

Unlimited-Size Diffusion Restoration

Yinhuai Wang Jiwen Yu Runyi Yu Jian Zhang[†]

Peking University Shenzhen Graduate School

{yinhuai, yujiwen, ingrid_yu}@stu.pku.edu.cn, zhangjian.sz@pku.edu.cn

Abstract

Recently, using diffusion models for zero-shot image restoration (IR) has become a new hot paradigm. This type of method only needs to use the pre-trained off-the-shelf diffusion models, without any finetuning, and can directly handle various IR tasks. The upper limit of the restoration performance depends on the pre-trained diffusion models, which are in rapid evolution. However, current methods only discuss how to deal with fixed-size images, but dealing with images of arbitrary sizes is very important for practical applications. This paper focuses on how to use those diffusion-based zero-shot IR methods to deal with any size while maintaining the excellent characteristics of zero-shot. A simple way to solve arbitrary size is to divide it into fixed-size patches and solve each patch independently. But this may yield significant artifacts since it neither considers the global semantics of all patches nor the local information of adjacent patches. Inspired by the Range-Null space Decomposition, we propose the Mask-Shift Restoration to address local incoherence and propose the Hierarchical Restoration to alleviate out-of-domain issues. Our simple, parameter-free approaches can be used not only for image restoration but also for image generation of unlimited sizes, with the potential to be a general tool for diffusion models. Code: https://github.com/wyuhuai/DDNM/tree/main/hq_demo.

1. Introduction

Recent progress in diffusion models [27, 29, 10, 25, 8, 2] has enlightened a lot works in solving Image Restoration (IR) tasks [34, 4, 28, 13, 12, 17, 26, 6, 7, 5, 24, 22, 35, 18, 37]. These diffusion-based IR methods can be roughly divided into supervised [24, 22, 35, 14] and zero-shot [34, 4, 28, 13, 12, 26, 17, 6, 7, 5]. Zero-shot methods reveal a new hot paradigm since they only need to use the pre-trained off-the-shelf diffusion model, and can directly handle various IR tasks without any finetuning. In this paper, we focus on


[†] Corresponding author. This work was supported in part by Shenzhen Research Project under Grant JCYJ20220531093215035 and Grant JSGGZD20220822095800001.



(a) Input LR image (64×32)



(b) SR result (1024×512) using DDNM, by patch



(c) SR result (1024×512) using DDNM, with **MSR & HiR**

Figure 1. Example of $16\times$ Super-Resolution (SR) that brings a 64×32 Low-Resolution (LR) image into 1024×512 SR results. (b) Simply dividing the result into eight 256×256 patches and using DDNM [34] to solve them independently will get poor results, because it neither considers the global semantics nor the boundary information of adjacent patches. (c) We propose Mask-Shift Restoration (MSR) to solve the boundary artifacts and Hierarchical Restoration (HiR) to address the lack of global semantics.

zero-shot methods [34, 4, 28, 13, 12, 26, 17, 6, 7, 5] which are concise, flexible, and in rapid progress.

Existing diffusion-based IR methods mainly focus on IR problems with fixed output sizes. But in real-world applications, the desired output size may be arbitrary, depending on the user’s demands. There are two main difficulties in applying these zero-shot IR methods to arbitrary output size: (1) The used diffusion models are usually pre-trained on fixed-size images, thus facing out-of-domain (OOD) issues when extending to arbitrary sizes; (2) The default network structure may not support arbitrary output size. The OOD issue can be solved by training the diffusion models with random cropped images. But the network structure constraint is hard to address. A common practice to bypass this constraint is to divide the input image into fixed-size patches and use the network to process each patch independently, then, concatenate the result patches as the final result, as shown in the middle of Fig. 1. However, this may lead to evident block artifacts and unreasonable restoration, because it neither considers the global semantics of all patches nor the local information of adjacent patches.

We observe that the neighboring correlation is well considered in inpainting tasks in DDNM [34], which inspired us to leave overlapped regions when dividing patches, then take the overlapped region as extra mask constraints when solving the following patches. We name this method Mask-Shift Restoration (MSR), which assures the coherence between patches and effectively eliminates boundary artifacts.

To further alleviate the OOD problem, we propose to first restore the result at a small size, then use the small result as a global prior for the final result. We name this method Hierarchical Restoration (HiR). Note that both MSR and HiR perfectly fit the zero-shot properties, and can be flexibly combined. The bottom of Fig. 1 shows the result using both MSR and HiR based on DDNM. From the perspective of Range-Null space Decomposition (RND), MSR and HiR are essentially adding extra linear constraints to the given inverse problem. This property makes it perfectly suitable for DDNM, which is exactly built on the principle of RND.

Our contribution includes:

1. We propose Mask-Shift Restoration (MSR), a simple but effective method to eliminate boundary artifacts when processing a large image in patches.
2. We propose Hierarchical Restoration (HiR) to alleviate the out-of-domain problem and the lack of global semantics when processing a large image in patches.
3. We provide typical pipelines for using MSR and HiR for diverse applications, including but not limited to image generation, super-resolution, colorization, inpainting, and denoising. It is worth noting that our proposed methods are parameter-free and training-free, and can be applied to diverse diffusion models and zero-shot restoration methods.

2. Preliminaries

2.1. Diffusion Models

Diffusion models have diverse interpretations [29, 25, 2, 16, 15], but in this paper, we put aside the mathematical meaning and introduce the diffusion model in the most concise and general way. Diffusion models [27, 29, 10, 25, 8, 2] define a T -step forward process and a T -step reverse process. The forward process adds random noise to data, while the reverse process constructs desired data samples from the noise. Specifically, the forward process yields a noisy image \mathbf{x}_t from a clean image \mathbf{x}_0 :

$$\mathbf{x}_t = a_t \mathbf{x}_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (1)$$

where $t \sim \{0, \dots, T\}$, a_t and σ_t are predefined scale factors, \mathcal{N} represents the Gaussian distribution.

The core of the reverse process is estimating the clean image \mathbf{x}_0 from the noisy image \mathbf{x}_t :

$$\mathbf{x}_{0|t} = \frac{1}{a_t} (\mathbf{x}_t - \sigma_t \epsilon_t) \quad (2)$$

which is a reverse of Eq. 1, with ϵ_t denotes the estimation of noise ϵ and $\mathbf{x}_{0|t}$ represents the estimation of \mathbf{x}_0 at time step t . Typically, a denoiser \mathcal{Z}_θ is used to yield ϵ_t :

$$\epsilon_t = \mathcal{Z}_\theta(\mathbf{x}_t, t) \quad (3)$$

Then we can use Eq. 1 to generate the previous state \mathbf{x}_{t-1} , with $\mathbf{x}_{0|t}$ as the estimation of \mathbf{x}_0 :

$$\mathbf{x}_{t-1} = a_{t-1} \mathbf{x}_{0|t} + \sigma_{t-1} \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (4)$$

With the above formulations, one can generate a clean image \mathbf{x}_0 from a random noise $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ by iterating Eq. 2 and Eq. 4 while decreasing t from T to 0.

Such a reverse process is the simplest form. Further, for Eq. 4, we can interpolate the newly added noise ϵ with the estimated previous noise ϵ_t under the premise of invariant total variance:

$$\mathbf{x}_{t-1} = a_{t-1} \mathbf{x}_{0|t} + \sigma_{t-1} (\eta_t \epsilon + \sqrt{1 - \eta_t^2} \epsilon_t), \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (5)$$

where η_t is an interpolation factor that controls the ratio of the newly introduced noise ϵ . Note that Eq. 5 describes a general form of reverse sampling methods. The critical difference between different sampling methods is the setting of η_t . For DDIM [25], η_t is a time-independent scalar; For DDPM [10] and Analytic-DPM [2], η_t is a time-dependent function.

To train the denoiser \mathcal{Z}_θ , one can randomly pick a clean image \mathbf{x}_0 from the dataset and pick a random time-step t to yield a noisy image \mathbf{x}_t using Eq. 1. Then, update the network parameters θ with the following gradient descent step [10], and repeat the whole process until converged.

$$\nabla_\theta \|\epsilon - \mathcal{Z}_\theta(\mathbf{x}_t, t)\|_2^2. \quad (6)$$

2.2. Denoising Diffusion Null-space Model (DDNM)

Recent progress shows that pre-trained diffusion models can be used to solve linear inverse problems in a zero-shot manner [34, 17, 4, 28, 12], without extra training or optimization. DDNM [34] explains the nature of such methods.

DDNM starts with noise-free linear image inverse problems. Given a degraded image $\mathbf{y} = \mathbf{A}\mathbf{x}$ where \mathbf{A} is a linear operator and \mathbf{x} is the original image, image restoration aims at yielding a result $\hat{\mathbf{x}}$ that satisfies two constraints:

$$\text{Consistency : } \mathbf{A}\hat{\mathbf{x}} \equiv \mathbf{y}, \quad \text{Realness : } \hat{\mathbf{x}} \sim q(\mathbf{x}), \quad (7)$$

where $q(\mathbf{x})$ denotes the distribution of the GT images.

Such a problem has a general solution that analytically satisfies the *Consistency* constraint:

$$\hat{\mathbf{x}} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_r. \quad (8)$$

where \mathbf{A}^\dagger is the pseudo-inverse of \mathbf{A} (satisfies $\mathbf{A}\mathbf{A}^\dagger \mathbf{A} \equiv \mathbf{A}$), and \mathbf{x}_r is the unknown null-space variable to be solved. Note that Eq. 8 originates from the Range-Null space Decomposition [34, 32, 3]. Another interpretation is that $\mathbf{A}^\dagger \mathbf{y}$ can be seen as a special solution of $\mathbf{A}\mathbf{x} = \mathbf{y}$ since $\mathbf{A}\mathbf{A}^\dagger \mathbf{y} \equiv \mathbf{A}\mathbf{A}^\dagger \mathbf{A}\mathbf{x} \equiv \mathbf{A}\mathbf{x} \equiv \mathbf{y}$; and $(\mathbf{I} - \mathbf{A}^\dagger \mathbf{A})\mathbf{x}_r$ can be seen as a general solution of $\mathbf{A}\mathbf{x} = \mathbf{0}$ since $\mathbf{A}(\mathbf{I} - \mathbf{A}^\dagger \mathbf{A})\mathbf{x}_r \equiv (\mathbf{A} - \mathbf{A})\mathbf{x}_r \equiv \mathbf{0}$ holds whatever \mathbf{x}_r is.

To conclude, Eq. 8 defined a solution that analytically satisfies the Consistency constraint but needs to find proper null-space variable \mathbf{x}_r to meet the Realness constraint. As we will get to later, the methods proposed in this paper heavily rely on the use of Eq. 8.

In DDNM [34], the critical step using diffusion models for inverse problems is taking each estimation $\mathbf{x}_{0|t}$ as the null-space variable \mathbf{x}_r in Eq. 8:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_{0|t}. \quad (9)$$

then use this consistent result $\hat{\mathbf{x}}_{0|t}$ for subsequent sampling:

$$\mathbf{x}_{t-1} = a_{t-1} \hat{\mathbf{x}}_{0|t} + \sigma_{t-1} (\eta_t \boldsymbol{\epsilon} + \sqrt{1 - \eta_t^2} \boldsymbol{\epsilon}_t), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}) \quad (10)$$

Algo. 1 shows the whole process of DDNM. See Appendix for DDNM with noisy situations.

Algorithm 1 Sampling process of DDNM

- 1: $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\boldsymbol{\epsilon}_t = \mathcal{Z}_\theta(\mathbf{x}_t, t)$
 - 4: $\mathbf{x}_{0|t} = \frac{1}{a_t} (\mathbf{x}_t - \sigma_t \boldsymbol{\epsilon}_t)$
 - 5: $\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_{0|t}$
 - 6: $\mathbf{x}_{t-1} = a_{t-1} \hat{\mathbf{x}}_{0|t} + \sigma_{t-1} (\eta_t \boldsymbol{\epsilon} + \sqrt{1 - \eta_t^2} \boldsymbol{\epsilon}_t), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$
 - 7: **return** \mathbf{x}_0
-

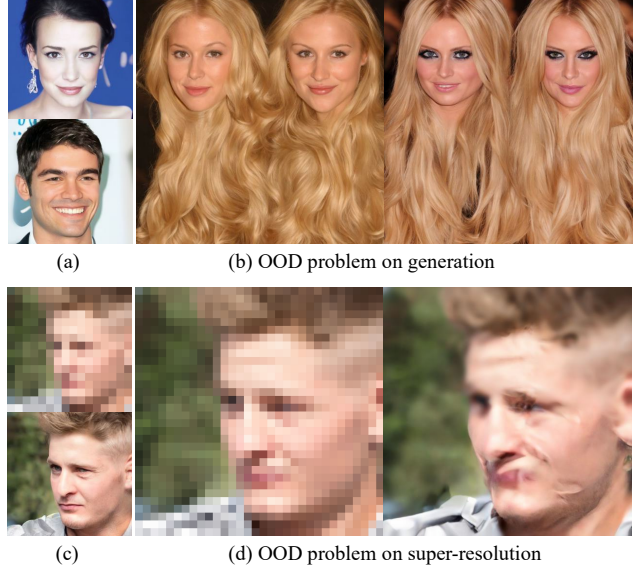


Figure 2. Out-Of-Domain (OOD) problem. (a) 256×256 images generated by diffusion model trained on aligned 256×256 CelebA dataset. (b) 512×512 images generated by the same diffusion model. We can see that the model can not generate bigger faces even enforce to generate a 512×512 image. (c) Applying the same diffusion model to DDNM [34] for $16 \times$ SR task yields good results of size 256×256 . (d) Applying the same diffusion model to DDNM for $16 \times$ SR task yields terrible results of size 512×512 . This is caused by the OOD problem.

Algorithm 2 Mask-Shift Restoration, based on DDNM

Additional Requirement: The already restored region $\hat{\mathbf{x}}_0$ and the corresponding mask \mathbf{A}_m .

- 1: $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\boldsymbol{\epsilon}_t = \mathcal{Z}_\theta(\mathbf{x}_t, t)$
 - 4: $\mathbf{x}_{0|t} = \frac{1}{a_t} (\mathbf{x}_t - \sigma_t \boldsymbol{\epsilon}_t)$
 - 5: $\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_{0|t}$
 - 6: $\tilde{\mathbf{x}}_{0|t} = \mathbf{A}_m \hat{\mathbf{x}}_0 + (\mathbf{I} - \mathbf{A}_m) \hat{\mathbf{x}}_{0|t}$
 - 7: $\mathbf{x}_{t-1} = a_{t-1} \tilde{\mathbf{x}}_{0|t} + \sigma_{t-1} (\eta_t \boldsymbol{\epsilon} + \sqrt{1 - \eta_t^2} \boldsymbol{\epsilon}_t), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$
 - 8: **return** \mathbf{x}_0
-

Algorithm 3 Hierarchical Restoration, based on DDNM

Additional Requirement: The low-resolution result $\tilde{\mathbf{x}}_0$ and the corresponding downsampler \mathbf{A}_{sr} and its pseudo-inverse \mathbf{A}_{sr}^\dagger . The already restored region $\hat{\mathbf{x}}_0$ and the corresponding mask \mathbf{A}_m .

- 1: $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\boldsymbol{\epsilon}_t = \mathcal{Z}_\theta(\mathbf{x}_t, t)$
 - 4: $\mathbf{x}_{0|t} = \frac{1}{a_t} (\mathbf{x}_t - \sigma_t \boldsymbol{\epsilon}_t)$
 - 5: $\tilde{\mathbf{x}}_{0|t} = \mathbf{A}_{sr}^\dagger \tilde{\mathbf{x}}_0 + (\mathbf{I} - \mathbf{A}_{sr}^\dagger \mathbf{A}_{sr}) \mathbf{x}_{0|t}$
 - 6: $\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \tilde{\mathbf{x}}_{0|t}$
 - 7: $\tilde{\mathbf{x}}_{0|t} = \mathbf{A}_m \hat{\mathbf{x}}_0 + (\mathbf{I} - \mathbf{A}_m) \hat{\mathbf{x}}_{0|t}$
 - 8: $\mathbf{x}_{t-1} = a_{t-1} \tilde{\mathbf{x}}_{0|t} + \sigma_{t-1} (\eta_t \boldsymbol{\epsilon} + \sqrt{1 - \eta_t^2} \boldsymbol{\epsilon}_t), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$
 - 9: **return** \mathbf{x}_0
-



Figure 3. Example of Mask-Shift Restoration for $4\times$ SR. Given an input LR image (a) with a non-square size, we first use DDNM to SR the first square patch and update the result (Step 1). Then, we shift the patch right, leaving some overlapped regions with the previous patch. Since the overlapped region is already restored, we set them as fixed and only solve the rest region (Step 2). To this end, we need an extra inpainting (mask) constraint, which is perfectly suitable for zero-shot methods like DDNM to handle. **Zoom in for the best view**

3. Method

We have introduced the basic principles of the diffusion model and DDNM. We can see that the limitation of the image processing size lies in the denoiser. Usually, the denoiser is pre-trained on fixed-size images. How do we use such pre-trained denoisers for unlimited-size image restoration? In the following part, we propose two methods to achieve this goal, both inherit the zero-shot property.

3.1. Process as a Whole Image

Typical diffusion models [27, 29, 10, 25, 8, 2] use U-Net structures [21] as the denoiser backbone. Theoretically, U-Net is a convolutional network and thus supports scalable input size.

Hence a simple solution is to directly change the model processing size. A similar approach has been widely adopted by Stable Diffusion [20] for flexible generated size. Despite supporting flexible input size, the denoiser trained on fixed image size may face Out-Of-Domain (OOD) problem when applied to other image sizes. As shown in Fig. 2, a diffusion model trained on CelebA 256×256 fails to generate desired 512×512 face images. One way to solve the OOD issue is to train the 256×256 denoiser with a random cropped dataset, rather than an aligned one. Interestingly, ImageNet and LAION-5B happen to be non-aligned datasets, and hence suffer relatively minor OOD issues.

3.2. Process as Patches

Directly changing the model processing size may work, but it still has the following limitations: (1) It may yield bad results when facing OOD problems, as shown in Fig. 2(b). (2) It still has limitations on image size, e.g., divisible by 32; (3) Large sizes, e.g., 1024×1024 , may cause unaffordable memory consumption; (4) The classifier guidance [8] can not be applied since it is usually designed for fixed input sizes; (5) Other potential network backbones [19] may not support flexible processing size.

How to use diffusion models with fixed processing sizes to solve arbitrary image sizes? A simple solution is dividing the input image \mathbf{y} into patches, solving each patch indepen-

dently, then concatenating the results. But this may cause evident boundary artifacts, as shown in the middle of Fig. 1. This is because each patch is solved independently and their connection is not considered.

3.3. Mask-Shift Restoration

Among the many image restoration tasks, inpainting is the typical one that considers the connection between the masked and unmasked region. Zero-shot methods like DDNM [34] and RePaint [17] show good performance in solving inpainting.

Our insight is that we can leave overlapped regions when dividing patches, then take these overlapped regions as an extra constraint when solving the following patches. The neat thing is that this constraint can be integrated into existing zero-shot methods [34, 4, 28, 13, 12, 26, 17, 6, 7, 5], with just one extra line of code!

Let's take a $4\times$ SR task for example, as shown in Fig. 3. Given an input image \mathbf{y}^{full} with size 64×96 , our aim is to get an SR result with size 256×384 . Here we set the degradation operator \mathbf{A} as the average-pooling downsampler, and its pseudo-inverse \mathbf{A}^\dagger as the replication upsampler [32]. Fig. 3(a) shows the result of $\mathbf{A}^\dagger\mathbf{y}^{full}$. We first divide $\mathbf{A}^\dagger\mathbf{y}^{full}$ into two square patches $\mathbf{A}^\dagger\hat{\mathbf{y}}$ and $\mathbf{A}^\dagger\mathbf{y}$ of size 256×256 . Note that $\mathbf{A}^\dagger\hat{\mathbf{y}}$ and $\mathbf{A}^\dagger\mathbf{y}$ has an overlap of size 256×128 .

We first use default DDNM to process $\mathbf{A}^\dagger\hat{\mathbf{y}}$ and get the SR result $\hat{\mathbf{x}}_0$ (Step 1 in Fig. 3). Note that $\mathbf{A}^\dagger\hat{\mathbf{y}}$ and $\mathbf{A}^\dagger\mathbf{y}$ has an overlap of size 256×128 , and this overlapped region is already restored in $\hat{\mathbf{x}}_0$. So when we use DDNM to solve $\mathbf{A}^\dagger\mathbf{y}$, we can take the restored overlapped region as a known part in an inpainting setting (Step 2 in Fig. 3). Specifically, we insert an extra inpainting constraint behind Eq. 9 in DDNM:

$$\bar{\mathbf{x}}_{0|t} = \mathbf{A}_m\hat{\mathbf{x}}_0 + (\mathbf{I} - \mathbf{A}_m)\hat{\mathbf{x}}_{0|t}. \quad (11)$$

where \mathbf{A}_m denotes the mask operator for overlapped region between $\mathbf{A}^\dagger\hat{\mathbf{y}}$ and $\mathbf{A}^\dagger\mathbf{y}$. The whole algorithm is summarized in Algo. 2, named as Mask-Shift Restoration (MSR).

As we can see from Fig. 3(c), the final result concatenated by the results of Step 1 and Step 2 does not show



Figure 4. Comparison on large scale inpainting. (b) and (c) yields unreasonable results since the patch is too small to capture global semantic information. In contrast, (c) yields a decent result due to the use of Hierarchical Restoration (HiR). **Zoom in for the best view.**




Figure 5. Example of Hierarchical Restoration for Inpainting. (a) We first do a $2\times$ downsampling for Fig. 4(a) and use MSR to restore a small result. (b) Then we use this small result as extra low-frequency guidance, and use MSR at the original size to yield the final result.

boundary artifacts. Similarly, we can iteratively use MSR to generate an unlimited-size image without boundary artifacts. Note that the overlapped region and the shifted direction can be arbitrary, and the supported task is also not limited to SR, but to all linear inverse problems.

3.4. Hierarchical Restoration

Though MSR assures local coherence, it owns a small receptive field when dealing with a large image. This may lead to a lack of grasp of global information, resulting in poor semantic information recovery. In Fig. 4(a) we show a masked image of size 512×768 , where any 256×256 patch can not cover the whole semantic subject. Fig. 4(b) shows the result using MSR based on DDNM. Though with good local coherence, it yields unreasonable semantic structures.

To extend the receptive field for better semantic restoration, we propose Hierarchical Restoration (HiR). HiR consists of two phases: a semantic restoration phase and a texture restoration phase.

Take Fig. 4(a) for example. For the semantic restoration phase, we first undergo a $2\times$ downsample to convert the 512×768 input into a 256×384 one, where a 256×256

patch can cover the whole semantic subject. Then we use MSR based on DDNM to get a 256×384 inpainting result \tilde{x}_0 , as shown in Fig. 5(a). This result is semantically reasonable and can be used as a low-frequency reference. For the texture restoration phase (Fig. 5(b)), we add an extra low-frequency constraint before Eq. 9:

$$\tilde{x}_{0|t} = \mathbf{A}_{\text{sr}}^\dagger \tilde{x}_0 + (\mathbf{I} - \mathbf{A}_{\text{sr}}^\dagger \mathbf{A}_{\text{sr}}) \mathbf{x}_{0|t}. \quad (12)$$

where \mathbf{A}_{sr} and $\mathbf{A}_{\text{sr}}^\dagger$ represent the average-pooling downsampler and its pseudo-inverse upsampler [32], respectively. Algo. 3 shows the whole algorithm of the second phase of HiR.

As we can see from Fig. 4(d), the use of HiR significantly improves semantic correctness. Note that the HiR is not limited to inpainting tasks, but is also useful for large-scale SR (Fig. 1(c)) and colorization (Fig. 7), etc.

3.5. Flexible Pipeline for Applications

Mask-Shift Restoration (MSR) can be seen as a general patch connection technology, and Hierarchical Restoration (HiR) can be seen as a general method to improve restora-

tion quality. The essence of both MSR and HiR is to determine part of the information via prior knowledge to narrow the solution space. In this paper, we implement MSR and HiR via the Range-Null space Decomposition, which is concise, effective, and mathematically elegant. Besides, there remain other possible ways to implement MSR and HiR, e.g., adding extra loss into optimization-based methods such as DPS. Hence the proposed MSR and HiR can be also used for other diffusion-based zero-shot IR methods, e.g., ILVR[4], RePaint[17], and DPS[5].

4. Experiment

In this section, we describe the configuration of the experiment in detail. All experiments use the denoiser pre-trained on ImageNet 256×256 , provided by guided-diffusion [8]. We use the classifier guidance [8] for sampling. Besides, the time-travel sampling [34, 17] is also used to improve the generative quality.

Given a desired result size, we divide it into patches from left to right, top to bottom. Each patch has a size 256×256 and has overlaps of 128 pixels with its neighbor patch, except for the boundary case. We solve the first patch using the original DDNM and solve the following patches in sequence (left to right, top to bottom) using MSR based on DDNM. Fig. 3 shows the results on $4 \times$ SR, with $T = 100$, time-travel length [34] $l = 10$, repeat times $r = 3$. In Fig. 6, we present qualitative comparisons between BSRGAN [36] and MSR-based DDNM. We experiment on $4 \times$ SR and noisy $4 \times$ SR of different sizes, where MSR-based DDNM uses $T = 250$, $l = 10$, and $r = 3$. For Fig. 1(c), Fig. 4(d), and Fig. 7 we use HiR based on DDNM.

5. Related Work

Range-Null space Decomposition (RND) [30] is a concept in linear algebra. When applied to linear inverse problems, RND explicitly defines the upper limit of recoverable information. Chen et al. [3] introduce RND into image inverse problems, and propose learning the range and null space respectively. Wang et al. [32] propose using GAN Prior to learn the Null-space and propose using average-pooling and its pseudo-inverse as a general tool for SR tasks. In DDNM [34], the authors propose using diffusion sampling to learn the Null-space and propose several practical operators for diverse applications.

Diffusion-based Zero-Shot Image Restoration Methods can be roughly divided into RND-based [34, 4, 28, 13, 12, 26, 17] and optimization-based [6, 7, 5]. The essence of these two branches lies in modifying only the sampling process while keeping the network unchanged. Specifically, they modify the intermediate image $\mathbf{x}_{0|t}$ or its noisy version \mathbf{x}_t . For a given input and a certain degradation operator, RND-based methods use RND to explicitly assure the data




Figure 6. Experiment on noisy $4 \times$ SR. Compared with BSRGAN [36], a supervised IR method, we can see that our method performs better in both realism and consistency. Due to the use of RND [34, 32, 3], our method can faithfully inherit the correct color and structure information in LR, while BSRGAN [36] fails (see the results of butterfly). **Zoom in for the best view**

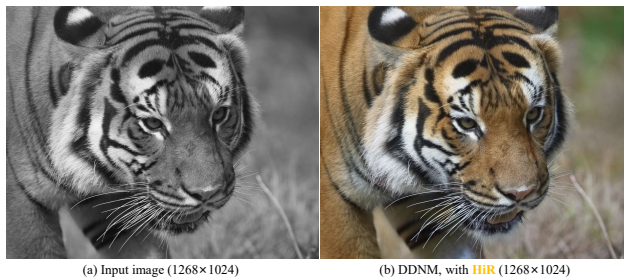


Figure 7. Colorization using HiR. **Zoom in for the best view**

consistency of $\mathbf{x}_{0|t}$ or \mathbf{x}_t , while optimization-based methods optimize $\mathbf{x}_{0|t}$ or \mathbf{x}_t toward the data consistency. Generally speaking, the RND-based methods perform better in linear inverse problems but can not solve non-linear problems. The optimization-based methods cost more on memory and inference time but can support any differentiable operator, even as a complex network [1].

6. Limitations & Discussions

Zero-shot IR methods [34, 4, 28, 13, 12, 26, 17, 6, 7, 5] using diffusion models certainly open up a promising new direction for IR problems. The method proposed in this paper further enables those methods to support unlimited image size. However, there remain some limitations to be

solved. Firstly, the calculation and time consumption are significantly more than those prevailing supervised methods. Secondly, the ceiling of performance depends on the pre-trained diffusion models. It may yield more interesting applications if applying our method to models like Imagen [23], but they are not open-sourced yet. On the other hand, widely used models like Stable Diffusion [20] are based on latent space, which makes it difficult to apply zero-shot methods. Thirdly, the degradation operator is explicitly needed, which makes it difficult for tasks like rain and haze removal.

Another interesting observation is that MSR can be seen as a general image connection method, where we can use different models to restore special crops, e.g., use face restoration models [9, 31, 32, 33, 11] for face crops, then fuse them with the background using MSR to avoid boundary artifacts.

References

- [1] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. *arXiv preprint arXiv:2302.07121*, 2023.
- [2] Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. In *International Conference on Learning Representations (ICLR)*, 2022.
- [3] Dongdong Chen and Mike E Davies. Deep decomposition learning for inverse imaging problems. In *European Conference on Computer Vision (ECCV)*. Springer, 2020.
- [4] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [5] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *International Conference on Learning Representations (ICLR)*, 2023.
- [6] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [7] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- [9] Yuchao Gu, Xintao Wang, Liangbin Xie, Chao Dong, Gen Li, Ying Shan, and Ming-Ming Cheng. Vqfr: Blind face restoration with vector-quantized dictionary and parallel decoder. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVIII*, pages 126–143. Springer, 2022.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 2020.
- [11] Yujie Hu, Yinhuai Wang, and Jian Zhang. Dear-gan: Degradation-aware face restoration with gan prior. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.
- [12] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [13] Bahjat Kawar, Jiaming Song, Stefano Ermon, and Michael Elad. Jpeg artifact correction using denoising diffusion restoration models. In *Neural Information Processing Systems (NeurIPS) Workshop on Score-Based Methods*, 2022.
- [14] Guan-Hong Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. I²sb: Image-to-image schrodinger bridge. *arXiv preprint arXiv:2302.05872*, 2023.
- [15] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *International Conference on Learning Representations (ICLR)*, 2023.
- [16] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [17] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [18] Cindy M Nguyen, Eric R Chan, Alexander W Bergman, and Gordon Wetzstein. Diffusion in the dark: A diffusion model for low-light text recognition. *arXiv preprint arXiv:2303.04291*, 2023.
- [19] William Peebles and Saining Xie. Scalable diffusion models with transformers. *arXiv e-prints*, pages arXiv–2212, 2022.
- [20] 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 (CVPR)*, pages 10684–10695, 2022.
- [21] 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 (MICCAI)*, pages 234–241. Springer, 2015.
- [22] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022.
- [23] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image

- diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- [24] 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*, 2022.
- [25] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021.
- [26] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations (ICLR)*, 2023.
- [27] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- [28] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations (ICLR)*, 2021.
- [29] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2020.
- [30] Marco Taboga. Range null-space decomposition. In *Lectures on matrix algebra*. <https://www.statlect.com/matrix-algebra/range-null-space-decomposition.>, 2021.
- [31] 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 (CVPR)*, 2021.
- [32] Yinhuai Wang, Yujie Hu, Jiwen Yu, and Jian Zhang. Gan prior based null-space learning for consistent super-resolution. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
- [33] Yinhuai Wang, Yujie Hu, and Jian Zhang. Panini-net: Gan prior based degradation-aware feature interpolation for face restoration. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2022.
- [34] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *International Conference on Learning Representations (ICLR)*, 2023.
- [35] Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alexandros G Dimakis, and Peyman Milanfar. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [36] 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 (ICCV)*, 2021.
- [37] Zixiang Zhao, Haowen Bai, Yuanzhi Zhu, Jianshe Zhang, Shuang Xu, Yulun Zhang, Kai Zhang, Deyu Meng, Radu Timofte, and Luc Van Gool. Ddfm: Denoising diffusion model for multi-modality image fusion. *arXiv preprint arXiv:2303.06840*, 2023.