

Boosting Diffusion Guidance via Learning Degradation-Aware Models for Blind Super Resolution

Shao-Hao Lu¹ Ren Wang² Ching-Chun Huang¹ Wei-Chen Chiu¹

¹National Yang Ming Chiao Tung University ²MediaTek Inc.

{shlu2240.cs11, chingchun, walon}@nycu.edu.tw ren.wang@mediatek.com

Abstract

Recently, diffusion-based blind super-resolution (SR) methods have shown great ability to generate high-resolution images with abundant high-frequency detail, but the detail is often achieved at the expense of fidelity. Meanwhile, another line of research focusing on rectifying the reverse process of diffusion models (i.e., diffusion guidance), has demonstrated the power to generate high-fidelity results for non-blind SR. However, these methods rely on known degradation kernels, making them difficult to apply to blind SR. To address these issues, we introduce degradation-aware models that can be integrated into the diffusion guidance framework, eliminating the need to know degradation kernels. Additionally, we propose two novel techniques—input perturbation and guidance scalar—to further improve our performance. Extensive experimental results show that our proposed method has superior performance over state-of-the-art methods on blind SR benchmarks.

1. Introduction

Single image blind super-resolution (blind-SR) aims to recover a high-resolution (HR) image from a low-resolution (LR) observation which undergoes an unknown degradation process, where the loss of high-frequency detail during the unknown degradation primarily leads to the inherent challenge in the task of blind-SR. Another well-known challenge comes from having multiple potential solutions of high-resolution images (i.e., multiple HR images would yield the identical LR images), and thus making the ill-posed nature of super-resolution problem. Although the renaissance of deep learning techniques in recent decade brings the magic leap to the task of super-resolution, especially the common practice stemmed from regression-based methods [15, 18] excels in fidelity, they often suffer from insufficient high-frequency image detail. This deficiency highlights the need for generative models, where the essence lies in drawing high-frequency detail from the im-

age prior distribution captured by generative models. Previous approaches have turned to Generative Adversarial Networks (GANs) [11, 31, 38] for such purpose, yet there comes other issues such as unstable training or mode collapse.

With the recent advances in diffusion-based generative models, some works [19, 26, 29, 34–36] have utilized pretrained diffusion models to address the blind super-resolution problem for better training stability, image synthesis quality, and more flexible controls of generative procedure. Specifically, some works [19, 29, 36] propose to directly guide the diffusion sampling with LR input, while some works [34, 35] use additional prompts for guidance. Though these methods can produce high-quality images, they often fall short in terms of fidelity. Several non-blind diffusion-based image restoration methods [14, 32] have attempted to address this challenge, demonstrating the ability to produce high-fidelity results. However, these methods have limitations in that the degradation operator must be both known and linear, whereas degradation operators are typically unknown in real-world scenarios.

In this work, we adopt the Denoising Diffusion Null-Space Model (DDNM) [32], a representative non-blind diffusion-based image restoration method, as our backbone framework. In order to eliminate the need to know degradation kernels in DDNM, we introduce several modifications. First, we use a deep neural network to approximate the degradation process, eliminating the need to provide the degradation operator. This network is called degradation model. The degradation model not only addresses unknown degradation kernels, but also potentially overcomes the limitation of kernel linearity, enhancing the versatility and effectiveness of our proposed method. Second, for the inverse matrix of the degradation operator in DDNM (noting that since degradation operator in DDNM is represented as matrix, its inverse is thus related to the degradation removal), we replace it with a restoration model. During the sampling process, the restoration model aims to recover an HR image from an LR image generated by our degradation model at each timestep. Last, to improve performance, both of the degradation model and the restoration model are condi-

tioned on the same degradation representations, which are extracted from the LR input by an encoder model. As a result, we call these two models degradation-aware models.

Moreover, we propose two additional techniques to enhance the capacity and performance of our method against the prediction errors upon the degradation: 1) *input perturbation* is introduced to reduce the performance drops associated with the inaccuracies in degradation estimation, and thus contributing to model stability; 2) *guidance scalar* is introduced to facilitate better integration between the restoration model and the pretrained diffusion model as combining the restoration model with diffusion sampling guidance could alter the data distribution predefined by the diffusion model, which leads to issues such as hallucinations and artifacts. With all the above-mentioned components, we present a robust method for the task of diffusion-based blind super-resolution, and extensive experiments demonstrate that our method has the capability to generate realistic details without sacrificing the fidelity. Compared to other state-of-the-art generative model-based blind-SR methods, our approach excels in both fidelity and perceptual quality.

Our contributions are summarized as follows:

- We propose degradation-aware models, which consists of a degradation model and a restoration model. With these two models, we eliminate the need to know degradation kernels in DDNM [32], making it applicable to blind-SR.
- We propose input perturbation and guidance scalar, which contribute to make our method more tolerant towards the inaccuracies in degradation estimation.
- Extensive experiments demonstrate that our method has superior performance over other state-of-the-art generative model-based blind-SR methods, in terms of both fidelity and perceptual quality.

2. Related Works

2.1. Image Super-Resolution

Image super-resolution (SR) aims to restore a high-resolution (HR) image from its degraded low-resolution (LR) observation. Early methods [3, 8, 9, 18] usually assume a predefined degradation operator such as bicubic downsampling. These methods often fail in real-world data while the degradation assumption is not accurate. Therefore, recent works have shifted their focus towards blind-SR, where the degradation is unknown. Some methods [15, 22, 30, 37, 41] use unsupervised learning techniques to learn degradation models, while others [21, 31, 38, 40] aim to synthesize LR-HR image pairs that resemble real-world data. However, capturing diverse real-world degradations remains challenging.

2.2. Diffusion-based Image Super-Resolution

Recent advances in diffusion models [13, 24, 26, 28] have shown promise in image synthesis tasks. Many image restoration (IR) methods [5, 10, 14, 32] use diffusion models as an image prior. Some non-blind IR methods [14, 32] demonstrate the ability to produce high-fidelity results with known degradation operators. Some works [5, 10] extend non-blind IR to blind IR by estimating the degradation operator, and thus knowing the degradation operator becomes unnecessary. However, all of these methods require the degradation operator to be linear, which should be oversimplified in real-world scenarios.

Most diffusion-based blind super-resolution methods [4, 19, 29, 34–36] do not need a known or estimated degradation operator. Instead, they usually use conditioning signals to guide diffusion models during sampling. Stable-SR [29] and Diff-BIR [19] utilize only the low-resolution observation as the conditioning signal. DASR [35] and SeeSR [34] further incorporate semantic information as an extra conditioning signal. All of these methods have shown highly promising capabilities in generating realistic image detail. However, they often fall short in terms of fidelity due to the inherent diffusion hallucination. Achieving high-fidelity results remains a significant challenge for most diffusion-based blind-SR methods.

3. Preliminaries

3.1. Diffusion Models

We follow the diffusion model defined in DDPM [13]. Given a natural image $\mathbf{x}_0 \sim q(\mathbf{x})$, DDPM uses a forward Markovian diffusion process $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ to repeatedly add Gaussian noise on it until $t = T$ by

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}), \quad (1)$$

where β_t is the pre-defined variance at timestep t . Note that if T is sufficiently large (e.g., 1000), we can have $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Based on the property of Markov chain, for any intermediate timestep $t \in \{1, 2, \dots, T\}$, the corresponding noisy distribution has an analytic form:

$$\begin{aligned} q(\mathbf{x}_t|\mathbf{x}_0) &= \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}) \\ &= \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \end{aligned} \quad (2)$$

where $\bar{\alpha}_t = \prod_{i=1}^t (1 - \beta_i)$ and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The reverse process aims at yielding the previous state \mathbf{x}_{t-1} from \mathbf{x}_t using the posterior distribution:

$$p(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_{0:t}) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_t(\mathbf{x}_t, \mathbf{x}_{0:t}), \sigma_t^2\mathbf{I}), \quad (3)$$

where

$$\mu_t(\mathbf{x}_t, \mathbf{x}_{0|t}) = \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}\mathbf{x}_{0|t} + \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}\mathbf{x}_t, \quad (4)$$

$$\sigma_t^2 = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t}\beta_t, \quad (5)$$

and

$$\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t}\epsilon_\theta(\mathbf{x}_t, t)), \quad (6)$$

where $\epsilon_\theta(\mathbf{x}_t, t)$ is the estimated noise ϵ at timestep t . By iteratively sampling \mathbf{x}_{t-1} from $p(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_{0|t})$, DDPM can generate realistic image $\mathbf{x}_0 \sim q(\mathbf{x})$ from random noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

3.2. Denoising Diffusion Null-Space Model

Denoising Diffusion Null-Space Model (DDNM) [32] offers a zero-shot framework for solving a range of arbitrary linear inverse problems. This innovative approach leverages a pre-trained off-the-shelf diffusion model as a generative prior, requiring no additional training or modifications to existing networks. By focusing solely on refining the null-space contents during the reverse diffusion process, DDNM has the capability to generate diverse results that satisfy both fidelity and perceptual quality.

Given a degraded image $\mathbf{y} \in \mathbb{R}^{d \times 1}$ and a degradation operator $\mathbf{A} \in \mathbb{R}^{d \times D}$, we could have $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^{D \times 1}$ is the ground-truth image. DDNM applies this relationship to rectify the $\mathbf{x}_{0|t}$ in Eq. (6) as

$$\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A})\mathbf{x}_{0|t}, \quad (7)$$

where $\hat{\mathbf{x}}_{0|t}$ can be seen as the composition of $\mathbf{x}_{0|t}$ and \mathbf{y} through range-null space decomposition. Then, they substitute the $\mathbf{x}_{0|t}$ in Eq. (4) by this rectified version and sample \mathbf{x}_{t-1} using Eq. (3).

The strength of DDNM lies in its ability to utilize the denoising diffusion model to fill in the null-space information, ensuring that the output adheres to the notion of *Realness*. This adherence is particularly effective in scenarios where the degradation is linear and known, leading to perfect *Consistency* due to the natural properties of the Range-Null space. However, a notable limitation arises when the degradation operator \mathbf{A} must be both known and linear, which is often not aligned with the complexities of real-world applications where degradations are typically unknown.

4. Method

We follow the Denoising Diffusion Null-Space Model (DDNM) [32] framework, considering its capability to generate high-quality images while preserving high fidelity. To run the DDNM algorithm, the first step involves obtaining

the degradation operator \mathbf{A} and its corresponding pseudo inverse \mathbf{A}^\dagger . We investigate replacing the degradation operator \mathbf{A} and its corresponding pseudo-inverse \mathbf{A}^\dagger with a DNN-based degradation model and a DNN-based restoration model, both of which are in a more general form. The purpose of the degradation model is to imitate the degradation process. During sampling, the restoration model will recover an HR image from an LR image generated by the degradation model at each timestep. Furthermore, to improve performance, both of the degradation model and the restoration model will be conditioned on the same degradation representations extracted from the LR input by an encoder model. In Sec. 4.1, we give a detailed formulation for all of these models. In Sec. 4.2, we present *input perturbation*, which aims to mitigate performance drops caused by inaccuracies in degradation estimation. In Sec. 4.3, we present *guidance scalar*, which aims to facilitate better integration between the restoration model and DDNM to avoid hallucinations and artifacts.

4.1. Degradation-Aware Models

Li *et al.* [15] introduce a method to encode latent degradation representations from degraded images through the joint learning of image degradation and restoration. The process begins with employing an encoder E to encode the degraded image \mathbf{y} into the degradation representations $E(\mathbf{y})$. This representation is modeled as a multi-channel latent map, capturing the 2D spatially varying degradation in a latent space. Subsequently, the framework introduces an image degradation model G_d and an image restoration model G_r to produce the degraded image \mathbf{y}' and the restored image \mathbf{x}' , respectively. The entire framework is trained in an end-to-end manner.

In the context of our work, we replace the \mathbf{A} and \mathbf{A}^\dagger in DDNM with G_d and G_r , with the aim of removing the limitation on known and linear degradation kernels. If we substitute the \mathbf{A} and \mathbf{A}^\dagger in Eq. (7) by G_d and G_r , we can rewrite Eq. (7) as

$$\hat{\mathbf{x}}_{0|t} = \mathbf{x}_{0|t} + (G_r(\mathbf{y}) - \mathbf{x}'_{0|t}), \quad (8)$$

where $\mathbf{x}'_{0|t} = G_r(G_d(\mathbf{x}_{0|t}))$. Note that Eq. (8) can be interpreted as guiding $\mathbf{x}_{0|t}$ towards $G_r(\mathbf{y})$.

Notably, the supervision for G_d and G_r in this framework are separated, indicating the independence of these two models. To ensure the pseudo-inverse relationship between G_d and G_r , we introduce a consistency loss during the joint training of E, G_d, G_r as

$$\mathcal{L}_c = \|\mathbf{y} - G_d(G_r(\mathbf{y}), E(\mathbf{y})), E(\mathbf{y})\|_1. \quad (9)$$

This additional loss term promotes stability and accelerates convergence during training.

It is expected that replacing \mathbf{A} and \mathbf{A}^\dagger with neural networks could easily yield promising results. However, due

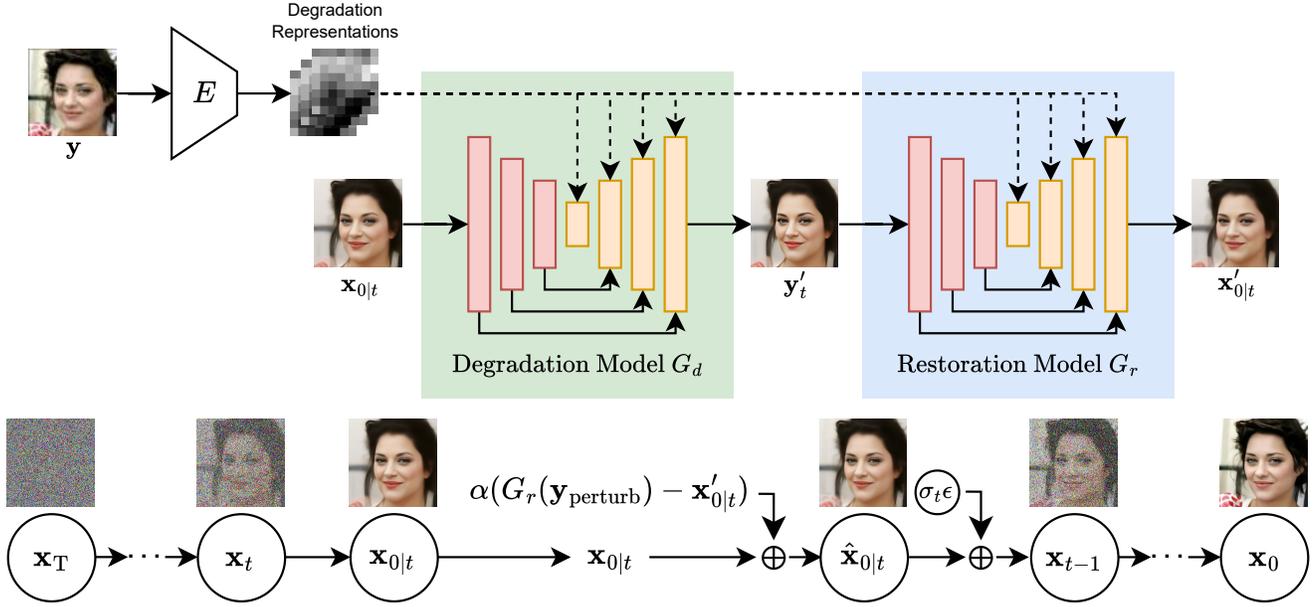


Figure 1. **Overview of our method.** At each timestep t of a reverse process, given the output of a pretrained diffusion model $\mathbf{x}_{0|t}$ and a degraded image \mathbf{y} , our degradation model G_d takes as input $\mathbf{x}_{0|t}$ and degradation representations $E(\mathbf{y})$, and produces an estimated degraded image \mathbf{y}'_t . Then, our restoration model G_r transforms \mathbf{y}'_t into a restored image $\mathbf{x}'_{0|t}$. We then rectify $\mathbf{x}_{0|t}$ to $\hat{\mathbf{x}}_{0|t}$ using $\mathbf{x}'_{0|t}$ and $G_r(\mathbf{y}_{\text{perturb}})$, where $\mathbf{y}_{\text{perturb}}$ is a perturbed degraded image, and add noise $\sigma_t \epsilon$ to obtain \mathbf{x}_{t-1} for the next timestep. Repeating this process, we get the final restored image \mathbf{x}_0 at timestep 0.

to the inherent estimation errors and the potential misalignment between G_d and G_r , the Range-Null space decomposition in DDNM may not perfectly hold. To address this issue, we present input perturbation and guidance scalar in the next sections. These refinements aim to further enhance the synergy between DDNM and the degradation-aware models G_d and G_r , and thereby address potential issues caused by degradation inconsistencies to improve the overall performance of our method.

4.2. Input Perturbation

In Eq. (8), \mathbf{y} represents the ground-truth degraded image, which can be thought of as obtained by $\mathbf{y} = G_d^{\text{gt}}(\mathbf{x})$, where G_d^{gt} is the ground-truth degradation operator and \mathbf{x} is the ground-truth high-resolution image. As a result, $G_r(\mathbf{y})$ and $G_r(G_d(\mathbf{x}_{0|t}))$ cooperate with different degradation operators, and this mismatch could lead to failure cases. Intuitively, we attempt to perform a perturbation on \mathbf{y} such that it were also obtained by applying G_d as the degradation operator.

However, directly transforming \mathbf{y} into $G_d(\mathbf{x})$ is infeasible due to the unknown ground-truth high-resolution image \mathbf{x} . To address this, we propose a perturbation approach using G_d and G_r sequentially as

$$\mathbf{y}_{\text{perturb}} = G_d(G_r(\mathbf{y})). \quad (10)$$

If we substitute the \mathbf{y} in Eq. (8) by $\mathbf{y}_{\text{perturb}}$, the degradation

operators in Eq. (8) become aligned. The choice of G_r is flexible and can be any off-the-shelf restoration model. In our method, we can simply use the G_r described in Sec. 4.1. By applying $\mathbf{y}_{\text{perturb}}$ to Eq. (8), our method is expected to become more robust.

4.3. Guidance Scalar

Note that $G_r(\mathbf{y})$ represents the output of a restoration model. We observe that applying $G_r(\mathbf{y})$ to Eq. (8) significantly accelerates the restoration process. This might bring a negative effect on fidelity as the diffusion model would take more time to hallucinate nonexistent contents. To counteract this, it becomes imperative to reduce the influence of $G_r(\mathbf{y})$ in Eq. (8), slowing down the restoration process to find an optimal balance. That is, we should be able to control the restoration strength during the diffusion process.

Recall that a characteristic of Eq. (8) is that it can be viewed as guiding $\mathbf{x}_{0|t}$ towards $G_r(\mathbf{y})$. However, in this case, the approximation error from $G_d(\mathbf{x}_{0|t})$ will make the update direction not correctly point towards $G_r(\mathbf{y})$. As a result, how to determine an appropriate weighting of this guidance is also important from this perspective.

To this end, we introduce a scalar $\alpha \in [0, 1]$ to rewrite Eq. (8) as

$$\hat{\mathbf{x}}_{0|t} = \mathbf{x}_{0|t} + \alpha(G_r(\mathbf{y}) - \mathbf{x}'_{0|t}), \quad (11)$$

Metric	MsdNet [15]	BSRGAN [38]	Real-ESRGAN [31]	LDL [16]	DASR [17]	FeMaSR [2]	DDNM [32]	LDM [26]	ResShift [36]	PASD [35]	DiffBIR [19]	SeeSR [34]	Ours
PSNR \uparrow	27.3933	22.0694	21.7337	23.1245	22.9961	20.2219	23.0214	21.7728	21.6208	22.7452	20.7647	21.8206	24.4681
SSIM \uparrow	0.7276	0.5188	0.5179	0.5355	0.5369	0.4601	0.5459	0.5108	0.5101	0.5348	0.4422	0.5009	0.6208
LPIPS \downarrow	0.3486	0.3392	0.3486	0.3498	0.3682	0.3242	0.4814	0.3238	0.3097	0.412	0.4106	0.3399	0.3069
DISTS \downarrow	0.2558	0.2634	0.2681	0.2746	0.2824	0.2573	0.3182	0.2474	0.2543	0.2855	0.2866	0.2599	0.2393
FID \downarrow	91.772	92.604	91.3774	94.2726	102.948	94.2519	106.1603	83.9125	88.3887	90.0368	101.1623	86.939	82.2803

Table 1. **Quantitative comparison of 4 \times upsampling on DIV2K-Val.** The best performance is in red while the second best is in blue.

Method	4 \times					8 \times				
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow
DDNM [32]	25.6975	0.8069	0.2074	0.1634	38.3359	20.1491	0.65	0.2469	0.2061	52.6844
SR3 [27]	26.0082	0.7392	0.1289	0.1209	19.4323	23.79	0.6995	0.1742	0.1505	25.6039
DiffBIR [19]	22.9068	0.7098	0.1769	0.1793	66.8081	19.0676	0.5621	0.264	0.2102	67.7493
SeeSR [34]	22.3977	0.7187	0.1808	0.1747	66.7649	19.2302	0.5827	0.2542	0.1974	67.6858
MsdNet [15]	28.8877	0.8317	0.2405	0.2048	59.757	25.2789	0.7527	0.3426	0.2555	76.6328
Ours	27.1513	0.7784	0.0913	0.1066	18.6689	22.8065	0.6893	0.1502	0.1399	23.9716

Table 2. **Quantitative comparison on CelebA-Val.** The best performance is in red while the second best is in blue.

where α is used to control the guidance weighting. Our experiments show that setting $\alpha = 0.3$ usually results in a suitable fidelity without introducing too many hallucinated contents.

5. Experiments

5.1. Experiment Settings

Data preparation. We evaluate our method on DIV2K unknown degradation dataset [1] and CelebA-HQ dataset [20].

To train our degradation and restoration models, we generate LR-HR paired data as described in Supplementary. For testing dataset, We use the official train, test split for DIV2K, following SeeSR [34] we randomly crop 1K HR patches (resolution: 256×256) and corresponding LR patches (resolution: 64×64) from the DIV2K unknown degradation validation set. We name this dataset as *DIV2K-Val*. For the CelebA-HQ datasets, we randomly sample 1K images as testing data and generate corresponding LR images using the same pipeline as that for generating training data, naming this dataset *CelebA-Val*. We experiment on a 4 \times degradation scale on *DIV2K-Val*, as only 4 \times degraded scale LR images are provided officially. For *CelebA-Val*, we experiment with both 4 \times and 8 \times degraded scales.

Implementation details. For the *DIV2K-Val* testing set, we use the 256×256 uncondition denoising model pretrained on ImageNet by Dhariwal *et al.* [6]. For the CelebA-Val testing set, we use 256 uncondition denoising model pretrained on CelebA release by Meng *et al.* [23]. For training the degradation and restoration models, we follow the settings of Li *et al.* [15] with the addition of a consistency loss, where the learning rate of the consistency loss is set to 0.1 of the pixel loss. During inference with the DDNM algorithm, we use DDPM as the diffusion sampling formula and adopt

the spaced DDPM sampling [25] with 100 timesteps. The guidance scalar is set to 0.3 throughout all the experiments.

Methods in comparison. We compare our method with several state-of-the-art generative model-based blind super-resolution (SR) methods. For the *DIV2K-Val* dataset, we compare our approach with GAN-based methods, including BSRGAN [38], Real-ESRGAN [31], LDL [16], FeMaSR [2] and DASR [17]. We also compare with diffusion-based methods, including LDM [26], ResShift [36], PASD [35], DiffBIR [19] and SeeSR [34]. For the *CelebA-Val* dataset, we compare with DiffBIR [19], SeeSR [34] and SR3 [27], where SR3 is trained with the same training dataset we use in this work. Since our method is strongly related to [32] and MsdNet [15], we also compare it with them. It is worth mentioning that for blind-SR tasks, the ground truth degradation kernel is unknown. Therefore, we use the default SR settings in DDNM experiments, where \mathbf{A} is set to average pooling and \mathbf{A}^\dagger is set to patch upsample. We use the publicly released codes and pretrained models (except SR3) of the competing methods for testing.

We employ several evaluation metrics to showcase our model’s capability in terms of fidelity and perceptual quality. For fidelity measures, we use PSNR and SSIM [33]. For perceptual quality measures, we use LPIPS [39] and DISTS [7]. Additionally, we use FID [12] to evaluate the distance between the distributions of original and restored images. These metrics provide a comprehensive evaluation of our model’s performance.

5.2. Comparison with State-of-the-Arts

Quantitative comparisons. We first present the quantitative comparison on *DIV2K-Val* in Tab. 1. The following observations can be made: 1) Compared to DDNM, where \mathbf{A} is set to average pooling and \mathbf{A}^\dagger is set to patch up-

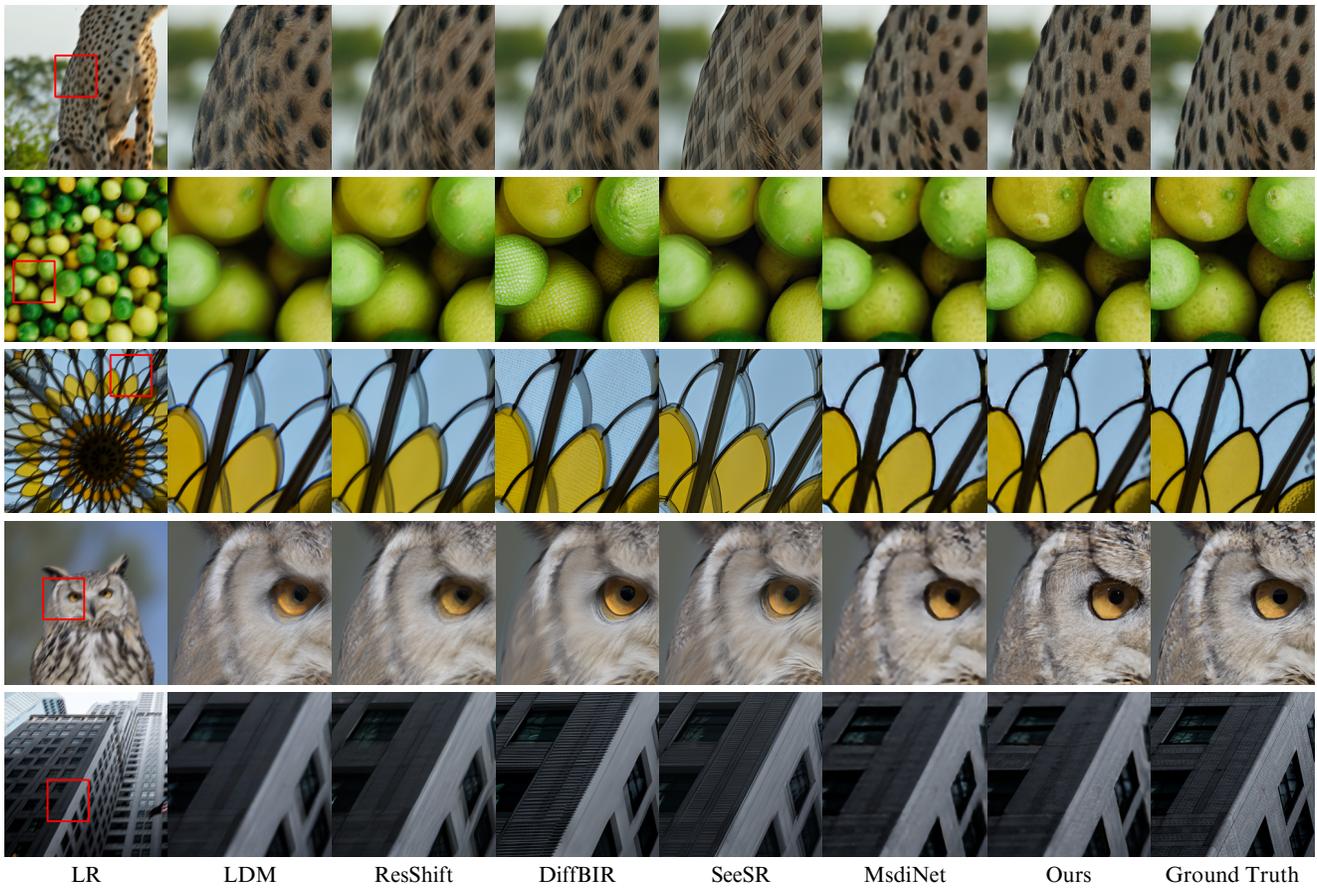


Figure 2. **Qualitative comparison of $4\times$ upsampling on *DIV2K-Val*.** The magnified areas are indicated with red boxes.

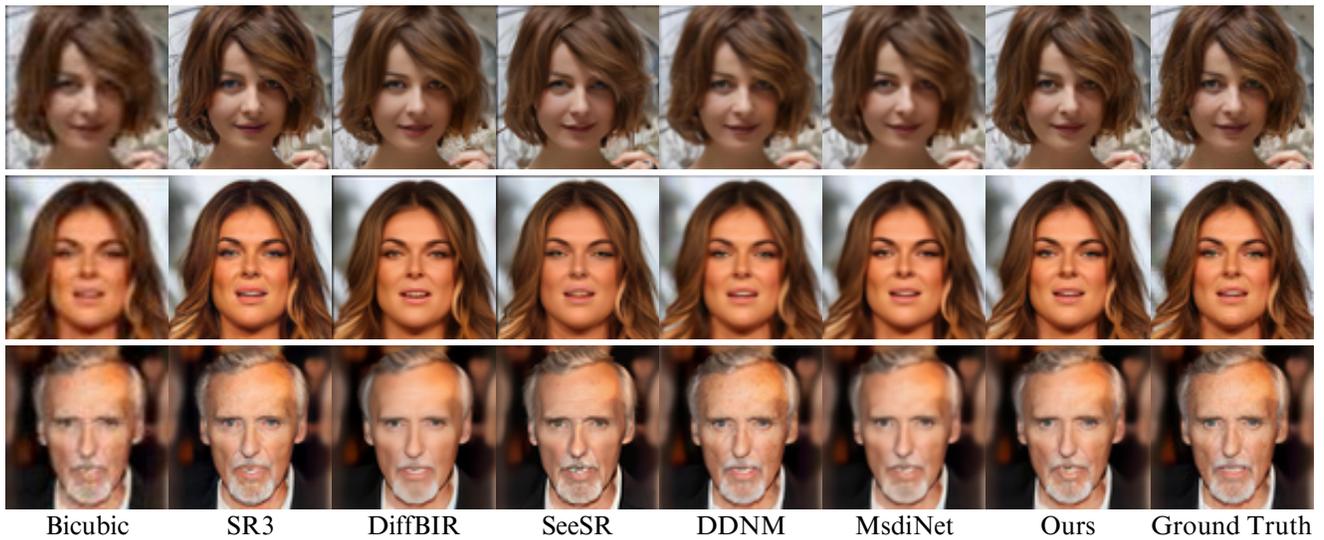


Figure 3. **Qualitative comparison of $4\times$ upsampling on *CelebA-Val*.**

sample, our method excels in all metrics, showcasing the effectiveness of G_d, G_r , input perturbation, and guidance

scalar. 2) Our method achieves the best scores on LPIPS, DISTS, and FID metrics, indicating that we effectively uti-

lize the capabilities of pretrained diffusion models to produce outputs with high perceptual quality. Additionally, we achieve the best PSNR and SSIM scores compared to other generative model-based blind-SR methods, successfully addressing the challenge of the fidelity weakness often associated with generative model-based blind-SR methods. 3) While MsdiNet has the best scores on fidelity metrics, it performs poorly on perceptual measurements. This is due to its DNN regression-based fidelity-focused learning objectives, which fall short in striking a balance between fidelity and perceptual quality.

We then present the quantitative comparison on *CelebA-Val* in Tab. 2. Our method achieves the best scores in image perceptual measurements. We believe our method strikes a better balance between fidelity and perceptual quality compared to DNN regression-based methods (MsdiNet) or other generative model-based blind-SR methods. Overall, compared with other generative-based blind-SR methods, our method achieves better results not only in fidelity but also in image perceptual measurements.

Qualitative comparisons. Fig. 2 presents qualitative comparisons between our method and other diffusion-based blind-SR methods on *DIV2K-Val* dataset. LDM and ResShift produce blurry results since they rely solely on the condition of low-resolution observations. DiffBIR and SeeSR can produce outputs with better image quality but often generate incorrect textures due to hallucinations from the diffusion prior. PASD results are blurry due to the challenge of correctly extracting semantic prompts from low-resolution patches. In comparison, our method, with the help of additional degradation and restoration models, can produce outputs that excel in both fidelity and perceptual quality.

Fig. 3 presents qualitative comparisons between our method and other methods on the *CelebA-Val* dataset. SR3 generates excessive high-frequency information. DiffBIR and SeeSR exhibit shortcomings in fidelity. DDNM, without ground truth \mathbf{A}^\dagger , and MsdiNet struggle to generate high-quality images. In contrast, our method, with the assistance of degradation and restoration models, achieves high fidelity and perceptual quality.

5.3. Ablation Studies

Effectiveness of input perturbation and guidance scalar.

To demonstrate the impact of input perturbation and the guidance scalar, we conducted detailed ablation studies on our proposed method using the *DIV2K-Val* dataset. The results are presented quantitatively in Tab. 3 and qualitatively in Fig. 5, both of which highlight the impact of each component. We compared the results of simply combining DDNM [32] with the extra degradation and restoration models against incorporating input perturbation, the guidance scalar, or both. In Tab. 3, the guidance scalar signifi-

Input Perturbation	Guidance Scalar	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
✓	✓	24.4681	0.6208	0.3069
-	✓	20.7559	0.5525	0.3544
✓	-	17.5641	0.3585	0.5147
-	-	13.5803	0.2281	0.6269

Table 3. Ablation study on input perturbation and guidance scalar.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
$\alpha = 0.1$	24.2709	0.5961	0.3892
$\alpha = 0.3$	24.4681	0.6208	0.3069
$\alpha = 0.5$	22.387	0.5648	0.341
$\alpha = 1.0$	17.5641	0.3585	0.5147

Table 4. Ablation study on different values of guidance scalar.

cantly benefits both fidelity and perceptual quality, improving PSNR by 7.2 dB and LPIPS by 0.27. The input perturbation also benefits fidelity and perceptual quality. Combining both techniques further amplifies these gains, highlighting their effectiveness in improving our method’s performance. Fig. 5 supports the same conclusion quantitatively.

Design choices of the degradation operator and its pseudo-inverse.

We conduct experiments using various approaches to model the degradation operator and its corresponding pseudo-inverse. We train a DNN-based degradation model and a DNN-based restoration model to replace the \mathbf{A} and \mathbf{A}^\dagger in the DDNM [32] algorithm, providing a more generalized approach to modeling the degradation and its corresponding pseudo-inverse. In our ablation experiments, we refer to this method as *implicit*, as the degradation representation is implicitly learned as a feature.

Another naive approach to modeling the degradation operator and its pseudo-inverse is by defining \mathbf{A} as a convolution with an explicit Gaussian kernel. This method assumes \mathbf{A} is linear and known, allowing \mathbf{A}^\dagger to be calculated using techniques like Singular Value Decomposition (SVD), similar to what DDNM proposes. However, relying on the assumption that \mathbf{A} is a convolution operator limits its expressiveness compared to methods that learn the degradation implicitly or through more complex models. To construct an experiment based on this approach, we trained an explicit kernel estimator using the same training data described in Sec. 5.1. Given a low-resolution input observation, the explicit kernel estimator predicts a corresponding explicit linear degradation kernel. We define convolution with the predicted kernel as \mathbf{A} and calculate its corresponding pseudo-inverse operator \mathbf{A}^\dagger using SVD. In our ablation study, we refer to this method as *explicit* as it explicitly defines the degradation process using a Gaussian kernel. To

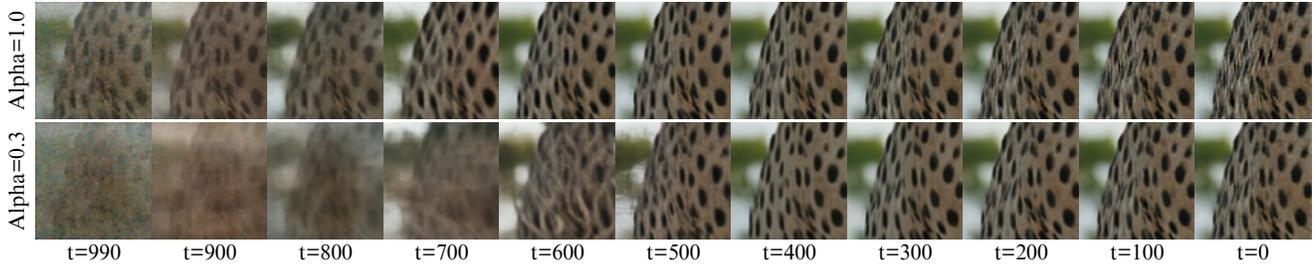


Figure 4. Ablation study on the speed of restoration process under different guidance scalar values.

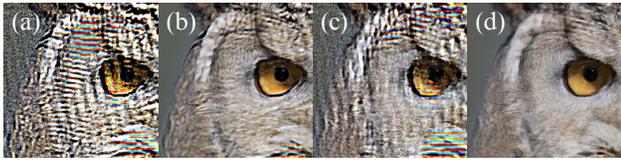


Figure 5. Ablation study on input perturbation and guidance scalar. (a) both off, (b) guidance scalar on, (c) input perturbation on, (d) both on.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
<i>implicit</i>	24.4681	0.6208	0.3069
<i>explicit</i>	20.955	0.3999	0.5985
<i>combine</i>	22.4154	0.531	0.4787

Table 5. Ablation study on various approaches to modeling degradation. All of the experiments are conducted with $4\times$ scale.

further experiment with the design choices of the \mathbf{A} and \mathbf{A}^\dagger in the DDNM [32] algorithm, we also explore combining the explicit version of \mathbf{A} with the implicit version of \mathbf{A}^\dagger . In our ablation study, we refer to this approach as *combine*. More details of the *explicit* and *combine* approaches are described in Supplementary.

We conduct the ablation study on *DIV2K-Val* dataset as shown in Tab. 5, we observe that the implicit version of \mathbf{A} and \mathbf{A}^\dagger outperforms the *combine* and *explicit* approaches. This superiority is attributed to the greater expressiveness of the implicit degradation representation compared to convolution with a Gaussian kernel. Convolution with a Gaussian kernel assumes a global degradation operator, which limits its ability to capture variations in degradation that can occur locally across the image in real-world scenarios. In contrast, our *implicit* approach can capture the 2D spatially varying degradation in a latent space.

Overall, the *implicit* setting of \mathbf{A} and \mathbf{A}^\dagger offers better expressiveness in modeling degradation, which makes our method more robust and reliable across diverse datasets and real-world scenarios.

Ablation study on the value of guidance scalar. Note that

the guidance scalar α in Eq. (11) is a critical parameter. If α is set too high, it accelerates the restoration process, leading to a hallucination problem in the pretrained diffusion model. Conversely, setting it too low reduces the influence of the guidance, leading to low-fidelity results. Note that setting it to zero equates to unconditional diffusion sampling.

We provide quantitative comparison in Tab. 4 and qualitative comparison in Fig. 4, both of which demonstrate the speed of restoration process under different scalar values. Both Tab. 4 and Fig. 4 show that by introducing the guidance scalar, we effectively slow down the restoration process, thereby mitigating the performance decrease caused by restoration acceleration.

6. Conclusions

We have proposed a novel diffusion-based blind super-resolution (SR) solution. By combining the the degradation-aware models with DDNM, we boost diffusion guidance to produce high quality images without sacrificing the fidelity. We propose the use of a more flexible and general form of the \mathbf{A} and \mathbf{A}^\dagger in DDNM, extending their roles to not only degradation but also restoration. The integration of input perturbation and guidance scalar further enhances the performance of our method, reducing the impact of degradation estimation errors and improving the synergy between our models and DDNM. Through these innovations, we demonstrate that our method excels in both fidelity and perceptual quality compared to state-of-the-art diffusion-based blind SR methods.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017. 5
- [2] Chaofeng Chen, Xinyu Shi, Yipeng Qin, Xiaoming Li, Xiaoguang Han, Tao Yang, and Shihui Guo. Real-world blind super-resolution via feature matching with implicit high-resolution priors. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1329–1338, 2022. 5

- [3] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12299–12310, 2021. 2
- [4] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. in 2021 ieee. In *CVF international conference on computer vision (ICCV)*, pages 14347–14356, 2021. 2
- [5] Hyungjin Chung, Jeongsol Kim, Sehui Kim, and Jong Chul Ye. Parallel diffusion models of operator and image for blind inverse problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6059–6069, 2023. 2
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 5
- [7] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5):2567–2581, 2020. 5
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 184–199. Springer, 2014. 2
- [9] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 2
- [10] Ben Fei, Zhaoyang Lyu, Liang Pan, Junzhe Zhang, Weidong Yang, Tianyue Luo, Bo Zhang, and Bo Dai. Generative diffusion prior for unified image restoration and enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9935–9946, 2023. 2
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 1
- [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] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [14] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022. 1, 2
- [15] Dasong Li, Yi Zhang, Ka Chun Cheung, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Learning degradation representations for image deblurring. In *European Conference on Computer Vision*, pages 736–753. Springer, 2022. 1, 2, 3, 5
- [16] Jie Liang, Hui Zeng, and Lei Zhang. Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5657–5666, 2022. 5
- [17] Jie Liang, Hui Zeng, and Lei Zhang. Efficient and degradation-adaptive network for real-world image super-resolution. In *European Conference on Computer Vision*, pages 574–591. Springer, 2022. 5
- [18] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 1, 2
- [19] Xinqi Lin, Jingwen He, Ziyang Chen, Zhaoyang Lyu, Ben Fei, Bo Dai, Wanli Ouyang, Yu Qiao, and Chao Dong. Diffbir: Towards blind image restoration with generative diffusion prior. *arXiv preprint arXiv:2308.15070*, 2023. 1, 2, 5
- [20] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 5
- [21] Fangzhou Luo, Xiaolin Wu, and Yanhui Guo. And: Adversarial neural degradation for learning blind image super-resolution. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [22] Shunta Maeda. Unpaired image super-resolution using pseudo-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 291–300, 2020. 2
- [23] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 5
- [24] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [25] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021. 5
- [26] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2, 5
- [27] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2022. 5
- [28] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2

- [29] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *arXiv preprint arXiv:2305.07015*, 2023. 1, 2
- [30] Longguang Wang, Yingqian Wang, Xiaoyu Dong, Qingyu Xu, Jungang Yang, Wei An, and Yulan Guo. Unsupervised degradation representation learning for blind super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10581–10590, 2021. 2
- [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, 2, 5
- [32] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*, 2022. 1, 2, 3, 5, 7, 8
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [34] Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 25456–25467, 2024. 1, 2, 5
- [35] Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. *arXiv preprint arXiv:2308.14469*, 2023. 1, 2, 5
- [36] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2, 5
- [37] Jiahui Zhang, Shijian Lu, Fangneng Zhan, and Yingchen Yu. Blind image super-resolution via contrastive representation learning. *arXiv preprint arXiv:2107.00708*, 2021. 2
- [38] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4791–4800, 2021. 1, 2, 5
- [39] 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
- [40] Shangchen Zhou, Kelvin Chan, Chongyi Li, and Chen Change Loy. Towards robust blind face restoration with codebook lookup transformer. *Advances in Neural Information Processing Systems*, 35:30599–30611, 2022. 2
- [41] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2