

Low-Rank Residual Diffusion Models

Junfu Tan
Georgia Institute of Technology
jtan370@gatech.edu

Jiang Yuan[†]
YunQue AGI
1677532160yuan@gmail.com

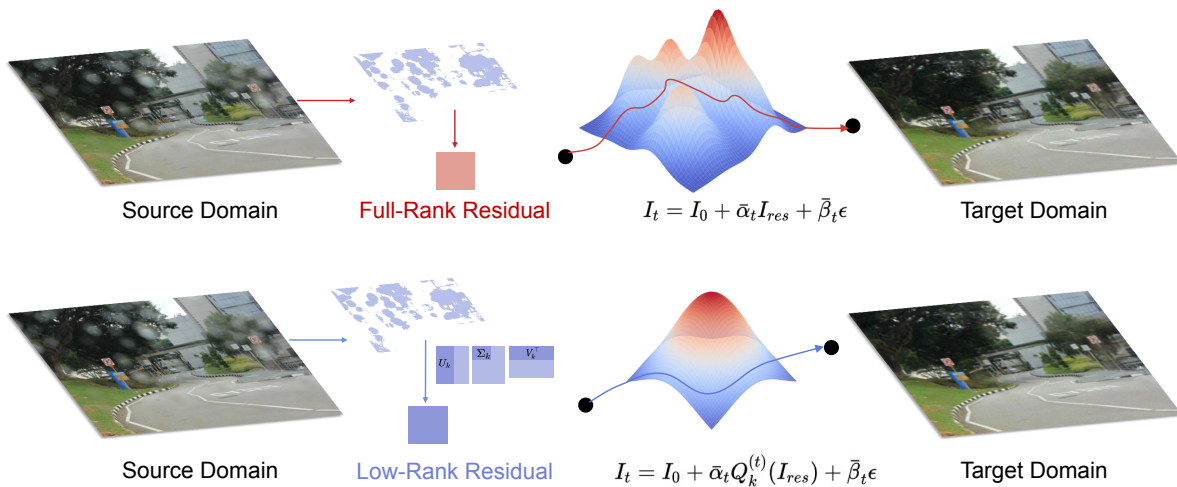


Figure 1. Residual Denoising Diffusion Models (RDDM) are diffusion models that learn the residual mapping between degraded source domains and clean target domains. In contrast, our method constrains the residual within a low-rank subspace, yielding more compact semantic representations and significantly more efficient sampling.

Abstract

Residual diffusion models have achieved remarkable progress in image restoration tasks. In near-domain transformations such as image deraining, however, we observe that the residuals exhibit an inherent low-rank structure. Motivated by this property, we propose the Low-Rank Residual Diffusion Model (LRDM), which performs diffusion within a compact low-rank residual subspace for efficient and structure-preserving restoration. We formalize this observation as the Low-Rank Residual Assumption and show that the variational lower bound becomes tighter when residuals lie in a low-rank space. Building on this insight, we introduce an Asymmetric Residual Diffusion Process that constrains the forward process in the low-rank domain while maintaining full-rank flexibility in the reverse process. To accommodate the varying complexity of residuals across diffusion timesteps, we further introduce an Adaptive Rank Selection mechanism that dynamically adjusts the rank during the diffusion process. Ex-

periments on deraining, deblurring, and deshading benchmarks show that LRDM surpasses full-rank diffusion baselines and achieves state-of-the-art performance, validating the advantage of modeling diffusion in a low-rank residual space. **Project Page:** <https://github.com/JF-Tan/LRDM>.

1. Introduction

Diffusion models [17, 39, 41–43] have emerged as a cornerstone of modern generative modeling, achieving state-of-the-art results in diverse domains such as text-to-image synthesis [37], image editing [2, 32], and computational photography [25]. Their iterative denoising process enables the generation of photorealistic images with unprecedented diversity and fidelity.

A particularly active line of research leverages conditional diffusion models for image-to-image (I2I) transformations [18, 58], especially for image restoration tasks. These tasks—such as deraining [8, 14], deblurring [33, 48], and deshading [20, 40]—aim to reconstruct clean images from degraded inputs. We denote the pixel-wise difference

[†]Jiang Yuan contributed equally to this work.

between the degraded domain and the clean domain as the *residual*. Recent advances have proposed various strategies to improve diffusion-based restoration. RDDM [27] introduces a residual-learning formulation within the diffusion process, while DiffUIR [56] adopts a selective hourglass pathway to handle multiple degradations, and DeblurDiff [21] predicts latent blur kernels to guide deblurring. However, in many practical restoration scenarios, the source and target domains are already highly similar in pixel space. We refer to such settings as *near-domain restoration* tasks, where the underlying residuals are intrinsically sparse and lie in a low-dimensional subspace. When diffusion is performed directly in this sparse residual space using standard full-rank formulations, the generative process must explore a largely empty semantic space, which makes sampling less efficient and can hinder stable reconstruction.

This observation suggests that explicitly modeling the low-dimensional structure of residuals may provide a more suitable inductive bias for *near-domain restoration* tasks. We propose the Low-Rank Residual Diffusion Model (LRDM), a novel generative paradigm for sparse residual spaces in *near-domain restoration* tasks as shown in Fig. 1. To provide a rigorous mathematical foundation, we first establish the Low-Rank Residual Hypothesis, demonstrating theoretically that confining the residual distribution to a low-rank subspace yields a tighter Variational Lower Bound (VLB). Grounded in this insight, LRDM introduces an Asymmetric Residual Diffusion Process. Specifically, the forward noise-injection process is strictly constrained to the low-rank manifold to prevent isotropic noise from polluting background invariants, while the reverse generative process is relaxed to the full-rank space to maintain high-fidelity reconstruction flexibility. Furthermore, to accommodate the dynamic nature of the diffusion trajectory, we design an Adaptive Rank Selection (ARS) mechanism. ARS dynamically calibrates the subspace rank across varying timesteps, ensuring that the model’s representational capacity optimally aligns with the evolving structural complexity of the residuals.

Our contributions are summarized as follows:

- (1) We reveal and theoretically justify the low-rank structure of residuals in *near-domain restoration* tasks, showing that a low-rank diffusion space yields a tighter variational bound.
- (2) We introduce LRDM, a diffusion framework with Adaptive Rank Selection that dynamically adjusts residual rank throughout the diffusion process.
- (3) Experiments on deraining, deblurring, and deshading show that LRDM delivers state-of-the-art fidelity and detail while using fewer sampling steps than full-rank diffusion models.

2. Related Work

2.1. Image Restoration

Image restoration aims to reconstruct clean, high-quality images from degraded observations. Classical approaches typically rely on hand-crafted priors such as non-local self-similarity [3], or gradient-domain statistics [22]. With the rise of deep learning, convolutional networks have become the dominant paradigm for tasks such as deraining [14, 53], deblurring [33, 44], and deshading [16]. These methods learn direct mappings from degraded to clean images, but often struggle to generalize to diverse degradations or handle fine-grained spatially varying artifacts.

More recent models introduce implicit priors via neural architectures, such as transformer-based restoration [52] and kernel-modulated networks for deblurring [55]. Although these models achieve strong performance, they remain constrained by deterministic one-shot prediction and lack the generative flexibility required to preserve uncertainty and fine-scale structures. This motivates the shift toward more expressive generative frameworks for restoration.

2.2. Diffusion Models for Image Restoration

Diffusion models have recently become a powerful tool for image restoration due to their strong generative priors and stable likelihood training [17, 41]. Early works adapt diffusion models to inverse problems by integrating classical degradation operators into the generative process, such as DDRM [20], DPS [11], and Palette [40]. These approaches exploit score-based sampling but still operate in the full image space, requiring large numbers of sampling steps to converge.

Subsequent improvements introduce task-specific structures. Whang et al. propose stochastic deblurring via diffusion with learned noise priors [48], while RDDM [27] formulates a residual-guided diffusion process to emphasize degradation-specific components. DiffUIR [56] deploys a selective hourglass architecture for multi-type degradations, and DeblurDiff [21] predicts latent blur kernels to guide reverse sampling. Despite their diverse model designs, all these methods retain a full-rank generative process, meaning the diffusion model repeatedly operates in the complete pixel or latent space. The residuals between domains are, in fact, sparse for near-domain tasks, predominantly exhibiting an inherent semantic pattern.

For *near-domain restoration* tasks, executing the diffusion process within a full-rank space introduces excessive degrees of freedom into the residual semantic space learned by the model. Addressing this long-overlooked inefficiency, this paper proposes a solution: replacing the full-rank core process with a low-rank residual diffusion mechanism specifically tailored for *near-domain restoration* tasks.

2.3. Low-Rank Methods in Vision

Low-rank modeling has a long-standing history in computer vision. Traditional matrix decomposition methods—such as PCA, RPCA [5], and low-rank matrix completion [4]—have been widely applied in video denoising, background subtraction, shadow removal, and image alignment. These approaches leverage the observation that many visual signals lie near low-dimensional subspaces, especially when the variation is shallow or structured.

In image restoration, low-rank assumptions have been used to capture patch redundancy [24], enforce structural consistency [15], or model global correlations in degraded images [54]. With deep learning, low-rank priors also appear in network compression, feature disentangling, and generative modeling [12, 57]. Among diffusion-based restoration methods, [34] represents the first attempt to deeply integrate low-rank priors into the diffusion pipeline, focusing specifically on the spectral low-rank property of hyperspectral images.

Our work introduces a novel perspective: in *near-domain restoration*, degradation residuals are inherently low-rank. By modeling diffusion within this subspace—via an asymmetric formulation and adaptive rank selection—we provide a principled, efficient alternative to existing diffusion paradigms.

3. Low-rank Residual Diffusion Models

3.1. Preliminary of Residual Diffusion Models

Residual Denoising Diffusion Models (RDDM) follow the standard paradigm of diffusion models [17, 41], which are typically defined by a forward diffusion process and the corresponding reverse denoising process. The forward process is a Markov Chain with a Gaussian transition in which data are gradually corrupted by Gaussian noise according to the variance schedules $\{\alpha_t, \beta_t\}_{t=1}^T : q(I_t|I_{t-1}) := \mathcal{N}(I_t; I_{t-1} + \alpha_t I_{res}, \beta_t^2 \mathbf{I})$, where $\{I_t\}_{t=1}^T$ are latent variables having the same dimensionality as data $I_0 \sim q(I_0)$. The term I_{in} denotes the degraded source domain image, and $I_{res} = I_{in} - I_0$ represents the domain gap, or *residual*, between the source and target domains. One notable property is that the forward process enables closed-form sampling of I_t at any desired timestep t :

$$I_t = I_0 + \bar{\alpha}_t I_{res} + \bar{\beta}_t \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (1)$$

where $\bar{\alpha}_t = \sum_{i=1}^t \alpha_i$ and $\bar{\beta}_t = \sqrt{\sum_{i=1}^t \beta_i^2}$. If $t = T$, $\bar{\alpha}_T = 1$ and $I_T = I_{in} + \bar{\beta}_T \epsilon$. The reverse process is another Markov Chain defined as:

$$q(I_{t-1}|I_t, I_0, I_{res}) = \mathcal{N}(I_{t-1}; I_0 + \bar{\alpha}_{t-1} I_{res} + \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2} \frac{I_t - (I_0 + \bar{\alpha}_t I_{res})}{\bar{\beta}_t}, \sigma_t^2 \mathbf{I}), \quad (2)$$

where $\sigma_t^2 = \eta \beta_t^2 \bar{\beta}_{t-1}^2 / \bar{\beta}_t^2$ and η controls whether the generation process is random ($\eta = 1$) or deterministic ($\eta = 0$). And using network to parameterize: $I_0^\theta = I_t - \bar{\alpha}_t I_{res}^\theta - \bar{\beta}_t \epsilon_\theta$:

$$I_{t-1} = I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{res}^\theta - (\bar{\beta}_t - \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}) \epsilon_\theta + \sigma_t \epsilon_t. \quad (3)$$

Thus, we can define $p_\theta(I_{t-1}|I_t) := q_\theta(I_{t-1}|I_t, I_0^\theta, I_{res}^\theta)$. The goal of RDDM is to optimize the KL divergence $D_{\text{KL}}(q(I_{t-1}|I_t, I_0, I_{res}) \| p_\theta(I_{t-1}|I_t))$, resulting in the following two loss terms:

$$\mathbb{E}[\|I_{res} - I_{res}^\theta(I_t, t, I_{in})\|^2], \quad (4)$$

$$\mathbb{E}[\|\epsilon - \epsilon_\theta(I_t, t, I_{in})\|^2], \quad (5)$$

where the training input image I_t is synthesized using I_0 , I_{res} , and ϵ by Eq. (1).

3.2. The Low-Rank Residual Hypothesis

Empirically (see Figs. 2 and 3), the residual between degraded and clean images exhibits a rapidly decaying singular-value spectrum: most energy is concentrated in a small number of leading components. As shown in Fig. 2, we observed that the specific residuals exhibit low-rank characteristics, i.e., a typical degradation pattern. Furthermore, the result of SVD decomposition and reconstruction using only a portion of the low-rank data is not significantly different from the original image.

Hypothesis 1 (Low-Rank Residual Hypothesis)

Let $I_{res} \in \mathbb{R}^D$ denote the residual, with an uncentered covariance matrix $\Sigma = \mathbb{E}[I_{res} I_{res}^\top]$. We hypothesize that for near-domain restoration, the eigenvalues of Σ ($\lambda_1 \geq \dots \geq \lambda_D \geq 0$) exhibit a rapid spectral decay. Consequently, there exists a rank- k subspace \mathcal{S}_k ($k \ll D$) spanning the principal components, such that the projection error is tightly bounded by an arbitrarily small threshold $\epsilon > 0$:

$$\mathbb{E}[\|I_{res} - \mathcal{Q}_k(I_{res})\|_2^2] = \sum_{i=k+1}^D \lambda_i \leq \epsilon, \quad (6)$$

where \mathcal{Q}_k is the orthogonal projection onto \mathcal{S}_k . In the ideal near-domain limit ($\epsilon \rightarrow 0$), the orthogonal complement vanishes: $\mathcal{Q}_k^\perp(I_{res}) = \mathbf{0}$.

Concretely, based on this assumption, we will explore the differences and connections between full-rank and low-rank models in residual diffusion models. Suppose Q is a

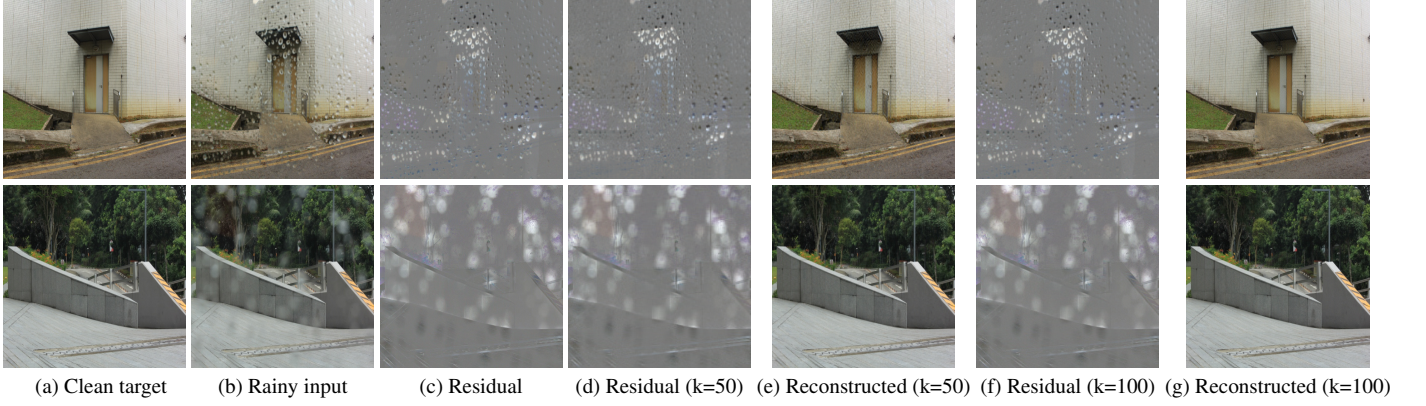


Figure 2. Visual comparisons of residual analysis and low-rank reconstruction with different ranks on the Raindrop dataset (image size 512×512). Columns (a)–(b) show the clear ground truth and rainy input images; (c) shows their residual; (d)–(e) and (f)–(g) present low-rank residuals and reconstructed images obtained via truncated SVD with $k=50$ and $k=100$, respectively.

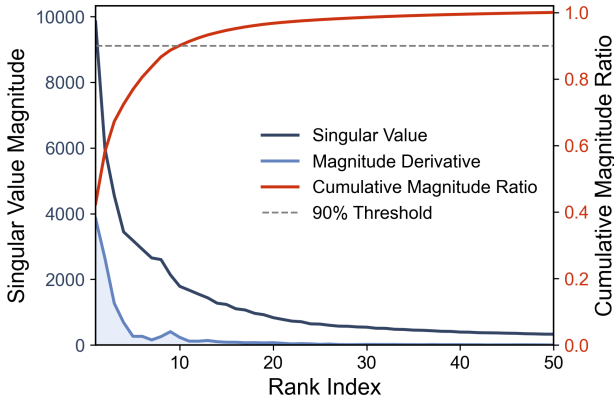


Figure 3. Joint visualization of singular values, their magnitude derivative, and cumulative magnitude ratio of the residual matrix between rainy and clean images. The sharp decay of singular values and the early saturation of the cumulative magnitude ($\approx 90\%$ within the top 10 ranks) indicate that most structural information is contained in a small number of dominant components, confirming the strong low-rank property of the residual signal. Results are averaged over samples from the Raindrop dataset with an image size of 512×512.

mapping from the original residual to the low-rank space \mathcal{S}_k . Therefore, we have $Q(I_{res}) = I_{res}$, $Q^\perp(I_{res}) = 0$, where Q^\perp is a mapping from the original residual to the complement of the low-rank space \mathcal{S}_k . We project $I_t - I_0$ onto two orthogonal subspaces:

$$\begin{aligned}
 I_{t,k} &\triangleq Q_k(I_t - I_0) \\
 &= Q_k(\bar{\alpha}_t I_{res} + \bar{\beta}_t \epsilon) \\
 &= \bar{\alpha}_t I_{res} + \bar{\beta}_t Q_k(\epsilon),
 \end{aligned} \tag{7}$$

$$\begin{aligned}
 I_{t,k}^\perp &\triangleq Q_k^\perp(I_t - I_0) \\
 &= Q_k^\perp(\bar{\alpha}_t I_{res} + \bar{\beta}_t \epsilon) \\
 &= \bar{\beta}_t Q_k^\perp(\epsilon).
 \end{aligned} \tag{8}$$

We observe that the $q(I_t|I_0, I_{res})$ process naturally decomposes into two independent Markov chains: (1) An RDDM process $q(I_{t,k}|I_0, I_{res})$ in the \mathcal{S}_k space containing all information of I_{res} ; (2) A standard DDPM (pure noise) process $q(I_{t,k}^\perp)$ in the \mathcal{S}_k^\perp space containing no information of I_{res} , starting from $I_{0,k}^\perp = 0$. Therefore, the true data distribution $q(I_t|I_0, I_{res})$ can be written as: $q(I_t|I_0, I_{res}) = q(I_{t,k}|I_0, I_{res}) \cdot q(I_{t,k}^\perp)$. Similarly, the posterior distribution $q(I_{t-1}|I_t, I_0, I_{res})$ can also be decomposed as:

$$q(I_{t-1}|I_t, I_0, I_{res}) = q(I_{t-1,k}|I_{t,k}, I_0, I_{res}) \cdot q(I_{t-1,k}^\perp|I_{t,k}^\perp). \tag{9}$$

Thus, the variational loss L_{t-1} can be expressed as [17]:

$$\begin{aligned}
 L_{t-1} &= \mathbb{E}_{q(I_0, I_t)} [D_{KL}(q(I_{t-1}|I_t, I_0, I_{res}) || p_\theta(I_{t-1}|I_t))] \\
 &= \mathbb{E}_q [D_{KL}(q(I_{t-1,k}|I_{t,k}, I_0, I_{res}) || p_\theta(I_{t-1,k}|I_t))] \\
 &\quad + \mathbb{E}_q [D_{KL}(q(I_{t-1,k}^\perp|I_{t,k}^\perp) || p_\theta(I_{t-1,k}^\perp|I_{t,k}^\perp))].
 \end{aligned} \tag{10}$$

We decompose the above equation into $L_{t-1} = L_{t-1,k}(\theta) + L_{t-1,k}^\perp(\theta)$. According to the fundamental property of KL divergence, we know that $D_{KL}(\cdot||\cdot) \geq 0$. Therefore, the second term $L_{t-1,k}^\perp(\theta)$ (i.e., the loss in the noisy space \mathcal{S}_k^\perp) as an expectation must be a non-negative value:

$$L_{t-1,k}^\perp(\theta) = \mathbb{E}_q [D_{KL}(\dots)] \geq 0. \tag{11}$$

This means that the total loss L_t of the full-rank model has a lower lower bound, namely $L_{t-1,k}(\theta)$:

$$L_{t-1}(\theta) = L_{t-1,k}(\theta) + L_{t-1,k}^\perp(\theta) \geq L_{t-1,k}(\theta). \tag{12}$$




Figure 4. (a) Rank evolution curves over time t , including linear (lin_increase, lin_decrease) and polynomial (poly_increase, poly_decrease with $p = 3$) trends, scaled to maximum rank 256; (b) Rank evolution $k(t)$ over time t for polynomial envelope functions with $p = 1, 2, 3, 4, 5$, showing increase and decrease trends, scaled to maximum rank 256.

Recall that optimizing a diffusion model entails maximizing the Variational Lower Bound (VLB), which is mathematically equivalent to minimizing the variational loss L_{t-1} . Since $L_{t-1,k}(\theta)$ represents the exact loss evaluated solely on the informative low-rank space \mathcal{S}_k , the inequality Eq. (12) provides a profound theoretical guarantee: under the Low-Rank Residual Hypothesis, restricting the generative process to the low-rank subspace intrinsically yields a strictly tighter VLB.

3.3. Asymmetry Residual Diffusion Process

Based on our low-rank residual hypothesis, instead of modeling the residuals I_{res} in the whole space \mathbb{R}^D , we redefine the diffusion process in a subspace of low-rank k -dimensional \mathcal{S}_k with low-rank residuals $I_{res}^{(k)} \triangleq Q_k(I_{res})$.

The forward process is a Markov chain whose Gaussian transition mean is constrained in a low-rank space: $q(I_t|I_{t-1}) := \mathcal{N}(I_t; I_{t-1} + \alpha_t I_{res}^{(k)}, \beta_t^2 \mathbf{I})$. The sample I_t (i.e., Eq. (1)) at any t time can be rewritten as:

$$I_t = I_0 + \bar{\alpha}_t Q_k(I_{res}) + \bar{\beta}_t \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (13)$$

where $\bar{\alpha}_t = \sum_{i=1}^t \alpha_i$ and $\bar{\beta}_t = \sqrt{\sum_{i=1}^t \beta_i^2}$. When $t = T$, $\bar{\alpha}_T = 1$ and $I_T = I_{in} + \bar{\beta}_T \epsilon$.

Although the forward noise addition process is low-rank, I_t itself is still full-rank due to the presence of the isotropic noise term $\bar{\beta}_t \epsilon$. Therefore, the true posterior probability is still full-rank. $q(I_{t-1} | I_t, I_0, I_{res}^{(k)})$ is a full-rank Gaussian with isotropic covariance determined by the forward noise schedule. Therefore, forcing the learned reverse kernel $p_\theta(I_{t-1} | I_t)$ to lie in the low-rank subspace \mathcal{S}_k would induce a fundamental *model mismatch*: a low-rank parameterization cannot represent the full-rank posterior and would necessarily discard components required to match the true conditional distribution.

To avoid this mismatch, LRDM adopts an **asymmetric** design: the forward process injects signal only in \mathcal{S}_k (regularizing the learned signal prior), while the reverse model p_θ

remains unrestricted and operates in the full ambient space \mathbb{R}^D . This design preserves the expressivity needed to approximate the full-rank posterior while benefiting from the tighter variational bound and inductive bias induced by the low-rank forward signal.

3.4. Adaptive Rank Selection Strategies

The Asymmetry Residual Diffusion Process introduced above implicitly assumes a fixed rank for the residual subspace throughout the forward process. This assumption is overly restrictive: the structural statistics of an image evolve over timesteps—early transitions are dominated by coarse, high-energy components, whereas later stages emphasize finer details and noise-like variations. To accommodate this progression, we decompose the residual representation I_{res} into subspaces of time-varying dimensionality and introduce a rank schedule $k(t)$ that adapts to the evolving diffusion dynamics. The forward diffusion process can be further rewritten using the time-dependent low-rank projection:

$$I_t = I_0 + \bar{\alpha}_t Q_k^{(t)}(I_{res}) + \bar{\beta}_t \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (14)$$

As illustrated in Fig. 4, we propose a family of adaptive rank selection strategies that parameterize the rank at each timestep $t \in [0, T]$ as a deterministic function of time. These schedules allow the model to flexibly modulate the low-rank constraint in accordance with the semantic complexity at different diffusion stages. We investigate four representative strategies:

Linear Increasing (LI): The rank grows linearly with time, gradually introducing higher-rank components as the process progresses. This schedule is defined as:

$$k_{LI}(t) = \left\lceil \frac{t}{T} \cdot R_m \right\rceil, \quad (15)$$

where t is the current timestep, T is the total number of timesteps, R_m is the maximum rank value, and $\lceil \cdot \rceil$ denotes the ceiling function.

Linear Decreasing (LD): The rank starts high and decreases linearly, prioritizing global and low-frequency structures in early stages while imposing stronger low-rank regularization later. Its formulation is:

$$k_{LD}(t) = \left\lceil \left(1 - \frac{t}{T}\right) \cdot R_m \right\rceil. \quad (16)$$

Polynomial Increasing (PI): The rank follows a non-linear growth curve, maintaining a low-rank approximation for most of the diffusion process and rapidly increasing rank near the end. It is derived from a polynomial envelope function:

$$\text{envelope}(d; p) = 1 + a \cdot d^p + b \cdot d^{p+1} + c \cdot d^{p+2}, \quad (17)$$




Figure 5. Visual comparison of deraining results on the Raindrop [35] and Rain1400 [14] datasets. Columns (a)–(b) show the rainy inputs and ground-truth clean images. Columns (c)–(f) present the outputs of representative baselines, including Restormer, X-Restormer, Diff-Plugin, and RDDM. Column (g) shows the results of our proposed LRDM. Across both datasets, LRDM produces cleaner structures, sharper textures, and more faithful restoration in challenging regions (e.g., water-streak areas and high-frequency details), demonstrating its effectiveness in modeling informative low-rank residuals within the diffusion framework.

Raindrop Removal (Raindrop)			Deraining (Rain1400)		
Method	PSNR \uparrow	SSIM \uparrow	Method	PSNR \uparrow	SSIM \uparrow
AttentGAN [35]	31.59	0.917	Uformer [47]	32.84	0.931
Quanetal <i>et al.</i> [36]	31.37	0.918	Restormer [52]	33.68	0.939
IDT [50]	31.87	0.931	DRSformer [8]	33.66	0.939
UDR-S ² [7]	32.64	0.942	UDR-S ² [7]	33.08	0.930
Restormer [52]	31.67	<u>0.958</u>	AirNet [23]	32.36	0.928
NAFNet [6]	31.03	0.952	DA-CLIP [30]	29.67	0.851
MPerceiver [1]	33.21	0.929	NAFNet [6]	32.38	0.943
RainDiff64* [52]	32.29	0.942	MPerceiver [52]	33.40	0.937
RainDiff128 [6]	32.43	0.933	EMResformer [6]	32.93	0.936
RDDM [27]	32.51	0.956	RDDM [27]	32.21	<u>0.952</u>
LRDM (Ours)	<u>33.09</u>	0.967	LRDM (Ours)	34.39	0.954

Table 1. Quantitative comparison (PSNR/SSIM) on **Raindrop** and **Rain1400** datasets. Bold text represents the best result, and horizontal lines represent the second-best result.

where $d = t/T$ (normalized timestep), p is the polynomial exponent, and coefficients are $a = -\frac{(p+1)(p+2)}{2}$, $b = p(p+2)$, $c = -\frac{p(p+1)}{2}$. The rank evolution is:

$$k_{PI}(t; p) = \left[\left(1 - \text{envelope} \left(\frac{t}{T}; p \right) \right) \cdot R_m \right]. \quad (18)$$

Polynomial Decreasing (PD): The inverse of PI, where the rank starts high and decays polynomially, emphasizing high-rank reconstruction early and enforcing rapid compression as the diffusion stabilizes. Using the same envelope function, its expression is:

$$k_{PD}(t; p) = \left[\text{envelope} \left(\frac{t}{T}; p \right) \cdot R_m \right]. \quad (19)$$

These adaptive schedules enable the diffusion model to flexibly balance low-rank structure preservation and high-frequency detail reconstruction throughout the process. Empirically, we find that the choice of scheduling function significantly influences both convergence stability and perceptual quality in Tab. 4, suggesting that rank dynamics play a crucial role in regulating the flow of information between different subspaces.

Method	RealBlur-J		RealBlur-R	
	PSNR	SSIM	PSNR	SSIM
Restormer [52]	28.96	0.878	36.19	0.957
X-Restormer [9]	28.87	0.878	36.27	<u>0.958</u>
HI-Diff [10]	29.15	0.889	36.28	<u>0.958</u>
UFPNet [13]	29.87	0.884	36.25	0.952
AdaRevD [31]	<u>30.12</u>	<u>0.894</u>	<u>36.53</u>	0.957
DiffBIR [26]	26.92	0.745	32.60	0.849
OSDiff [49]	26.83	0.801	33.54	0.905
Diff-Plugin [29]	25.77	0.771	32.64	0.849
FideDiff [28]	28.96	0.869	36.01	0.942
LRDM (Ours)	30.21	0.933	37.92	0.976

Table 2. Quantitative comparison on **RealBlur-J** and **RealBlur-R**.

4. Experiments and Analysis

4.1. Datasets and Metrics

We evaluate our Low-Rank Residual Diffusion Models (LRDM) with adaptive rank modulation across a diverse set of vision restoration tasks to demonstrate its robustness and generalization ability. For image deraining, we use the Raindrop [35] and Rain1400 [14] datasets. Image deblurring experiments are conducted on GoPro [33], RealBlur-J [38], and RealBlur-R [38]. For shadow removal, we adopt the ISTD dataset [45]. For image inpainting, we adopt the CelebA-HQ dataset [19]. These datasets cover both synthetic and real-world degradations, providing a comprehensive testbed for evaluating the effectiveness of our method.

Following standard evaluation protocols [27], we report restoration fidelity and structural quality using PSNR and SSIM [46].

4.2. Quantitative Comparison

We evaluate LRDM against recent state-of-the-art methods on two deraining benchmarks (Raindrop [35], Rain1400

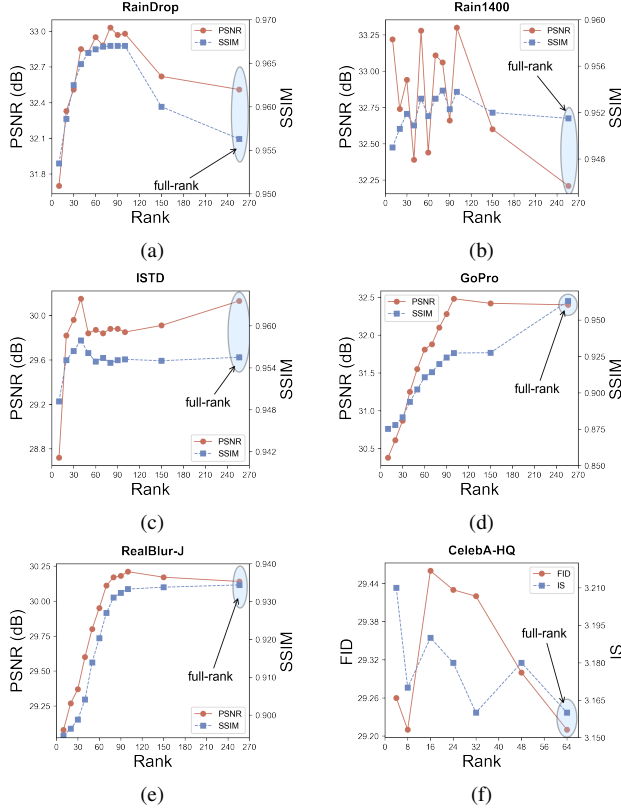


Figure 6. Analysis of PSNR/SSIM versus rank on (a) Raindrop, (b) Rain1400, (c) ISTD, (d) GoPro, (e) RealBlur-J, and (f) CelebA-HQ. Across all datasets, our method achieves optimal performance at medium ranks, consistently outperforming the full-rank baseline. This demonstrates that an appropriate low-rank constraint introduces a beneficial inductive bias, effectively enhancing restoration across diverse degradation types and domains.




Figure 7. Comparison of image restoration results under different fixed rank settings and without low rank (w.o. LR), where GT is the ground truth image, Input is the input image, r represents the rank, and PSNR values are used to evaluate the quality of each restoration result.

[14] and two real-world deblurring benchmarks (RealBlur-J [38], RealBlur-R [38]). PSNR and SSIM are used as the evaluation metrics; higher values indicate better restoration.

Tab. 1 summarizes the quantitative results on Raindrop and Rain1400. On Raindrop, LRDM achieves the

best SSIM (**0.967**) and the second-highest PSNR (**33.09 dB**), while ranking first among all diffusion-based models. On Rain1400, LRDM delivers the best overall performance, achieving **34.39 dB PSNR** and **0.954 SSIM**, again surpassing all diffusion-based baselines as well as transformer/CNN methods. These results highlight that enforcing a low-rank constraint on the residual pathway enables diffusion models to better preserve structures and suppress irrelevant high-frequency noise.

Tab. 2 summarizes results on RealBlur-J and RealBlur-R. LRDM obtains the best reported scores in our comparisons (**30.21 dB PSNR / 0.933 SSIM** on RealBlur-J and **37.92 dB PSNR / 0.976 SSIM** on RealBlur-R). Compared to several diffusion-based methods, LRDM shows noticeable improvements; against top non-diffusion baselines the margins are smaller but still positive. Overall, the quantitative results indicate that low-rank residual modeling is broadly effective across both synthetic and real degradations.

As shown in Fig. 5, we provide a visual comparison with representative baseline methods on the Raindrop and Rain1400 datasets. The results show that LRDM removes rain streaks while better preserving background details.

4.3. Fixed-rank Ablation

We first study the effect of using a fixed low-rank approximation for the residual across multiple datasets. Fig. 6 plots PSNR/SSIM as a function of rank for Raindrop, Rain1400, ISTD, GoPro, RealBlur-J and CelebA-HQ. A consistent pattern emerges: performance rises rapidly from very low ranks, peaks in a mid-range, and then saturates or slightly declines toward full rank.

For example, on Raindrop PSNR increases from **31.7 dB** (rank 10) to a peak near **33.03 dB** (rank 80), then slightly drops at full rank. Rain1400 shows a similar behaviour with an optimum near rank 100. ISTD and GoPro follow the same trend though with different saturation points. These observations support the hypothesis that a modest low-rank subspace captures most task-relevant residual structure while suppressing redundant noise. The qualitative results in Fig. 7 further confirm this finding. We can intuitively see that a fixed low rank (e.g., $r=20$ or $r=40$) is sufficient to achieve recovery performance comparable to, or even better than, a full-rank (w.o. LR) baseline in terms of PSNR.

4.4. Adaptive Rank Modulation

Motivated by the fixed-rank findings, we evaluate adaptive rank schedules that vary $k(t)$ over diffusion timesteps. We compare four representative schedules: Linear Increasing (LI), Linear Decreasing (LD), Polynomial Increasing (PI) and Polynomial Decreasing (PD).

Tab. 3 reports the summary on Raindrop and Rain1400. Dynamic schedules consistently outperform fixed-rank baselines, with the polynomial-decreasing (PD) schedule

Modulation	Raindrop		Rain1400	
	PSNR	SSIM	PSNR	SSIM
lin_increase	<u>32.95</u>	0.9670	<u>33.01</u>	<u>0.9535</u>
lin_decrease	32.92	0.9668	32.79	0.9534
poly_increase	32.69	0.9650	32.88	0.9516
poly_decrease	33.03	<u>0.9669</u>	33.39	0.9540

Table 3. Performance comparison of different modulation rank methods on Raindrop [35] and Rain1400 [14] datasets.

Modulation and Orders	Raindrop		Rain1400		
	PSNR	SSIM	PSNR	SSIM	
poly_decrease	(p=1)	32.57	0.9651	34.39	0.9542
	(p=2)	33.09	0.9670	33.54	0.9538
	(p=3)	33.03	0.9669	33.39	0.9540
	(p=4)	32.92	0.9666	33.22	0.9535
	(p=5)	33.03	0.9667	33.34	0.9538
poly_increase	(p=1)	32.74	0.9664	32.94	0.9541
	(p=2)	32.83	0.9663	32.97	0.9534
	(p=3)	32.69	0.9650	32.88	0.9516
	(p=4)	32.56	0.9639	33.11	0.9501
	(p=5)	32.48	0.9615	33.60	0.9481

Table 4. Performance comparison of different polynomial dynamic scheduling orders on Raindrop and Rain1400 datasets.

(specifically at $p = 3$) producing the best overall balance. PD reaches 33.03 dB / 0.9669 SSIM on Raindrop and 33.39 dB / 0.9540 SSIM on Rain1400.

We further analyze the effect of the polynomial order p in Tab. 4. In practice, lower-order schedules yield better performance, providing an effective balance between early-stage low-rank regularization and sufficient capacity for detail recovery in later steps. This pattern indicates that a mildly decreasing rank schedule (with respect to the diffusion timesteps) aligns best with the diffusion process’s inherent coarse-to-fine progression.

4.5. Sampling Steps and Efficiency

We assess rank modulation under different sampling steps on Rain1400. As shown in Fig. 8, low-rank configurations—particularly the PD schedule—achieve higher PSNR in the small-step regime (1–3 steps) compared with higher-rank LRDM variants that use more sampling steps. Remarkably, PD attains 34.65 dB with only 2 sampling steps, even surpassing some full-rank configurations that rely on larger sampling budgets. As expected, this advantage narrows as the number of steps increases, but within the practical low-step regime (≤ 5), adaptive low-rank modulation offers a compelling balance between quality and efficiency.




Figure 8. PSNR performance versus sampling steps on the Rain1400 dataset. Our adaptive polynomial modulation strategy achieves significantly higher PSNR at few sampling steps (1-5) compared to fixed-rank models, demonstrating superior convergence speed.

Model	Rain1400		GoPro		ISTD	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
RDDM [27]	32.21	0.952	32.40	0.963	30.13	0.956
RDDM + Low-Rank	33.30	0.954	32.48	0.927	30.15	0.958
ResShift [51]	25.41	0.713	26.38	0.803	24.92	0.719
ResShift + Low-Rank	25.44	0.716	27.18	0.809	25.23	0.724

Table 5. Ablation study on the feasibility of the proposed low-rank space across three datasets. Low-rank constraints improve PSNR across all datasets and maintain competitive SSIM.

4.6. Ablation on Low-Rank Feasibility

To validate the general applicability of low-rank constraints, we integrate our low-rank module into two different baselines [27, 51] and evaluate on Rain1400, GoPro and ISTD shown in Tab. 5. Adding the low-rank constraint improves PSNR/SSIM for diffusion-based baselines, with only a slight trade-off in GoPro’s SSIM. For ResShift the gains are more varied across datasets but remain positive overall. These controlled experiments indicate that the proposed low-rank space is a generally useful inductive bias for *near-domain restoration*.

5. Conclusion

We presented Low-Rank Residual Diffusion Model (LRDM), a diffusion framework that models the restoration process within a low-rank residual subspace. By constraining diffusion dynamics and adaptively adjusting the rank across timesteps, LRDM achieves efficient and structure-preserving image restoration. This work reveals the intrinsic low-rank nature of near-domain transformations and provides a new perspective for designing efficient and interpretable diffusion models.

Limitation: We acknowledge that our low-rank hypothesis inherently relies on the *near-domain* assumption. Therefore, LRDM’s applicability is currently restricted in far-domain tasks with extreme structural deformations.

References

- [1] Yuang Ai, Huaibo Huang, Xiaoqiang Zhou, Jiexiang Wang, and Ran He. Multimodal prompt perceiver: Empower adaptiveness generalizability and fidelity for all-in-one image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25432–25444, 2024. 6
- [2] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18208–18218, 2022. 1
- [3] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 60–65 vol. 2, 2005. 2
- [4] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Found Comput Math* 9, 9(6):717–772, 2009. 3
- [5] Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *J. ACM*, 58(3), 2011. 3
- [6] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European conference on computer vision*, pages 17–33, 2022. 6
- [7] Sixiang Chen, Tian Ye, Jinbin Bai, Erkang Chen, Jun Shi, and Lei Zhu. Sparse sampling transformer with uncertainty-driven ranking for unified removal of raindrops and rain streaks. In *IEEE/CVF International Conference on Computer Vision*, pages 13106–13117, 2023. 6
- [8] Xiang Chen, Hao Li, Mingqiang Li, and Jinshan Pan. Learning a sparse transformer network for effective image deraining. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5896–5905, 2023. 1, 6
- [9] Xiangyu Chen, Zheyuan Li, Yuandong Pu, Yihao Liu, Jiantao Zhou, Yu Qiao, and Chao Dong. A comparative study of image restoration networks for general backbone network design. In *European conference on computer vision*, 2023. 6
- [10] Zheng Chen, Yulun Zhang, Liu Ding, Xia Bin, Jinjin Gu, Linghe Kong, and Xin Yuan. Hierarchical integration diffusion model for realistic image deblurring. In *International Conference on Neural Information Processing Systems*, 2023. 6
- [11] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [12] Emily Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *International Conference on Neural Information Processing Systems - Volume 1*, page 1269–1277, 2014. 3
- [13] Zhenxuan Fang, Fangfang Wu, Weisheng Dong, Xin Li, Jinjian Wu, and Guangming Shi. Self-supervised non-uniform kernel estimation with flow-based motion prior for blind image deblurring. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18105–18114, 2023. 6
- [14] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 6, 7, 8
- [15] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted nuclear norm minimization with application to image denoising. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2862–2869, 2014. 3
- [16] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1777–1786, 2020. 2
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851, 2020. 1, 2, 3, 4
- [18] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1125–1134, 2017. 1
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 6
- [20] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, pages 23593–23606, 2022. 1, 2
- [21] Lingshun Kong, Jiawei Zhang, Dongqing Zou, Jimmy Ren, Xiaohe Wu, Jiangxin Dong, and Jinshan Pan. Deblurdiff: Real-world image deblurring with generative diffusion models. arXiv preprint arXiv:2502.03810, 2025. 2
- [22] Dilip Krishnan and Rob Fergus. Fast image deconvolution using hyper-laplacian priors. In *Advances in Neural Information Processing Systems*, 2009. 2
- [23] Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *IEEE/CVF conference on computer vision and pattern recognition*, pages 17452–17462, 2022. 6
- [24] Jianwei Li, Xiaoou Chen, Dongqing Zou, Bo Gao, and Wei Teng. Conformal and low-rank sparse representation for image restoration. In *IEEE International Conference on Computer Vision (ICCV)*, pages 235–243, 2015. 3
- [25] Tong Li, Hansen Feng, Lizhi Wang, Lin Zhu, Zhiwei Xiong, and Hua Huang. Stimulating diffusion model for image denoising via adaptive embedding and ensembling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1
- [26] Xinqi Lin, Jingwen He, Ziyang Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. Diffbir: Toward blind image restoration with generative diffusion prior. In *European Conference on Computer Vision (ECCV)*, page 430–448, 2024. 6
- [27] Jiawei Liu, Qiang Wang, Huijie Fan, Yinong Wang, Yandong Tang, and Liangqiong Qu. Residual denoising diffusion models. In *IEEE/CVF Conference on Computer Vision*

- and *Pattern Recognition (CVPR)*, pages 2773–2783, 2024. 2, 6, 8
- [28] Xiaoyang Liu, Zhengyan Zhou, Zihang Xu, Jiezhong Cao, Zheng Chen, and Yulun Zhang. Fidediff: Efficient diffusion model for high-fidelity image motion deblurring. In *International Conference on Learning Representations (ICLR)*, 2026. 6
- [29] Yuhao Liu, Zhanghan Ke, Fang Liu, Nanxuan Zhao, and Rynson W.H. Lau. Diff-plugin: Revitalizing details for diffusion-based low-level tasks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 6
- [30] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models for universal image restoration. *arXiv preprint arXiv:2310.01018*, 3(8), 2023. 6
- [31] Xintian Mao, Qingli Li, and Yan Wang. Adarevd: Adaptive patch exiting reversible decoder pushes the limit of image deblurring. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25681–25690, 2024. 6
- [32] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022. 1
- [33] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3207–3215, 2017. 1, 2, 6
- [34] Li Pang, Xiangyu Rui, Long Cui, Hongzhong Wang, Deyu Meng, and Xiangyong Cao. Hir-diff: Unsupervised hyper-spectral image restoration via improved diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3005–3014, 2024. 3
- [35] Rui Qian, Robby T. Tan, Wenhan Yang, Jiajun Su, and Jiaying Liu. Attentive generative adversarial network for rain-drop removal from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2482–2491, 2018. 6, 8
- [36] Yuhui Quan, Shijie Deng, Yixin Chen, and Hui Ji. Deep learning for seeing through window with raindrops. In *IEEE/CVF International Conference on Computer Vision*, pages 2463–2471, 2019. 6
- [37] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. In *arXiv preprint arXiv:2204.06125*, 2022. 1
- [38] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *European Conference on Computer Vision (ECCV)*, 2020. 6, 7
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 1
- [40] Chitwan Saharia, William Chan, Huiwen Chang, Chris A Lee, Jonathan Ho, Tim Salimans, David J Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. 1, 2
- [41] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 1, 2, 3
- [42] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, 2019.
- [43] 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*, 2021. 1
- [44] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Jiaya Jia. Scale-recurrent network for deep image deblurring. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8174–8182, 2018. 2
- [45] Jifeng Wang, Xiang Li, and Jian Yang. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [46] 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. 6
- [47] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17683–17693, 2022. 6
- [48] Jay Whang, Mauricio Delbracio, Hossein Talebi, Peyman Milanfar, and Hieu Kim. Deblurring via stochastic refinement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1760–1769, 2022. 1, 2
- [49] Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. *arXiv preprint arXiv:2406.08177*, 2024. 6
- [50] Jie Xiao, Xueyang Fu, Aiping Liu, Feng Wu, and Zheng-Jun Zha. Image de-raining transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):12978–12995, 2022. 6
- [51] 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*, 36:13294–13307, 2023. 8
- [52] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 2, 6
- [53] He Zhang and Vishal M. Patel. Density-aware single image de-raining using a multi-stream dense network. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 695–704, 2018. 2

- [54] 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. [3](#)
- [55] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Deep plug-and-play super-resolution for arbitrary blur kernels. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1671–1681, 2019. [2](#)
- [56] Dian Zheng, Xiao-Ming Wu, Shuzhou Yang, Jian Zhang, Jian-Fang Hu, and Wei-Shi Zheng. Selective hourglass mapping for universal image restoration based on diffusion model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25445–25455, 2024. [2](#)
- [57] Jiapeng Zhu, Ruili Feng, Yujun Shen, Deli Zhao, Zhengjun Zha, Jingren Zhou, and Qifeng Chen. Low-rank subspaces in GANs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [3](#)
- [58] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017. [1](#)