGAN [5]	27.24	0.7607	0.1611	4.79	35.85
GCFSR_{adv} (Ours)	27.81	0.7711	0.1210	3.84	27.90
	8× (64 ² → 512 ²)				
	PSNR ↑	SSIM ↑	LPIPS ↓	NIQE ↓	FID ↓
GFPGAN [30]	25.63	0.7069	0.1736	4.26	42.88
GPEN [35]	25.94	0.7038	0.1848	4.31	47.32
HiFaceGAN [34]	25.47	0.6774	0.2294	3.61	61.52
DFDNet [22]	25.03	0.7023	0.2313	4.60	52.98
PSFRGAN [5]	24.90	0.6871	0.2085	4.58	47.10
GCFSR_{adv} (Ours)	26.02	0.7139	0.1704	4.14	39.99

Effects of learning objectives. Our GCFSR is very easy to train and could achieve state-of-the-art performance for large-factor image super resolution. In this section, we will show that our GCFSR could obtain surprisingly good results for small-factor image super resolution (4×, 8×) by only using adversarial learning (without pixel-wise loss or perceptual loss). Here we compare our GCFSR_{adv} with blind face restoration (BFR) methods: GFPGAN [30], GPEN [35], DFDNet [22], PSFRGAN [5], and HiFaceGAN [34]. Note that most of the abovementioned methods only release models for output size 512 × 512. Therefore, for the convenience of comparison, we train our GCFSR_{adv} on resized 512 × 512 FFHQ training dataset and test it on resized 512 × 512 CelebA-HQ dataset. The upscaling factor s is randomly sampled from {4, 8, 16, 32}. The quantitative results are presented in Table 4. It observed that our GCFSR_{adv} achieves superior performance to blind face restoration methods. We also show visual comparisons in the supplementary file.

Effects of style modulation. Here we demonstrate the effectiveness of the style modulation module. As we mentioned before, the latent codes w are estimated by the encoder and will be used for style modulation in the generator. Unlike GAN inversion methods which utilize latent codes w to generate both global attributes (e.g., poses) and finer details, our w is more related to the latter. From Figure 7, the results with style modulation have more realistic details (e.g., eyelash and hair) and less artifacts (e.g., bad case for mouth generation). More results are in supplementary file. Besides, we provide the quantitative results in the supplementary file. In general, the style modulation module improves the overall performance in most metrics.

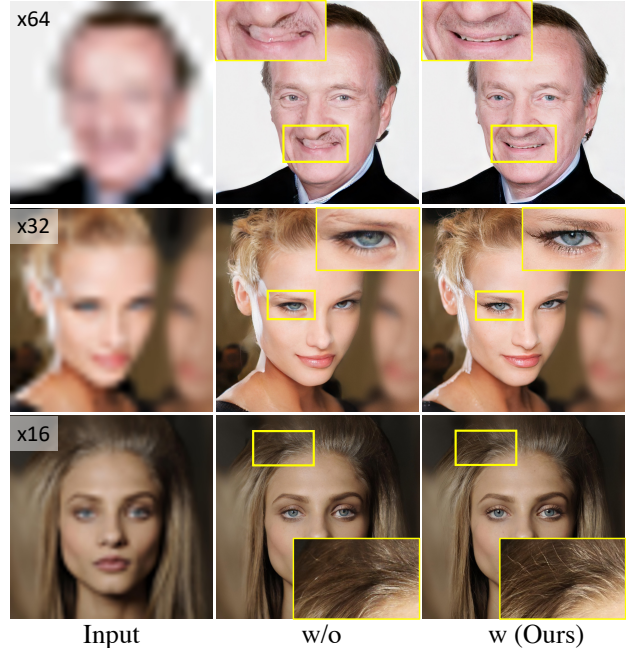


Figure 7. Visual comparison of GCFSR with and without style modulation module on CelebA-HQ for 64× (first row), 32× (second row) and 16× (third row) SR. **Zoom in for best view.**

5. Conclusion

We have presented a face SR framework called GCFSR without any additional priors but could handle very large-factor face SR (up to 64×). GCFSR has an encoder-generator architecture and is end-to-end trainable with fast convergence. In particular, the proposed style modulation module helps generate realistic face details, while the feature modulation module dynamically fuses the multi-level encoded features and generated ones under control of the conditional upscaling factor. In this way, our GCFSR could reconstruct faithful images with promising identity information and provide the flexibility for user adjustment.

Limitations Our work has several limitations. First, this work only investigates the single dimensional modulation regarding the upscaling factor. While in real-world scenarios, multi-dimensional modulation across multiple degradations should be considered. Second, although our framework shows extraordinary performance on face SR task, its generalization on general SR is waiting to be studied.

Acknowledgement This work is partially supported by the National Natural Science Foundation of China (61906184), the Joint Lab of CAS-HK, the Shenzhen Research Program(RCJC20200714114557087), the Shanghai Committee of Science and Technology, China (Grant No. 21DZ1100100).

References

- [1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018. [2](#)
- [2] 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*, pages 109–117, 2018. [2](#)
- [3] Qingxing Cao, Liang Lin, Yukai Shi, Xiaodan Liang, and Guanbin Li. Attention-aware face hallucination via deep reinforcement learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 690–698, 2017. [2](#)
- [4] Kelvin CK Chan, Xintao Wang, Xiangyu Xu, Jinwei Gu, and Chen Change Loy. Glean: Generative latent bank for large-factor image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14245–14254, 2021. [2](#), [3](#), [5](#), [6](#)
- [5] Chaofeng Chen, Xiaoming Li, Lingbo Yang, Xianhui Lin, Lei Zhang, and Kwan-Yee K Wong. Progressive semantic-aware style transformation for blind face restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11896–11905, 2021. [2](#), [7](#), [8](#)
- [6] Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, and Jian Yang. Fsrnet: End-to-end learning face super-resolution with facial priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2492–2501, 2018. [2](#)
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. [5](#)
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 38(02):295–307, 2016. [1](#)
- [9] Chao Dong, Chen Change Loy, and Xiaoou Tang. Accelerating the super-resolution convolutional neural network. In *European conference on computer vision*, pages 391–407. Springer, 2016. [1](#)
- [10] Rinon Gal, Dana Cohen, Amit Bermano, and Daniel Cohen-Or. Swagan: A style-based wavelet-driven generative model. *arXiv preprint arXiv:2102.06108*, 2021. [1](#)
- [11] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3012–3021, 2020. [2](#), [6](#)
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. [5](#)
- [13] 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*, pages 1689–1697, 2017. [2](#)
- [14] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [5](#)
- [15] 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. [1](#)
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. [1](#), [5](#), [7](#)
- [17] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *arXiv preprint arXiv:2106.12423*, 2021. [1](#)
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. [1](#), [2](#)
- [19] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. [2](#), [5](#)
- [20] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. [1](#), [2](#), [3](#), [4](#), [5](#)
- [21] Deokyun Kim, Minseon Kim, Gihyun Kwon, and Dae-Shik Kim. Progressive face super-resolution via attention to face landmark. In *the 30th British Machine Vision Conference (BMVC) 2019*. the 30th British Machine Vision Conference (BMVC) 2019, 2019. [2](#)
- [22] Xiaoming Li, Chaofeng Chen, Shangchen Zhou, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. Blind face restoration via deep multi-scale component dictionaries. In *European Conference on Computer Vision*, pages 399–415. Springer, 2020. [7](#), [8](#)
- [23] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2437–2445, 2020. [2](#), [6](#)
- [24] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. [5](#)
- [25] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [2](#)
- [26] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding

- in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021. [2](#)
- [27] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [3](#)
- [28] Ziyi Shen, Wei-Sheng Lai, Tingfa Xu, Jan Kautz, and Ming-Hsuan Yang. Deep semantic face deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8260–8269, 2018. [2](#)
- [29] Oncel Tuzel, Yuichi Taguchi, and John R Hershey. Global-local face upsampling network. *arXiv preprint arXiv:1603.07235*, 2016. [2](#)
- [30] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9168–9178, 2021. [2](#), [3](#), [6](#), [7](#), [8](#)
- [31] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1905–1914, 2021. [1](#)
- [32] 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*, pages 606–615, 2018. [1](#)
- [33] 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 (ECCV) workshops*, pages 0–0, 2018. [1](#), [2](#), [3](#)
- [34] Lingbo Yang, Shanshe Wang, Siwei Ma, Wen Gao, Chang Liu, Pan Wang, and Peiran Ren. Hifacegan: Face renovation via collaborative suppression and replenishment. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1551–1560, 2020. [8](#)
- [35] Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. Gan prior embedded network for blind face restoration in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 672–681, 2021. [2](#), [3](#), [6](#), [7](#), [8](#)
- [36] 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. [2](#)
- [37] Xin Yu and Fatih Porikli. Ultra-resolving face images by discriminative generative networks. In *European conference on computer vision*, pages 318–333. Springer, 2016. [2](#)
- [38] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. *arXiv preprint arXiv:2103.14006*, 2021. [1](#)
- [39] Kaipeng Zhang, Zhanpeng Zhang, Chia-Wen Cheng, Winston H 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. [2](#)
- [40] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [5](#)
- [41] Wenlong Zhang, Yihao Liu, Chao Dong, and Yu Qiao. Ranksrgan: Generative adversarial networks with ranker for image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3096–3105, 2019. [1](#)
- [42] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European conference on computer vision*, pages 592–608. Springer, 2020. [2](#)
- [43] Shizhan Zhu, Sifei Liu, Chen Change Loy, and Xiaoou Tang. Deep cascaded bi-network for face hallucination. In *European conference on computer vision*, pages 614–630. Springer, 2016. [2](#)