

Toward Interactive Modulation for Photo-Realistic Image Restoration

Haoming Cai¹ *Jingwen He¹ *Yu Qiao^{1,2}Chao Dong^{1,3} †

¹Key Laboratory of Human-Machine Intelligence-Synergy Systems,
Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

²Shanghai AI Lab, Shanghai, China

³SIAT Branch, Shenzhen Institute of Artificial Intelligence and Robotics for Society

{hm.cai, jw.he, yu.qiao, chao.dong}@siat.ac.cn

Abstract

Modulating image restoration level aims to generate a restored image by altering a factor that represents the restoration strength. Previous works mainly focused on optimizing the mean squared reconstruction error, which brings high reconstruction accuracy but lacks finer texture details. This paper presents a Controllable Unet Generative Adversarial Network (CUGAN) to generate high-frequency textures in the modulation tasks. CUGAN consists of two modules - base networks and condition networks. The base networks comprise a generator and a discriminator. In the generator, we realize the interactive control of restoration levels by tuning the weights of different features from different scales in the Unet architecture. Moreover, we adaptively modulate the intermediate features in the discriminator according to the severity of degradations. The condition networks accept the condition vector (encoded degradation information) as input, then generate modulation parameters for both the generator and the discriminator. During testing, users can control the output effects by tweaking the condition vector. We also provide a smooth transition between GAN and MSE effects by a simple transition method. Extensive experiments demonstrate that the proposed CUGAN achieves excellent performance on image restoration modulation tasks.

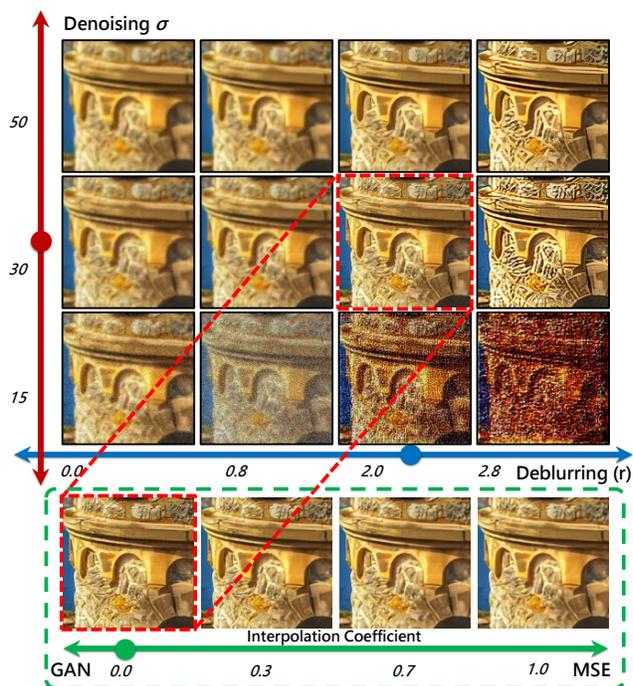


Figure 1. The input image is with blur $r=2$ (kernel width=2) and $\sigma=30$. Through the modulation on deblurring and denoising, users could obtain a well restored image with rich texture details. Besides, users are allowed to modulate on the green sliding bar to achieve smooth transition between GAN and MSE effects.

1. Introduction

Deep learning methods have achieved great success in many low-level vision tasks, such as image denoising, deblurring, and super-resolution. Various network architectures and training strategies have been continuously improving the reconstruction quality (e.g., PSNR). Later on, to pursue visually pleasing results, the generative adversarial net-

work (GAN) [13] is introduced to encourage the network to produce natural looking images. For most image restoration tasks, a deep model learns a deterministic mapping and outputs a fixed result for a pre-determined degradation type/level. For instance, many image restoration deep models are trained for a single level (e.g., Gaussian noise $\sigma=50$). However, this deterministic mapping is not flexible, as users cannot continuously modify the restoration effect or strength based on personal preference. Moreover, when the degradation model mismatches the degraded input, the out-

*Denotes equal contribution

†Denotes corresponding author (e-mail: chao.dong@siat.ac.cn)

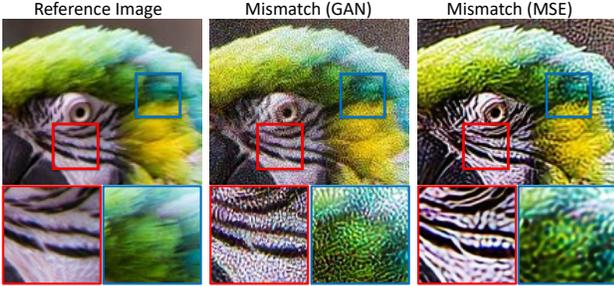


Figure 2. Artifacts incurred by applying GAN-based and PSNR-oriented model with mismatched restoration level. Models are trained on blur $r2$ +noise $\sigma30$, while the input image is with blur $r1$ +noise $\sigma50$.

put image will contain severe artifacts, especially for GAN-based models, as shown in Figure 2. Therefore, developing modulation models that can flexibly handle a range of degradations by sliding bars is essential and practical.

In recent years, several modulation methods [11, 16, 30, 31] have been proposed to adapt existing deep models to other objectives. Specifically, they can generate continuous restoration results between the pre-defined start level and end level (e.g. denoising $\sigma15 \rightarrow \sigma50$). Furthermore, CResMD [17] proposes a multi-dimension modulation framework that allows jointly modulation for different degradations. However, the above modulation methods are all PSNR-oriented, which will produce over-smoothed results without sufficient high-frequency details. To obtain modulation outputs with photo-realistic effect, this work focuses on interactive modulation for GAN-based image restoration.

The main challenges behind PSNR-oriented and GAN-based modulation are different. For PSNR-oriented modulation methods, such as CResMD, severe and mild degradations have different magnitude orders on the MSE loss, which will lead to the "unbalanced learning" problem. This phenomenon encourages models to focus on restoring images with severe degradations while ignoring the mild ones. On the contrary, in GAN-based modulation, the generator may ignore those severe degradations if we directly apply a vanilla GAN. Specifically, for a vanilla discriminator, an image restored from severe degradations will look like the fake one compared with that from mild ones. This incorrect judgment will lead to the vanishing generator gradient on those severe degradations. Therefore, we need specific discriminator to clarify images restored from various degradations. This discriminator could adjust its judgment criterion based on the degradation of the input image.

In this paper, we introduce a novel GAN-based modulation framework for photo-realistic image restoration. The proposed controllable Unet generative adversarial network (CUGAN) adopts the Unet architecture. We realize the in-

teractive control of the reconstruction result by tuning the weights of different features from different scales or within one scale in the Unet framework. To tackle vanishing generator gradient on severe degradations, we apply global feature modulation (GFM) [18] in the discriminator. Based on the severity of degradations, the global feature modulation adaptively modulates intermediate features of the discriminator. The above modulation is achieved by the condition networks embedded in the CUGAN. The condition networks accept a condition vector that encodes the degradation information for each input image. Then, the condition networks generate the tuning weights for the generator and the parameters of global feature modulation (GFM) for the discriminator. During testing, users can control the restoration strength by tweaking the condition vector. We also provide the trade-off method between GAN and MSE effect, allowing users to adjust the output effect in another dimension. Different from DNI [31], our interpolation works on both the base network G and the condition network C_G . Figure 1 shows a qualitative example for the whole modulation process.

To verify the effectiveness of the proposed method, we conduct extensive experiments on modulation for image restoration with multiple degradations. Experimental results show that the proposed CUGAN could achieve modulation with high visual quality, high reconstruction accuracy, or a compromise between them. We also demonstrate its effectiveness on image restoration in real-world scenarios. In this work, our main contributions can be summarized as follows:

- We propose a novel GAN-based modulation framework, named CUGAN, to obtain photo-realistic restoration results by interactive modulation.
- We propose a discriminator equipped with global feature modulation to ease the vanishing generator gradient on severe degradations.
- We test the restoration performance on selected datasets and real-world images. Extensive experiments show the effectiveness of CUGAN in modulated image restoration.

2. Related Work

Perceptual Image Restoration. With the rapid development of deep learning in recent years, deep learning techniques have been widely explored to tackle image restoration problems, such as image super-resolution, denoising, deblurring, and compression artifacts reduction [10, 14, 15, 34, 36, 37]. Focusing on optimizing PSNR, the above methods tend to generate blurry images. In contrast, perceptual image restoration targets to obtain better perceptual results with more texture details by applying GAN [13].

Ledig et al. [24] propose SRGAN that could generate photo-realistic images in SR task. In the PIRM2018-SR Challenge [3], ESRGAN [32] achieves state-of-the-art performance by improving the network architecture and loss functions. Benefiting from a learnable ranker, RankSRGAN [39] could generate visually pleasant images that favor different perceptual metrics.

Explorative Image Restoration. Despite that deep-learning-based methods have achieved high qualitative performance, most of them deal with image restoration problem by learning a deterministic mapping. To allow users adjusting the restoration effects, Bahat et al. [2] propose an editing module that could interactively influence the texture or brightness. With normalizing flows [9, 23], SR-Flow [27] takes a step forward to model the conditional distribution of all possible SR reconstructions given an LR facial input. Furthermore, DeepSEE [5] also adopts normalizing flow to leverage semantic maps for explorative facial super-resolution.

Modulation for Image Restoration. Although existing methods could allow user to adjust restored effect, most image restoration deep networks are trained on one specific degradation level. Therefore, it is tiring to train N various models for N various degradations. DNI [31] and AdaFM [16] find the high similarity on kernels between models trained on various levels. Based on this observation, DNI directly interpolate parameters between two related networks to attain a smooth control of diverse imagery effects. While AdaFM adopts a more efficient way: utilize depth-wise convolution layers to modulate the intermediate features instead of changing all convolution filters. Different from above two interpolation-based methods, the work in CFSnet [30] adaptively learns the interpolation coefficients and uses them to couple intermediate features from the main branch and tuning branch.

However, In real-world scenarios, images contain multiple degradations, such as blur and noise. Since different types of degradations are coherently related, modulation for multiple degradations should be conducted jointly, not independently. To address this, CResMD [17] regards the modulation as a conditional image restoration problem, and proposed a framework that accepts both corrupted images and their degradation information as input. Thus, image restoration is conditioned on the restoration/degradation information. Modulation for image restoration is essential in practical usages not merely because that it enables a unified framework to handle multiple degradations with arbitrary levels. More importantly, it provides a flexible interaction on the restoration strength.

3. Methods

Our goal is to design a GAN-based restoration model that takes in both the degraded image and desired restora-

tion information as inputs and outputs visually pleasing restored image. The restoration information, which is equal to the degradation information, can be regarded as sliding bars for users to modulate during testing. We achieve such a photo-realistic modulation model by the following approach. Given any corrupted image I_i^{dis} distorted with some degradations, we want to restore it to a clear image I_i^{res} , which is close to the ground truth image I^{gt} . To allow image restoration for different degradations, we accept a condition vector z ($z^k \in [0, 1]$). In particular, each dimension of z represents a certain degradation type while its value encodes the degradation level. Therefore, our task is to find a function F , such that: $F(I_i^{dis}, z_i) \rightarrow I^{gt}$, where $i = 1, 2, \dots, N$, and N indicates the number of all predefined degradations. In addition, we introduce a conditional discriminator, that aims to distinguish between the restored image I_i^{res} and the ground truth image I^{gt} under condition of z_i . In particular, the discriminator accepts I_i^{res} , I^{gt} , and z_i as inputs, and outputs the two probabilities for restored result and ground truth image, respectively.

3.1. Base Networks

The base networks consist of two parts: generator G and discriminator D . The generator G accepts the input image and outputs the restored result, while D aims to discriminate the restored result from the ground truth.

Generator. The architecture of the generator G is shown in Figure 3. Basically, it is a Unet framework incorporated with Residual Blocks. G consists of three scales, namely $scale_1$, $scale_2$, and $scale_3$. This multi-scale architecture could help achieve better results for GAN-based modulation compared with a plain network (Please see the experimental results in Ablation Study). For each scale, there are two residual blocks in the left and the other two in the right. Each residual block contains two 3×3 convolution layers, and a ReLU activation function between them. In $scale_1$, there are two 3×3 convolution layers at two ends. As for $scale_2$ and $scale_3$, we begin by using 2×2 strided convolution to downscale the features and end up with 2×2 transposed convolution for upscaling. The number of channels for each layer from $scale_1$, $scale_2$, $scale_3$ are set to 64, 128, and 256, respectively. Besides, each scale has a short connection that could skip its successive scale. Note that we also add a global connection to ease the difficulty for restoration on mild degradations.

Discriminator. The discriminator D is to discriminate ground truth images from restored images. It accepts 64×64 image patches, and outputs the predictions. The architecture of D is shown in Figure 3. We follow the work in [24], and design a discriminator that progressively down-sample the feature maps to a feature vector. Specifically, we use several strided convolution layers to half the sizes of feature maps each time. On the other hand, we increase

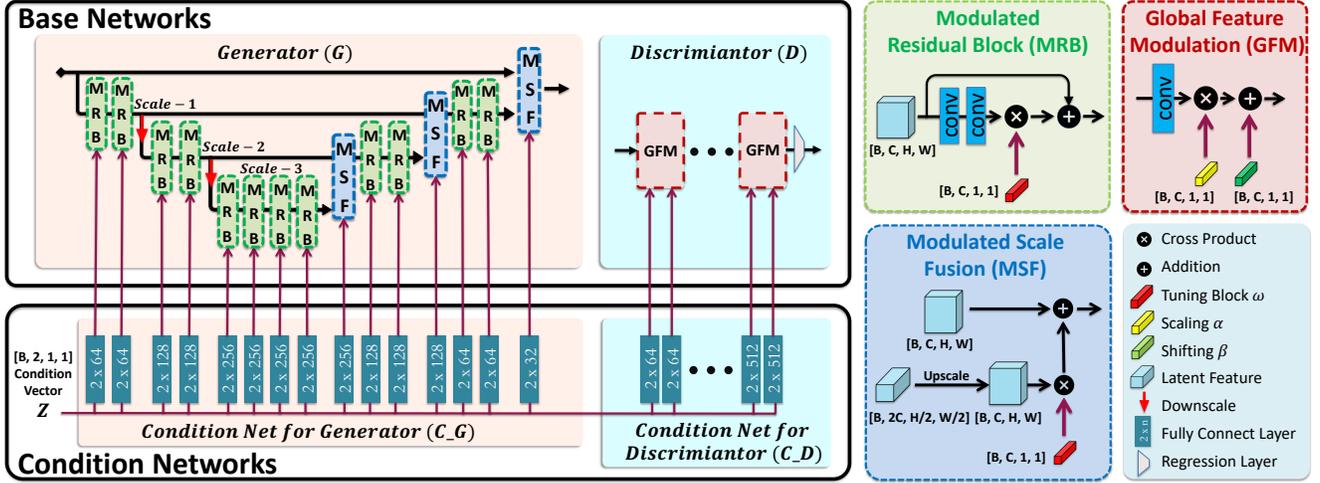


Figure 3. The framework of CUGAN. CUGAN consists of 2 base networks as well as 2 condition networks. The 2 base networks are the generator G and the discriminator D . The 2 condition networks are condition network for generator (C_G) and condition network for discriminator (C_D). During training and testing process, the condition networks C_G and C_D accept degradation information and generate tuning weights for G and D . In the modulation Unet, we provide Modulated Residual Block (MRB) and Modulated Scale Fusion (MSF) to modulate features from different scales or within one scale. As for the modulation discriminator, we provide Global Feature Modulation (GFM) to scale and shift intermediate feature maps.

the channels of feature maps by using convolution layers with doubled number of filters. There are 10 convolution layers in total. LeakyReLU [33] activation is adopted between convolution layers. After we obtaining the feature maps with 512 channels, we use global average pooling and two fully-connected layers (regression) to output the final probability.

3.2. Condition networks

The condition networks mainly contain two parts, one for the generator and the other for the discriminator, namely as C_G and C_D , respectively. Each condition network accepts a condition vector that encodes the degradation information for the input image, and outputs the parameters of modulation operations that will be used to modulate the intermediate feature in the base networks G and D .

First, the degradation information of each corrupted image x should be encoded into the condition vector z . Specifically, for each degradation type, the corresponding degradation level is scaled to a value within range $[0, 1]$. For instance, given an input image with blur level $r = 1$ ($r \in [0, 4]$) and noise level $\sigma = 10$ ($\sigma \in [0, 50]$), we could obtain a corresponding condition vector $z = [0.25, 0.20]$ by computing $[1/4, 10/50]$.

The architectures of condition networks are shown in Figure 3. Each condition network consists of several independent fully-connected layers (FC). To generate the tuning weights for the generator, the i -th fully-connected layer transforms the condition vector z of the input image x to the tuning weight w_i for the i -th modulation module in the generator. The formulation can be written as follows:

$$w_i = F_i^w(z),$$

where $F_i^w(\cdot)$ denotes the function of i -th fully-connected layer. The dimension of w_i is the same as the channel number of the corresponding feature maps to be modulated.

As for the i -th modulation module in the discriminator, we use two independent fully-connected layers to respectively generate the parameters of scaling and shifting operations for global feature modulation (GFM). In particular, we have:

$$\alpha_i = F_i^\alpha(z), \beta_i = F_i^\beta(z),$$

where $F_i^\alpha(\cdot)$ and $F_i^\beta(\cdot)$ denote the i -th fully-connected layers for generating parameters of scaling (α) and shifting (β) operators, respectively. Note that the dimensions of α_i and β_i are equal to the channel number of the intermediate feature maps in the i -th modulation module.

3.3. Modulation in Base Networks

In this section, we introduce our modulation strategies adopted in the generator and discriminator, which are illustrated in Figure 3. The parameters of modulation operations are generated by the condition networks.

For the generator, we impose controls on different features from different scales as well as within one scale. As we mentioned above, there are three scales in the generator. In different scales, the feature maps have different spatial sizes from high resolution to low resolution. In particular, for two successive scales: $scale_m$ and $scale_{m+1}$, we have:

$$x'_m = F_{m+1}^G(x_m) + x_m,$$

where $F_{m+1}^G(\cdot)$ denotes the transformation for $scale_{m+1}$ in the generator G , x_m is the feature maps in $scale_m$ before entering the $scale_{m+1}$, and x'_m is the feature maps obtained after the addition of the feature maps from $scale_{m+1}$.

Then, we formulate the modulated scale fusion (MSF) on features from $scale_m$ and $scale_{m+1}$:

$$x'_m = w_i * F_{m+1}^G(x_m) + x_m,$$

where w_i is the tuning weight generated by the i -th fully connected layer of the condition network for generator. Besides, the dimension of w_i is the same as the number of feature maps. This controlling strategy could dynamically change the weights for lower-resolution features. Intuitively, when the degradation information of the input image is extremely mild, the tuning weights will be changed close to zero, since there is no need to receive rich contextual information from low resolutions.

Within one scale, we formulate modulated residual block (MRB) to achieve control. To be more exact, we have:

$$x'_i = w_i * f(x_i) + x_i,$$

where $f(\cdot)$ denotes the transformation in the residual branch, x_i and x'_i are the original and modulated feature maps, respectively. w_i is the corresponding tuning weight.

In the discriminator, we adopt global feature modulation (GFM) to modulate the intermediate feature maps. To be more specific, we scale and then shift the intermediate features based on the condition. Practically, we perform modulation after batch normalization. In the i -th modulation module, we obtain the modulated feature maps by:

$$x'_i = \alpha_i * x_i + \beta_i,$$

where x_i and x'_i are the original and modulated feature maps. The multiplier α_i and addition operator β_i are respectively generated by two independent fully connected layers from the condition network for discriminator.

3.4. Interpolation between modulation models.

The proposed CUGAN can achieve modulation for photo-realistic image restoration. To meet different user flavors, we use the deep network interpolation strategy (DNI) [31] to achieve smooth transition between GAN and MSE effects. Specifically, we first train a PSNR-oriented modulation model with MSE loss, and obtain the networks G_{MSE} and C_{MSE} . For GAN training, we directly finetune all parameters of the PSNR-oriented modulation model, and obtain the networks G_{GAN} and C_{GAN} . Then, we interpolate all the corresponding parameters of these two modulation models to obtain an interpolated modulation model, whose parameters are:

$$\begin{aligned} \theta_G^{interp} &= (1 - \alpha)\theta_G^{GAN} + \alpha\theta_G^{MSE}, \\ \theta_C^{interp} &= (1 - \alpha)\theta_C^{GAN} + \alpha\theta_C^{MSE}, \end{aligned}$$

where θ_G^{interp} , θ_G^{MSE} and θ_G^{GAN} are the parameters of the base networks G_{interp} , G_{MSE} and G_{GAN} , respectively,

θ_C^{interp} , θ_C^{PSNR} and θ_C^{GAN} are the parameters of the condition networks C_{interp} , C_{PSNR} and C_{GAN} , respectively, and $\alpha \in [0, 1]$ is the interpolation coefficient.

For a given corrupted image x , we first specify an appropriate condition vector for photo-realistic image restoration. Then, we change the coefficient α from 0 to 1 to obtain smooth transition between GAN and MSE effects without artifacts (see Figure 1, 4).

4. Experiments

4.1. Implementation Details

We use DIV2K training dataset [1], which is a high-quality dataset and widely used in image restoration. We follow the work in [17] to generate the degraded input images. Specifically, we add Gaussian blur and Gaussian noise sequentially to every training image in DIV2K dataset with random levels. The covariance range of Gaussian noise is $\sigma \in [0, 50]$, and the range of kernel widths for 21×21 Gaussian blur is $r \in [0, 4]$. We uniformly sample the degradation levels with a stride of 0.1 and 1 for blur and noise, respectively. Besides, the degradation information for each degraded image is encoded into condition vector z .

The training process consists of two stages. First, we pre-train the PSNR-oriented modulation model by only using MSE loss, named as CUnet. The base network G and the corresponding condition network C_G are jointly trained. The learning rate is initialized as 5×10^{-4} and decayed by a factor of 2 every 2×10^5 iterations of update. The well-trained model CUnet will serve as the starting point for the next GAN training. In the second stage, the base networks, G and D , as well as the condition networks, C_G and C_D are jointly trained with the objective:

$$\mathcal{L}_{total} = \mathcal{L}_{percep} + 0.005\mathcal{L}_{GAN} + 0.01\mathcal{L}_{MSE}$$

where \mathcal{L}_{MSE} is the MSE loss, \mathcal{L}_{GAN} denotes the standard GAN loss, \mathcal{L}_{percep} represents the perceptual loss. The perceptual loss is the L1 distance between two activated features obtained from VGG19-54¹. The initial learning rate is set to 5×10^{-4} and will be halved at [50k, 100k, 200k, 300, 400k]. The number of mini-batch is set to 16, and the crop size is 64. For optimization, we use Adam [22] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Moreover, we augment the training data with random horizontal flips and 90-degree rotations. All the models are implemented with the PyTorch framework and trained on NVIDIA 1080Ti GPUs.

4.2. Evaluation for Modulation Performance

Settings. LIVE1 [29] and CBSD68 [28] are chosen to evaluate our models. We evaluate the modulation results on

¹We use pre-trained 19-layer VGG network, where 54 indicates features output by the 4th convolution before the 5th maxpooling layer

Degradation Settings		Single Degradation				Two Degradations					
blur		0	0	2	4	1	1	2	2	4	4
noise		30	50	0	0	15	30	30	50	30	50
condition vector		[0.0, 0.6]	[0.0, 1.0]	[0.5, 0.0]	[1.0, 0.0]	[0.15, 0.30]	[0.15, 0.60]	[0.50, 0.60]	[0.50, 1.00]	[1.00, 0.60]	[1.00, 1.00]
Metric	Model Name										
LPIPS ↓	UGAN	0.0490	0.0957	0.0484	0.1331	0.0697	0.1140	0.2068	0.2557	0.3154	0.3510
	CUGAN	0.0522	0.0966	0.0525	0.1463	0.0714	0.1178	0.2019	0.2498	0.3075	0.3350
	distance ↓	0.0032	0.0009	0.0041	0.0132	0.0017	0.0038	-0.0049	-0.0059	-0.0079	-0.0160
DISTS ↓	UGAN	0.0603	0.0901	0.0499	0.1013	0.0716	0.0987	0.1430	0.1705	0.1994	0.2110
	CUGAN	0.0639	0.0964	0.0525	0.1081	0.0727	0.1003	0.1392	0.1657	0.1982	0.2162
	distance ↓	0.0036	0.0063	0.0026	0.0068	0.0011	0.0016	-0.0038	-0.0048	-0.0012	0.0052

Table 1. Quantitative LPIPS/DISTS results for GAN-based modulation testing on LIVE1. ↓ means the lower the better. LPIPS/DISTS distance are bolded to stress the modulation results.

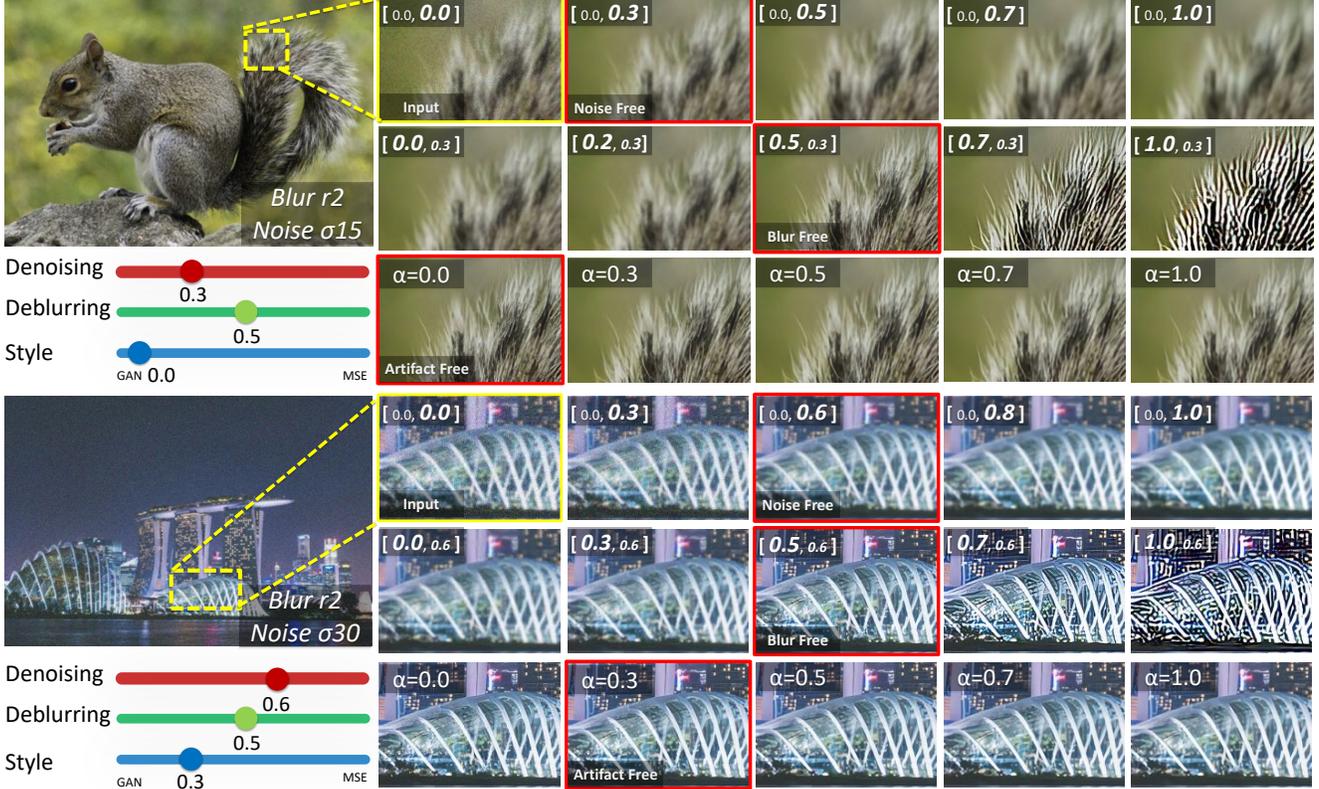


Figure 4. Qualitative results of Modulation for Image Restoration. For each group of images, three rows illustrate operations of denoising, deblurring, and interpolation, respectively. The colored-boxed images are the chosen results of the corresponding modulation process.

	Single Degradation				Two Degradations				
	blur	0	0	2	4	1	1	2	4
noise		30	50	0	0	15	30	50	50
Unet (1.364M)		30.62	28.28	30.24	26.85	29.12	27.41	24.57	23.03
CUnet (1.370M)		30.50	28.16	20.07	28.63	29.04	27.35	24.55	23.02
PSNR Distance ↓		0.12	0.12	0.17	0.22	0.08	0.06	0.02	0.01
CResMD (2.50M)		30.43	28.06	30.09	26.53	29.00	27.28	24.48	22.95

Table 2. Quantitative PSNR results for MSE-based modulation testing on CBSD68. The total parameter of the model is presented in the bracket. ↓ means the lower the better. PSNR distance are bolded to stress the modulation results.

single degradation and two degradations. For single degradation, we choose denoising σ_{30} , σ_{50} and deblurring r_2 , r_4 . For two degradations, we choose deblurring+denoising, $[r_1, \sigma_{15}]$ $[r_2, \sigma_{15}]$, $[r_2, \sigma_{30}]$, $[r_2, \sigma_{50}]$, $[r_4, \sigma_{30}]$, $[r_4, \sigma_{50}]$. We train two baseline models, baseline-MSE named Unet (CUnet w/o condition network) and baseline-GAN

named UGAN (CUGAN w/o condition networks), for each of chosen degradations. Specifically, we remove the condition networks and only use the base networks to train these baseline models. The implementation details remain the same with CUnet and CUGAN, except that every baseline model is only trained on one specific degradation level.

Based on the evaluation results in PIPAL dataset [21], LPIPS [38] and DISTS [8] have higher consistency with human ratings than NIQE, PI, and MA on images from the GAN model. Therefore, in our experiments, PSNR is used to evaluate the performance of PSNR-oriented models, while DISTS and LPIPS are used to evaluate the GAN-based models. Given the ground truth images, we calculate the PSNR of CUnet and Unet and compute the PSNR distance. Then we calculate LPIPS and DISTS of CUGAN

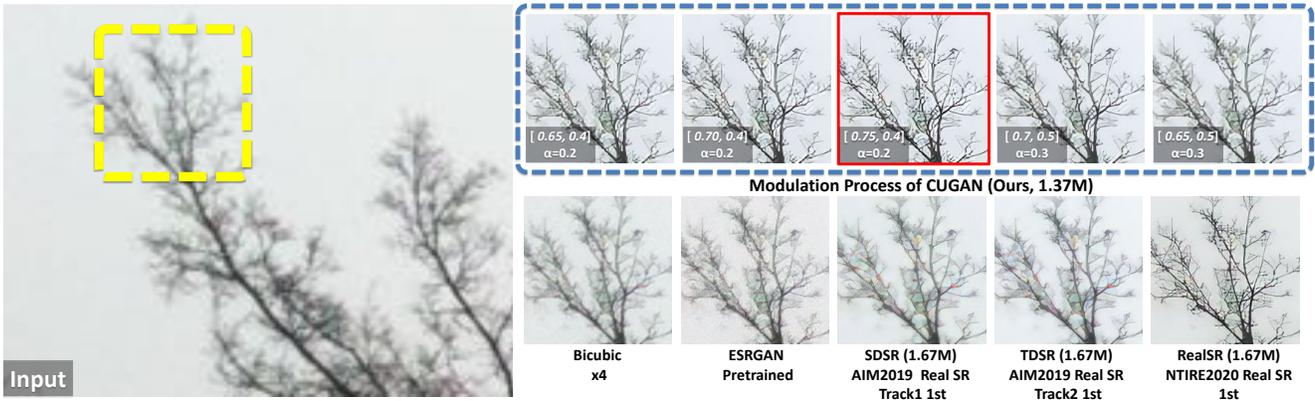


Figure 5. The qualitative comparison for blind SR task on DPED dataset [19]. The images in the first row show the modulation process of CUGAN. The corresponding condition vectors and interpolation coefficients α are denoted. The red-boxed image is the visually best result found by modulation. The number of total parameters of each model is presented in the bracket.

and UGAN and compute the LPIPS distance as well as the DISTs distance. Lower PSNR/LPIPS/DISTS distance indicates better modulation results.

Modulation for Image Restoration. Here, we evaluate the modulation performance of the proposed CUGAN quantitatively and qualitatively. The ideal performance of CUGAN and CUnet is to approach the performance of baseline models on every degradation level as close as possible. In Table 1, we evaluate the GAN-based modulation across various degradations on LIVE1 dataset. For the single degradation, the LPIPS/DISTS distances are all smaller than 0.0041/0.0068. This result demonstrates that our CUGAN could achieve modulation with high quality. As for two degradations (blur+noise), the modulation performance of CUGAN is even better than UGAN on several degradations: $[r2, \sigma30]$, $[r2, \sigma50]$, $[r4, \sigma30]$, $[r4, \sigma50]$. The quantitative results indicate that our CUGAN has better potential and ability on image restoration when degradations are more complex and severe. The first two rows of Figure 4 show the qualitative results, which present the smooth and stable transitions by modulation through the condition vector. In Table 2, we provide the PSNR-oriented modulation evaluated on the CBS68 dataset. The PSNR distances for all degradations are below 0.22dB. Furthermore, CUnet outperforms CResMD on almost all degradations. Specifically, CUnet surpasses CResMD on deblurring $r4$ with 0.32dB, and denoising $\sigma50$ with 0.22dB, which indicates a significant improvement. Moreover, CResMD contains 2.5 million parameters while ours contains 1.37 million parameters. These results prove the effectiveness and superiority of our proposed multi-dimension modulation framework.

Smooth transition between MSE and GAN effects. The proposed CUGAN provides the default modulation results with GAN effects. Meanwhile, CUnet provides reconstruction style-MSE effect, which is consistent with high reconstruction accuracy (high PSNR). We apply linear interpolation for all parameters of CUnet and CUGAN to achieve transition between GAN and MSE effect. In Figure 4, the third row illustrates the trade-off between these

methods	a = 15	a = 25	a = 30	a = 50
BM3D	31.08	28.57	27.76	25.62
TNRD	31.42	28.92	27.66	25.97
DnCNN-B	31.61	29.16	28.36	26.23
IRCNN	31.63	29.15	28.26	26.19
FFDNet	31.63	29.19	28.26	26.29
CUnet	34.53	31.80	30.89	28.48

Table 3. Comparison with different image denoising methods in term of PSNR on gray CBS68. The highest PSNR on various noise degrees are bolded.

two effects. After we have obtained a photo-realistic restored image by modulation through the condition vector, we could further smooth the image by changing an interpolation coefficient α from 0 to 1. In the transition, users could make the unpleasing textures/artifacts gradually disappear to obtain an artifact-free output.

4.3. Comparison on Image Denoising.

To show the competitive denoising ability of the proposed model, we compare CUnet trained for Section 4.2 with several state-of-the-art denoising methods: BM3D [7], TNRD [6], DnCNN [34], IRCNN [35], FFDNet [36]. We test them on gray CBS68 using the mean PSNR as the quantitative metric. To obtain degraded images, gaussian noise of different levels (e.g., 15, 25, 30, 40, and 50) are added to clean images. Then, we modulate the condition vector corresponding to the noise level to obtain restored images. Table 3 provides the quantitative comparison, which presents the competitive performance of our method on image denoising. CUnet provides competitive PSNR performance comparing with those state-of-the-art denoising methods.

4.4. Modulation on Real-World Images

To further validate the generalization ability of the proposed CUGAN, we evaluate the effectiveness of modulation on real-world images SR and denoising.

Qualitative Results on DPED. In real-world image SR, the LR images usually have unknown degradations, such as complicated downsampling kernel and noise. Therefore,

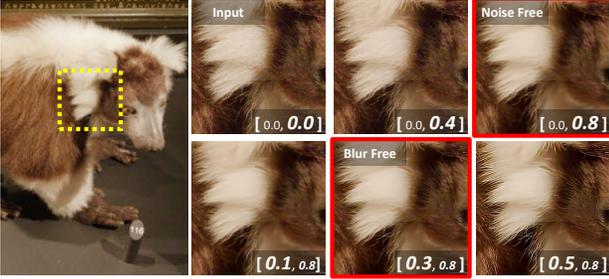


Figure 6. Modulation for image restoration in real-world images. The input image with unknown noise is from NIND [4] data set. Two rows present denoising and deblurring, respectively.

the non-blind GAN-based SR methods fail to produce satisfactory results. The proposed CUGAN could deal with this situation through modulation. In Figure 5, given an LR input from DPED [19], we first up-sample it by bicubic interpolation. Then, we obtain the restored result by changing the condition vector and interpolation coefficient. CUGAN has a competitive performance compared with those winner methods [12,20] in the AIM2019 [26] and NTIRE2020 [25] real-world image SR challenge with fewer parameters.

Qualitative Results on NIND. As for real-world image denoising, visual results are provided in Figure 6. The input image from a real-world denoising dataset NIND [4] contains noticeable but unknown noise. We first use our CUGAN to eliminate the noise by changing the condition vector from $[0.0, 0.0]$ to $[0.0, 0.8]$. The restored image is noise-free but slightly over-smoothing. Then, we sharpen this restored image by changing the first element from 0.0 to 0.3. After the modulation process, we can obtain a noise-free as well as a non-blurry restored result.

4.5. Ablation Study

Effect of the multi-scale architecture in the generator. To verify the effectiveness of the proposed multi-scale architecture, we train three models CUGAN-1 (Plain Structure), CUGAN-2, and CUGAN-3 (Ours) under the same experimental settings. Specifically, these three models have one scale, two scales, and three scales, respectively. Each model has 12 residual blocks, and each scale has the same number of residual blocks. The convergence curves on LIVE1 dataset for degradations $[r1, \sigma30]$ and $[r2, \sigma15]$ are presented in Figure 7. We use LPIPS to evaluate the performance. In Figure 7, we can observe that CUGAN-3 (Ours) surpasses CUGAN-1 and CUGAN-2 by a large margin on those degradations. Therefore, the multi-scale architecture is more effective in GAN-based modulation compared with plain networks. More convergence curves are in the supplementary.

Effect of modulation in the discriminator. In this section, we investigate the effectiveness of the modulation in the discriminator. Specifically, we train the GAN-based

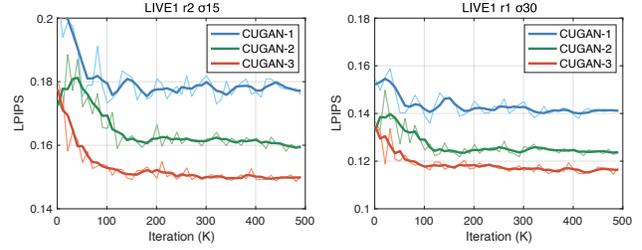


Figure 7. Convergence curves achieved by different models with different number of scale in term of LPIPS on LIVE1 dataset.

modulation model without the global feature modulation (GFM) in the discriminator and compare it with our CUGAN on different degradations evaluated by LIVE1 dataset. The performance of these two models in terms of LPIPS is shown in Figure 8. We can observe that our CUGAN equipped with GFM outperforms the other one on all those degradations. It is consistent with our assumption that the condition information is essential for the discriminator to make proper discrimination on restored images from different degradations.

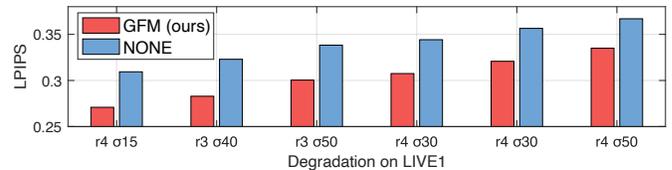


Figure 8. Results achieved by CUGAN and CUGAN without GFM.

5. Conclusion

In this work, we propose a novel GAN-based modulation framework named CUGAN. CUGAN consists of a generator and a discriminator, which are all controlled by conditions vector. The key idea of the CUGAN is introducing a modulation for users to freely control the strength of restoration and texture reconstruction with a photo-realistic effect. We apply modulated scale fusion (MSF) and modulated residual block (MRB) in the generator to achieve interactive modulation. Moreover, the Global Feature Modulation (GFM) is brought into the discriminator to ease the vanishing generator gradient. Although CUGAN could realize modulation across multiple degradations, the modulation strategy can be more effective and efficient. Better solutions are expected for future research.

Acknowledgements. This work was supported in part by the Shanghai Committee of Science and Technology, China (Grant No. 20DZ1100800), in part by the National Natural Science Foundation of China under Grant (61906184), Science and Technology Service Network Initiative of Chinese Academy of Sciences (KFJSTSQYZX-092), Shenzhen Institute of Artificial Intelligence and Robotics for Society.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 126–135, 2017. 5
- [2] Yuval Bahat and Tomer Michaeli. Explorable super resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2716–2725, 2020. 3
- [3] Yochai Blau, Roey Mechrez, Radu Timofte, Tomer Michaeli, and Lihi Zelnik-Manor. The 2018 pirm challenge on perceptual image super-resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 3
- [4] Benoit Brummer and Christophe De Vleeschouwer. Natural image noise dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 8
- [5] Marcel C Buhler, Andrés Romero, and Radu Timofte. Deepsee: deep disentangled semantic explorative extreme super-resolution. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 3
- [6] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1256–1272, 2016. 7
- [7] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 7
- [8] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *arXiv preprint arXiv:2004.07728*, 2020. 6
- [9] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. 3
- [10] 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*, pages 576–584, 2015. 2
- [11] Qingnan Fan, Dongdong Chen, Lu Yuan, Gang Hua, Nenghai Yu, and Baoquan Chen. Decouple learning for parameterized image operators. 2018. 2
- [12] Manuel Fritsche, Shuhang Gu, and Radu Timofte. Frequency separation for real-world super-resolution. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3599–3608. IEEE, 2019. 8
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1, 2
- [14] Jinjin Gu, Hannan Lu, Wangmeng Zuo, and Chao Dong. Blind super-resolution with iterative kernel correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [15] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. *2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [16] Jingwen He, Chao Dong, and Yu Qiao. Modulating image restoration with continual levels via adaptive feature modification layers. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 3
- [17] Jingwen He, Chao Dong, and Yu Qiao. Interactive multi-dimension modulation with dynamic controllable residual learning for image restoration, 2020. 2, 3, 5
- [18] Jingwen He, Yihao Liu, Yu Qiao, and Chao Dong. Conditional sequential modulation for efficient global image retouching. In *European Conference on Computer Vision*, pages 679–695. Springer, 2020. 2
- [19] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3277–3285, 2017. 7, 8
- [20] Xiaozhong Ji, Yun Cao, Ying Tai, Chengjie Wang, Jilin Li, and Feiyue Huang. Real-world super-resolution via kernel estimation and noise injection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 466–467, 2020. 8
- [21] Gu Jinjin, Cai Haoming, Chen Haoyu, Ye Xiaoxing, Jimmy S Ren, and Dong Chao. Pipal: a large-scale image quality assessment dataset for perceptual image restoration. In *European Conference on Computer Vision*, pages 633–651. Springer, 2020. 6
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [23] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018. 3
- [24] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 3
- [25] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Ntire 2020 challenge on real-world image super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 494–495, 2020. 8
- [26] Andreas Lugmayr, Martin Danelljan, Radu Timofte, Manuel Fritsche, Shuhang Gu, Kuldeep Purohit, Praveen Kandula, Maitreya Suin, AN Rajagoapalan, Nam Hyung Joon, et al. Aim 2019 challenge on real-world image super-resolution: Methods and results. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3575–3583. IEEE, 2019. 8
- [27] Andreas Lugmayr, Martin Danelljan, Luc Van Gool, and Radu Timofte. Srfow: Learning the super-resolution space with normalizing flow. In *European Conference on Computer Vision*, pages 715–732. Springer, 2020. 3

- [28] Stefan Roth and Michael J Black. Fields of experts: A framework for learning image priors. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 860–867. IEEE, 2005. 5
- [29] Hamid R Sheikh, Muhammad F Sabir, and Alan C Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on image processing*, 15(11):3440–3451, 2006. 5
- [30] Wei Wang, Ruiming Guo, Yapeng Tian, and Wenming Yang. Cfsnet: Toward a controllable feature space for image restoration. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2, 3
- [31] Xintao Wang, Ke Yu, Chao Dong, Xiaoou Tang, and Chen Change Loy. Deep network interpolation for continuous imagery effect transition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1692–1701, 2019. 2, 3, 5
- [32] 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 *The European Conference on Computer Vision Workshops (ECCVW)*, September 2018. 3
- [33] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015. 4
- [34] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. 2, 7
- [35] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3929–3938, 2017. 7
- [36] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. 2, 7
- [37] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Deep plug-and-play super-resolution for arbitrary blur kernels. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1671–1681, 2019. 2
- [38] 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. 6
- [39] 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 (ICCV)*, October 2019. 3