

Residual Denoising Diffusion Models

Jiawei Liu^{1,2,3}, Qiang Wang^{1,4}, Huijie Fan^{1,2*}, Yinong Wang⁵, Yandong Tang^{1,2}, Liangqiong Qu^{5*}

¹State Key Laboratory of Robotics, Shenyang Institute of Automation, Chinese Academy of Sciences

²Institutes for Robotics and Intelligent Manufacturing, Chinese Academy of Science

³University of Chinese Academy of Sciences ⁴Shenyang University ⁵The University of Hong Kong

{liujiawei@sjtu.edu.cn, qiangwang@sjtu.edu.cn, huijie.fan@siar.ac.cn, yinongwang@sjtu.edu.cn, yandongtang@sjtu.edu.cn, liangqiongqu@hku.hk}

{qiangwang@sjtu.edu.cn, huijie.fan@siar.ac.cn, yinongwang@sjtu.edu.cn, yandongtang@sjtu.edu.cn, liangqiongqu@hku.hk}

Abstract

We propose residual denoising diffusion models (RDDM), a novel dual diffusion process that decouples the traditional single denoising diffusion process into residual diffusion and noise diffusion. This dual diffusion framework expands the denoising-based diffusion models, initially uninterpretable for image restoration, into a unified and interpretable model for both image generation and restoration by introducing residuals. Specifically, our residual diffusion represents directional diffusion from the target image to the degraded input image and explicitly guides the reverse generation process for image restoration, while noise diffusion represents random perturbations in the diffusion process. The residual prioritizes certainty, while the noise emphasizes diversity, enabling RDDM to effectively unify tasks with varying certainty or diversity requirements, such as image generation and restoration. We demonstrate that our sampling process is consistent with that of DDPM and DDIM through coefficient transformation, and propose a partially path-independent generation process to better understand the reverse process. Notably, our RDDM enables a generic UNet, trained with only an L1 loss and a batch size of 1, to compete with state-of-the-art image restoration methods. We provide code and pre-trained models to encourage further exploration, application, and development of our innovative framework (<https://github.com/nachifur/RDDM>).

1. Introduction

In real-life scenarios, diffusion often occurs in complex forms involving multiple, concurrent processes, such as the dispersion of multiple gases or the propagation of different types of waves or fields. This leads us to ponder whether the denoising-based diffusion models [17, 51] have limitations in focusing solely on denoising. Current diffusion-




Figure 1. Denoising diffusion process - DDPM [17] (a) and our residual denoising diffusion process (b). For image restoration, we introduce residual diffusion to represent the diffusion direction from the target image to the input image.

based image restoration methods [22, 39, 48, 49, 82] extend the diffusion model to image restoration tasks by using degraded images as a condition input to implicitly guide the reverse generation process, without modifying the original denoising diffusion process [17, 51]. However, the reverse process starting from noise seems to be unnecessary, as the degraded image is already known. The forward process is non-interpretability for image restoration, as the diffusion process does not contain any information about the degraded image, as shown in Fig. 1(a).

In this paper, we explore a novel dual diffusion process and propose Residual Denoising Diffusion Models (RDDM), which can tackle the non-interpretability of a single denoising process for image restoration. In RDDM, we decouple the previous diffusion process into residual diffusion and noise diffusion. Residual diffusion prioritizes certainty and represents a directional diffusion from the target image to the conditional input image, and noise diffusion emphasizes diversity and represents random perturbations in the diffusion process. Thus, our RDDM can unify different tasks that require different certainty or diversity, e.g., image generation and restoration. Compared to denoising-based diffusion models for image restoration, the residuals in RDDM clearly indicate the forward diffusion direction and explicitly guide the reverse generation process for image restoration, as shown in Fig. 1(b).

Specifically, we redefine a new forward process that al-

*Corresponding author.

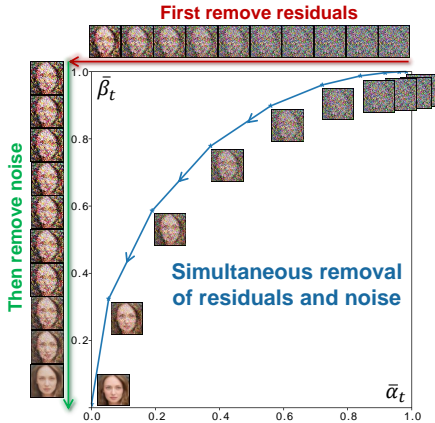


Figure 2. Decoupled dual diffusion framework. The previous forward diffusion process is decoupled into residual diffusion and noise diffusion, while in the reverse process, the simultaneous sampling can be decoupled into first removing the residuals and then removing noise.

lowers simultaneous diffusion of residuals and noise, wherein the target image progressively diffuses into a purely noisy image for image generation or a noise-carrying input image for image restoration. Unlike the previous denoising diffusion model [17, 51], which uses one coefficient schedule to control the mixing ratio of noise and images, our RDDM employs two independent coefficient schedules to control the diffusion speed of residuals and noise. We found that this independent diffusion property is also evident in the reverse generation process, e.g., readjusting the coefficient schedule within a certain range during testing does not affect the image generation results, and removing the residuals firstly, followed by denoising (see Fig. 2), can also produce semantically consistent images. Our RDDM is compatible with widely used denoising diffusion models, i.e., our sampling process is consistent with that of DDPM [17] and DDIM [51] by transforming coefficient schedules. In addition, our RDDM natively supports conditional inputs, enabling networks trained with only an ℓ_1 loss and a batch size of 1 to compete with state-of-the-art image restoration methods. We envision that our models can facilitate a unified and interpretable image-to-image distribution transformation methodology, highlighting that residuals and noise are equally important for diffusion models, e.g., the residual prioritizes certainty while the noise emphasizes diversity. The contributions of this paper are summarized as follows:

- We propose a novel dual diffusion framework to tackle the non-interpretability of a single denoising process for image restoration by introducing residuals. Our residual diffusion represents a directional diffusion from the target image to the conditional input image.
- We introduce a partially path-independent generation process that decouples residuals and noise, highlighting their roles in controlling directional residual shift (certainty)

and random perturbation (diversity), respectively.

- We design an automatic objective selection algorithm to choose whether to predict residuals or noise for unknown new tasks.
- Extensive experiments demonstrate that our method can be adapted to different tasks, e.g., image generation, restoration, inpainting and translation, focusing certainty or diversity, and involving paired or unpaired data.

2. Related Work

Denoising diffusion models (e.g., DDPM [17], SGM [52, 53], and DDIM [51]) were initially developed for image generation. Subsequent image restoration methods [14, 39, 48] based on DDPM and DDIM feed a degraded image as a conditional input to a denoising network, e.g., DvSR [62], SR3 [49], and WeatherDiffusion [82], which typically require large sampling steps and batch sizes. Additionally, the reverse process starting from noise in these methods seems unnecessary and inefficient for image restoration tasks. Thus, SDEdit [41], ColdDiffusion [2], InDI [11], and I2SB [29] propose generating a clear image directly from a degraded image or noise-carrying degraded image. InDI [11] and I2SB [29], which also present unified image generation and restoration frameworks, are the most closely related to our proposed RDDM. Specifically, the forward diffusion of InDI, I2SB, and our RDDM consistently employs a mixture of three terms (i.e., input images I_{in} , target images I_0 , and noise ϵ), extending beyond the denoising-based diffusion model [17, 51] which incorporates a mixture of two terms (i.e., I_0 and ϵ). However, InDI [11] and I2SB [29] opt for estimating the target image or its linear transformation term to replace the noise estimation, akin to a special case of our RDDM (SM-Res). In contrast, we introduce residual estimation while also embracing noise for both generation and restoration tasks. Our RDDM can further extend DDPM [17], DDIM [51], InDI [11], and I2SB [29] to independent double diffusion processes, and pave the way for the multi-dimensional diffusion process. We highlight that residuals and noise are equally important, e.g., the residual prioritizes certainty while the noise emphasizes diversity. In addition, our work is related to coefficient schedule design [44, 48], variance strategy optimization [3, 4, 24, 44], superimposed image decomposition [12, 81], curve integration [47], stochastic differential equations [53], and residual learning [15] for image restoration [1, 32, 56, 70, 72, 75]. See Appendix A.5 for detailed comparison.

3. Background

Denoising diffusion models [17, 50] aim to learn a distribution $p_\theta(I_0) := \int p_\theta(I_{0:T}) dI_{1:T}^1$ to approximate a tar-

¹To understand diffusion from an image perspective, we use I instead of x in DDPM [17].




Figure 3. The proposed residual denoising diffusion model (RDDM) is a unified framework for image generation and restoration (a shadow removal task is shown here). We introduce residuals (I_{res}) in RDDM, redefining the forward diffusion process to involve simultaneous diffusion of residuals and noise. The residuals ($I_{res} = I_{in} - I_0$) diffusion represents the directional diffusion from the target image I_0 to the degraded input image I_{in} , while the noise (ϵ) diffusion represents the random perturbations in the diffusion process. In RDDM, I_0 gradually diffuses into $I_T = I_{in} + \epsilon$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. In the third columns, I_T is a purely noisy image for image generation since $I_{in} = 0$, and a noise-carrying degraded image for image restoration as I_{in} is the degraded image.

get data distribution $q(I_0)$, where I_0 are target images and I_1, \dots, I_T ($T = 1000$) are latent images of the same dimension as I_0 . In the forward process, $q(I_0)$ is diffused into a Gaussian noise distribution using a fixed Markov chain,

$$q(I_{1:T}|I_0) := \prod_{t=1}^T q(I_t|I_{t-1}), \quad (1)$$

$$q(I_t|I_{t-1}) := \mathcal{N}(I_t; \sqrt{\alpha_t}I_{t-1}, (1 - \alpha_t)\mathbf{I}), \quad (2)$$

where $\alpha_{1:T} \in (0, 1]^T$. $q(I_t|I_{t-1})$ can also be written as $I_t = \sqrt{\alpha_t}I_{t-1} + \sqrt{1 - \alpha_t}\epsilon_{t-1}$. In fact, it is simpler to sampling I_t from I_0 by reparameterization [25, 26],

$$I_t = \sqrt{\bar{\alpha}_t}I_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad (3)$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. The reverse process is also a Markov chain starting at $p_\theta(I_T) \sim \mathcal{N}(I_T; \mathbf{0}, \mathbf{I})$,

$$p_\theta(I_{0:T}) := p_\theta(I_T) \prod_{t=1}^T p_\theta(I_{t-1}|I_t), \quad (4)$$

$$p_\theta(I_{t-1}|I_t) := \mathcal{N}(I_{t-1}; \mu_\theta(I_t, t), \Sigma_t \mathbf{I}), \quad (5)$$

where $p_\theta(I_{t-1}|I_t)$ is a learnable transfer probability (the variance schedule Σ_t is fixed). A simplified loss function [17] is derived from the maximum likelihood of $p_\theta(I_0)$, i.e., $L(\theta) := \mathbb{E}_{I_0 \sim q(I_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\epsilon - \epsilon_\theta(I_t, t)\|^2]$. The estimated noise ϵ_θ can be used to represent μ_θ in $p_\theta(I_{t-1}|I_t)$, thus I_{t-1} can be sampled from $p_\theta(I_{t-1}|I_t)$ step by step.

4. Residual Denoising Diffusion Models

Our goal is to develop a dual diffusion process to unify and interpret image generation and restoration. We modify the representation of $I_T = \epsilon$ in traditional DDPM to $I_T = I_{in} + \epsilon$ in our RDDM, where I_{in} is a degraded image (e.g., a shadow, low-light, or blurred image) for image restoration and is set to 0 for image generation. This modification is compatible with the widely used denoising diffusion model, e.g., $I_T = 0 + \epsilon$ is the pure noise (ϵ) for generation. For image restoration, I_T is a noisy-carrying degraded image ($I_{in} + \epsilon$), as shown in the third column in Fig. 3.

The modified forward process from I_0 to $I_T = I_{in} + \epsilon$ involves progressively degrading I_0 to I_{in} , and injecting noise ϵ . This naturally results in a dual diffusion process, a residual diffusion to model the transition from I_0 to I_{in} and a noise diffusion. For example, the forward diffusion process from the shadow-free image I_0 to the noisy carrying shadow image I_T involves progressively adding shadows and noise, as shown in the second row in Fig. 3.

In the following subsections, we detail the underlying theory and the methodology behind our RDDM. Inspired by residual learning [15, 31, 32], we redefine each forward diffusion process step in Section 4.1. For the reverse process, we present a training objective to predict the residuals and noise injected in the forward process in Section 4.2. In Section 4.3, we propose three sampling methods, i.e., residual prediction (SM-Res), noise prediction (SM-N), and “residual and noise prediction” (SM-Res-N).

4.1. Directional Residual Diffusion Process with Perturbation

To model the gradual degradation of image quality and the increment of noise, we define the single forward process step in our RDDM as follows:

$$I_t = I_{t-1} + I_{res}^t, \quad I_{res}^t \sim \mathcal{N}(\alpha_t I_{res}, \beta_t^2 \mathbf{I}), \quad (6)$$

where I_{res}^t represents a directional mean shift (residual diffusion) with random perturbation (noise diffusion) from state I_{t-1} to state I_t , the residuals I_{res} in I_{res}^t is the difference between I_{in} and I_0 (i.e., $I_{res} = I_{in} - I_0$), and two independent coefficient schedules α_t and β_t control the residual and noise diffusion, respectively. In fact, it is simpler to sample I_t from I_0 (like Eq. 3),

$$\begin{aligned} I_t &= I_{t-1} + \alpha_t I_{res} + \beta_t \epsilon_{t-1}, \\ &= I_{t-2} + (\alpha_{t-1} + \alpha_t) I_{res} + (\sqrt{\beta_{t-1}^2 + \beta_t^2}) \epsilon_{t-2} \\ &= \dots \\ &= I_0 + \bar{\alpha}_t I_{res} + \bar{\beta}_t \epsilon, \end{aligned} \quad (7)$$

where $\epsilon_{t-1}, \dots, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\bar{\alpha}_t = \sum_{i=1}^t \alpha_i$ and $\bar{\beta}_t = \sqrt{\sum_{i=1}^t \beta_i^2}$. If $t = T$, $\bar{\alpha}_T = 1$ and $I_T = I_{in} + \bar{\beta}_T \epsilon$. $\bar{\beta}_T$ can control the intensity of noise perturbation for image restoration (e.g., $\bar{\beta}_T^2 = 0.01$ for shadow removal), while $\bar{\beta}_T^2 = 1$ for image generation. From Eq. 6, the joint probability distributions in the forward process can be defined as:

$$q(I_{1:T}|I_0, I_{res}) := \prod_{t=1}^T q(I_t|I_{t-1}, I_{res}), \quad (8)$$

$$q(I_t|I_{t-1}, I_{res}) := \mathcal{N}(I_t; I_{t-1} + \alpha_t I_{res}, \beta_t^2 \mathbf{I}). \quad (9)$$

Eq. 7 defines the marginal probability distribution $q(I_t|I_0, I_{res}) = \mathcal{N}(I_t; I_0 + \bar{\alpha}_t I_{res}, \bar{\beta}_t^2 \mathbf{I})$. In fact, the forward diffusion of our RDDM is a mixture of three terms (i.e., I_0 , I_{res} , and ϵ), extending beyond the widely used denoising diffusion model that is a mixture of two terms, i.e., I_0 and ϵ . A similar mixture form of three terms can be seen in several concurrent works, e.g., InDI [11], I2SB [29], IR-SDE [40], and ResShift [67].

4.2. Generation Process and Training Objective

In the forward process (Eq. 7), residuals (I_{res}) and noise (ϵ) are gradually added to I_0 , and then synthesized into I_t , while the reverse process from I_T to I_0 involves the estimation of the residuals and noise injected in the forward process. We can train a residual network $I_{res}^\theta(I_t, t, I_{in})$ to predict I_{res} and a noise network $\epsilon_\theta(I_t, t, I_{in})$ to estimate ϵ . Using Eq. 7, we obtain the estimated target images $I_0^\theta = I_t - \bar{\alpha}_t I_{res}^\theta - \bar{\beta}_t \epsilon_\theta$. If I_0^θ and I_{res}^θ are given, the generation process is defined as,

$$p_\theta(I_{t-1}|I_t) := q_\sigma(I_{t-1}|I_t, I_0^\theta, I_{res}^\theta), \quad (10)$$

where the transfer probability $q_\sigma(I_{t-1}|I_t, I_0, I_{res})^2$ from I_t to I_{t-1} is,

$$q_\sigma(I_{t-1}|I_t, I_0, I_{res}) = \mathcal{N}(I_{t-1}; I_0 + \bar{\alpha}_{t-1} I_{res} + \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2} \frac{I_t - (I_0 + \bar{\alpha}_t I_{res})}{\bar{\beta}_t}, \sigma_t^2 \mathbf{I}), \quad (11)$$

where $\sigma_t^2 = \eta \bar{\beta}_t^2 \bar{\beta}_{t-1}^2 / \bar{\beta}_t^2$ and η controls whether the generation process is random ($\eta = 1$) or deterministic ($\eta = 0$). Using Eq. 10 and Eq. 11, I_{t-1} can be sampled from I_t via:

$$I_{t-1} = I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{res}^\theta - (\bar{\beta}_t - \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}) \epsilon_\theta + \sigma_t \epsilon_t, \quad (12)$$

where $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. When $\eta = 1$, our RDDM has the sum-constrained variance, while DDPM has preserving variance (see Appendix A.4). When $\eta = 0$ (i.e., $\sigma_t = 0$), the sampling process is deterministic,

$$I_{t-1} = I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{res}^\theta - (\bar{\beta}_t - \bar{\beta}_{t-1}) \epsilon_\theta. \quad (13)$$

²Eq. 11 does not change $q(I_t|I_0, I_{res})$ in Appendix A.2.

We derive the following simplified loss function for training (Appendix A.1):

$$L_{res}(\theta) := \mathbb{E} \left[\lambda_{res} \|I_{res} - I_{res}^\theta(I_t, t, I_{in})\|^2 \right], \quad (14)$$

$$L_\epsilon(\theta) := \mathbb{E} \left[\lambda_\epsilon \|\epsilon - \epsilon_\theta(I_t, t, I_{in})\|^2 \right], \quad (15)$$

where the hyperparameters $\lambda_{res}, \lambda_\epsilon \in \{0, 1\}$, and the training input image I_t is synthesized using I_0 , I_{res} , and ϵ by Eq. 7. I_t can also be synthesized using I_{in} (replace I_0 in Eq. 7 by $I_0 = I_{in} - I_{res}$),

$$I_t = I_{in} + (\bar{\alpha}_t - 1) I_{res} + \bar{\beta}_t \epsilon. \quad (16)$$

4.3. Sampling Method Selection Strategies

For the generation process (from I_t to I_{t-1}), I_t and I_{in} are known, and thus I_{res} and ϵ can represent each other by Eq. 16. From Eq. 14, 15, 16, we propose three sampling methods as follows.

SM-Res. When $\lambda_{res} = 1$ and $\lambda_\epsilon = 0$, the residuals I_{res}^θ are predicted by a network, while the noise ϵ_θ is represented as a transformation of I_{res}^θ using Eq. 16.

SM-N. When $\lambda_{res} = 0$ and $\lambda_\epsilon = 1$, the noise ϵ_θ is predicted by a network, while the residuals I_{res}^θ are represented as a transformation of ϵ_θ using Eq. 16.

SM-Res-N. When $\lambda_{res} = 1$ and $\lambda_\epsilon = 1$, both the residuals and the noise are predicted by networks.

To determine the optimal sampling method for real-world applications, we give empirical strategies and automatic selection algorithms in the following.

Empirical Research. Table 1 presents that the SM-Res shows better results for image restoration but offers a poorer FID for generation. On the other hand, the SM-N yields better frechet inception distance (FID in [16]) and inception scores (IS), but is ineffective in image restoration (e.g., PSNR 11.34 for shadow and 16.30 for low-light). This may be due to the inadequacy of using ϵ_θ to represent I_{res}^θ in Eq. 16 for restoration tasks. We attribute these inconsistent results to the fact that **residual predictions prioritize certainty, whereas noise predictions emphasize diversity**. In our experiments, we use SM-N for image generation, SM-Res for low-light (LOL [61]), and SM-Res-N for other image restoration tasks. For an unknown new task, we empirically recommend using SM-N for those requiring greater diversity and SM-Res for tasks that demand higher certainty.

Automatic Objective Selection Algorithm (AOSA). To automatically choose between SM-Res or SM-N for an unknown task, we develop an automatic sampling selection algorithm in Appendix B.2. This algorithm requires only a single network and learns the hyperparameter in Eq. 15, enabling a gradual transition from combined residual and noise training (akin to SM-Res-N) to individual prediction (SM-Res or SM-N). This plug-and-play training strategy requires less than 1000 additional training iterations and is

Sampling Method	Generation (CelebA)		Shadow removal (ISTD)			Low-light (LOL)		Deraining (RainDrop)	
	FID (↓)	IS (↑)	MAE(↓)	PSNR(↑)	SSIM(↑)	PSNR(↑)	SSIM(↑)	PSNR(↑)	SSIM(↑)
SM-Res	31.47	1.73	<u>4.76</u>	<u>30.72</u>	<u>0.959</u>	25.39	0.937	<u>31.96</u>	<u>0.9509</u>
SM-N	23.25	2.05	81.01	11.34	0.175	16.30	0.649	19.15	0.7179
SM-Res-N	<u>28.90</u>	<u>1.78</u>	4.67	30.91	0.962	<u>23.90</u>	<u>0.931</u>	32.51	0.9563

Table 1. Sampling method analysis. The sampling steps are 10 on the CelebA 64×64 [36] dataset, 5 on the ISTD [57] dataset, 2 on the LOL [61] dataset, and 5 on the RainDrop [45] dataset.

fully compatible with the current denoising-based diffusion methods [17]. Our RDDM using AOSA has the potential to provide a unified and interpretable methodology for modeling, training, and inference pipelines for unknown target tasks.

Comparison with Other Prediction Methods. Our SM-N is similar to DDIM [51] (or DDPM [17]), which only estimates the noise, and is consistent with DDPM and DDIM by transforming the coefficient/variance schedules in Eq. 12 (the proof in Appendix A.3),

$$\bar{\alpha}_t = 1 - \sqrt{\bar{\alpha}_{DDIM}^t}, \bar{\beta}_t = \sqrt{1 - \bar{\alpha}_{DDIM}^t}, \quad (17)$$

$$\sigma_t^2 = \sigma_t^2(DDIM).$$

In fact, current research has delved into numerous diffusion forms that extend beyond noise estimation. For example, IDDPM [44] proposes that it is feasible to estimate noise (ϵ_θ), clean target images (I_0^θ), or the mean term (μ_θ) to represent the transfer probabilities (i.e., $p_\theta(I_{t-1}|I_t)$ in Eq. 5). The score-based generative model (SGM) [52] and Schrödinger Bridge (I2SB [29]) estimate the score of noisy data (i.e., the sum of residuals and noise $\sum_{i=1}^t I_{res}^i$). Cold-Diffusion [2] and InDI [11] estimate the clean target images (I_0). Rectified Flow [35] predicts the residuals (I_{res}) to align with the image linear interpolation process without noise diffusion (i.e., $I_T = I_{in}$). A detailed comparison can be found in Appendix A.5.

These previous/concurrent works choose to estimate the noise, the residual, the target image, or its linear transformation term. In contrast, we introduce residual estimation while also embracing noise for both generation and restoration. Residuals and noise have equal and independent status, which is reflected in the forward process (Eq. 7), the reverse process (Eq. 13), and the loss function (Eq. 15). This independence means that the noise diffusion can even be removed and only the residual diffusion retained to model the image interpolation process (when $\bar{\beta}_T = 0$ in Eq. 7, RDDM degenerates to Rectified Flow [35]). In addition, this property derives a decoupled dual diffusion framework in Section 5.

5. Decoupled Dual Diffusion Framework

Upon examining DDPM from the perspective of RDDM, we discover that DDPM indeed involves the simultaneous

³ $\bar{\alpha}_{DDIM}^t$ here is α_t of DDIM [51].

Schedules	FID (↓)	IS (↑)
Linear (DDIM [51])	28.39 ⁴	2.05
Scaled linear [48]	28.15	2.00
Squared cosine [44]	47.21	2.64
α_t (mean), β_t^2 (mean)	38.35	2.22
α_t (linearly increasing), β_t^2 (linearly increasing)	40.03	<u>2.45</u>
α_t (linearly decreasing), β_t^2 (linearly decreasing)	<u>27.82</u>	2.26
α_t (linearly decreasing), β_t^2 (linearly increasing)	23.25	2.05

Table 2. Coefficient schedules analysis on CelebA (64×64) [36]. In our RDDM, the residual diffusion and noise diffusion are decoupled, so one may design a better schedule in the decoupled coefficient space, e.g., α_t (linearly decreasing), β_t^2 (linearly increasing). To be fair, all coefficient schedules were retrained using the same network structure, training, and evaluation. The sampling method is SM-N with 10 sampling steps using Eq. 13.

diffusion of residuals and noise, which is evident as Eq. 48 becomes equivalent to Eq. 44 in Appendix A.3. We find that it is possible to decouple these two types of diffusion. Section 5.1 presents a decoupled forward diffusion process. In Section 5.2, we propose a partially path-independent generation process and decouple the simultaneous sampling into first removing the residuals and then removing noise (see Fig. 6(d) and Fig. 17). This decoupled dual diffusion framework sheds light on the roles of deresidual and denoising in the DDPM generation process.

5.1. Decoupled Forward Diffusion Process

Our defined coefficients (α_t, β_t^2) offer a distinct physical interpretation. In the forward diffusion process (Eq. 7), α_t controls the speed of residual diffusion and β_t^2 regulates the speed of noise diffusion. In the reverse generation process (Eq. 13), $\bar{\alpha}_t$ and $\bar{\beta}_t$ are associated with the speed of removing residual and noise, respectively. In fact, there are no constraints on α^t and β_t^2 in Eq. 7, meaning that the residual diffusion and noise diffusion are independent of each other. Utilizing this decoupled property and the difference between these two diffusion processes, we should be able to design a better coefficient schedule, e.g., α_t (linearly decreasing) and β_t^2 (linearly increasing) in Table 2. This aligns with the intuition that, during the reverse generation pro-

⁴Our RDDM is implemented based on the popular diffusion repository github.com/lucidrains/denoising-diffusion-pytorch. Differences in network structure and training details may lead to poorer FID. We have verified sampling consistency with DDIM [51] in Table 3(a) and Appendix A.3.

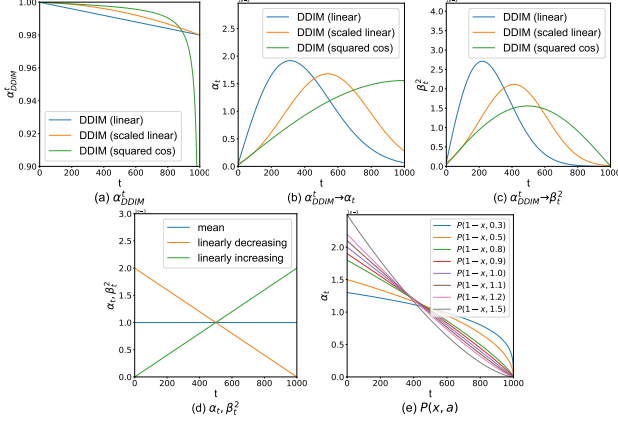


Figure 4. Coefficient transformation from DDIM [51] to RDDM using Eq. 17. (a) We show several schedules for α_{DDIM}^t , e.g., linear [51], scaled linear [48], and squared cosine [44]. (b) We transform α_{DDIM}^t into α_t in our RDDM. (c) We transform α_{DDIM}^t into β_t^2 in our RDDM. (d) A few simple schedules. (e) $P(x, a)$ is a normalized power function (see Eq. 18). "mean", "linearly increasing", and "linearly decreasing" in (d) can be denoted as $P(x, 0)$, $P(x, 1)$ and $P(1 - x, 1)$, respectively. See Algorithm 1 in Appendix A.3 for more details of (b) and (c).

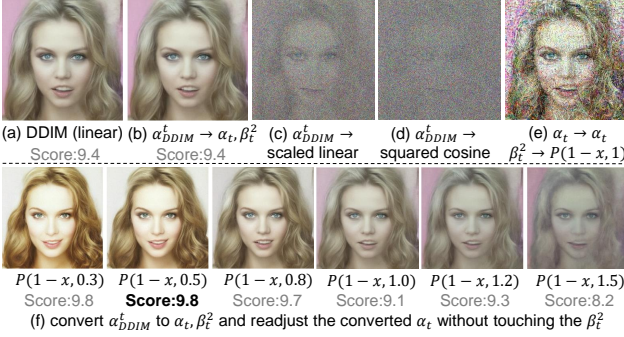


Figure 5. Analysis of readjusting coefficient schedules. We find that changing the α_t schedule barely affects the denoising process in (f) and edited faces may have higher face scores when assessed using AI face scoring software⁵. These images were generated using a pre-trained UNet on the CelebA (256 × 256) dataset [36] with 10 sampling steps.

cess (from T to 0), the estimated residuals become increasingly accurate while the estimated noise should also weaken progressively. Therefore, when t is close to 0, the deresidual pace should be faster and the denoising pace should be slower. Since our α_t and β_t^2 represent the speed of diffusion, we name the curve in Fig. 4 (b-d) the *diffusion speed curve*.

5.2. Partially Path-independent Generation Process

In the original DDPM [17] or DDIM [51], when the α_{DDIM}^t schedule changes, it is necessary to retrain the denoising network because this alters the diffusion process [44, 48]. As shown in Fig. 5(c)(d), directly changing the α_{DDIM}^t schedule causes denoising to fail. Here, we

propose a path-independent generation process, i.e., modifying the diffusion speed curve does not cause the image generation process to fail. We try to readjust the diffusion speed curve in the generation process. First, we convert the α_{DDIM}^t schedule of a pre-trained DDIM into the α_t and β_t^2 schedules of our RDDM using Eq. 17 (from Fig. 5(a) to Fig. 5(b)). We then readjust the converted α_t schedules using the normalized power function ($P(x, a)$ in Fig. 5(f)), without touching the β_t^2 schedule that controls noise diffusion, as shown in Fig. 5(f). $P(x, a)$ is defined as (a is a parameter of the power function),

$$P(x, a) := x^a / \int_0^1 x^a dx, \text{ where } x = t/T. \quad (18)$$

These schedule modifications shown in Fig. 5 lead to the following key findings.

1. Fig. 5(f) shows that modifying the residual diffusion speed curve (α_t) leads to a drastic change in the generation results, probably due to I_{res}^θ being represented as a transformation of ϵ_θ using Eq. 16.
2. As the time condition t represents the current noise intensity in the denoising network ($\epsilon_\theta(I_t, t, 0)$), modifying the noise diffusion speed curve (β_t^2) causes t to deviate from accurately indicating the current noise intensity, leading to denoising failure, as shown in Fig. 5(e).

Nonetheless, we believe that, corresponding to the decoupled forward diffusion process, there should also be a path-independent reverse generation process. To develop a path-independent generation process, we improve the generation process based on the above two key findings:

1. Two networks are used to estimate I_{res}^θ and ϵ_θ separately, i.e., SM-Res-N-2Net in Appendix B.2.
2. $\bar{\alpha}_t$ and $\bar{\beta}_t$ are used for the time conditions embedded in the network, i.e., $I_{res}^\theta(I_t, t, 0) \rightarrow I_{res}^\theta(I_t, \bar{\alpha}_t \cdot T, 0)$, $\epsilon_\theta(I_t, t, 0) \rightarrow \epsilon_\theta(I_t, \bar{\beta}_t \cdot T, 0)$.

These improvements lead to a partially path-independent generation process, as evidenced by the results shown in Fig. 6(c).

Analysis of Partially Path-independence via Green’s Theorem. “Path-independence” reminds us of Green’s theorem in curve integration [47]. From Eq. 13, we have:

$$I_t - I_{t-1} = (\bar{\alpha}_t - \bar{\alpha}_{t-1})I_{res}^\theta + (\bar{\beta}_t - \bar{\beta}_{t-1})\epsilon_\theta, \quad (19)$$

$$dI(t) = I_{res}^\theta(I(t), \bar{\alpha}(t) \cdot T, 0)d\bar{\alpha}(t) + \epsilon_\theta(I(t), \bar{\beta}(t) \cdot T, 0)d\bar{\beta}(t), \quad (20)$$

where $I(t) = I(0) + \bar{\alpha}(t)I_{res} + \bar{\beta}(t)\epsilon$. Given inputs $I(t)$ and $\bar{\alpha}(t)$, the denoising network learns to approximate the noise ϵ in $I(t)$ by estimating ϵ_θ . If this network is trained well and robust enough, it should be able to avoid the interference of

⁵<https://ux.xiaoice.com/beautyv3>

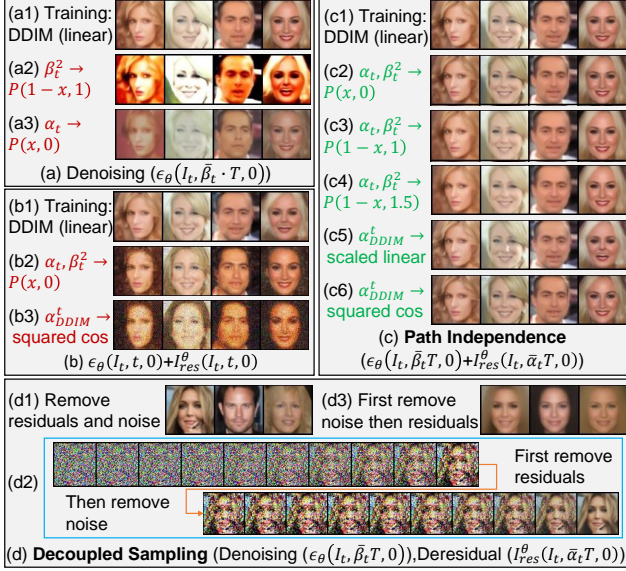


Figure 6. Partially path-independent generation process. (a1) We trained a denoising network using the DDIM linear schedule [51]. (a2-a3) We modified the α_t and β_t^2 schedules during testing. (b) We trained two networks to remove noise and residuals. In contrast to the sharply varying images in (a2-a3) and the noisy images in (b2-b3), (c) shows that we constructed a path independent generation process where modifications to the diffusion speed curve can generate a noise-free image with little variation in image semantics. (d) The simultaneous sampling in (d1) or (c) can be decomposed into first removing residuals and then noise (d2), or removing noise and then residuals (d3). In (d3), diversity is significantly reduced because noise is removed first.

the residual terms $\bar{\alpha}(t)I_{res}$ in $I(t)$. This also applies to a robust residual estimation network. Thus, we have

$$\frac{\partial I_{res}^\theta(I(t), \bar{\alpha}(t) \cdot T)}{\partial \bar{\beta}(t)} \approx 0, \frac{\partial \epsilon_\theta(I(t), \bar{\beta}(t) \cdot T)}{\partial \bar{\alpha}(t)} \approx 0. \quad (21)$$

If the equation in Formula 21 holds true, it serves as a necessary and sufficient condition for path independence in curve integration, which provides an explanation for why Fig. 6(c) achieves a partially path-independent generation process. The path-independent property is related to the network’s resilience to disturbances and applies to disturbances that vary within a certain range. However, excessive disturbances can lead to visual inconsistencies, e.g., readjusting α_t and β_t^2 to $P(x, 5)$. Thus, we refer to this generative property as partially path-independent. We also investigated two reverse paths to gain insight into the implications of the proposed partial path independence. In the first case, the residuals are removed first, followed by the noise: $I(T) \xrightarrow{-I_{res}} I(0) + \bar{\beta}_T \epsilon \xrightarrow{-\bar{\beta}_T \epsilon} I(0)$. The second case involves removing the noise first and then the residuals: $I(T) \xrightarrow{-\bar{\beta}_T \epsilon} I_{in} \xrightarrow{-I_{res}} I(0)$. The first case (Fig. 6(d2)) shows

(a) CelebA (FID)	5 steps	10 steps	15 steps	20 steps	100 steps
DDIM	69.60	40.45	32.67	30.61	23.66
DDIM→RDDM	69.60	40.41	32.71	30.77	24.92

(b) Shadow Removal	MAE(↓)			SSIM(↑)			PSNR(↑)		
	S	NS	ALL	S	NS	ALL	S	NS	ALL
DSC [19] ¶	9.48	6.14	6.67	0.967	-	-	33.45	-	-
FusionNet [13]	7.77	5.56	5.92	0.975	0.880	0.945	34.71	28.61	27.19
BMNet [79]	7.60	4.59	5.02	0.988	0.976	0.959	35.61	32.80	30.28
DMTN [31]	7.00	4.28	4.72	0.990	0.979	0.965	35.83	33.01	30.42
Ours (RDDM)	6.67	4.27	4.67	0.988	0.979	0.962	36.74	33.18	30.91

(c) Low-light	PSNR(↑)	SSIM(↑)	LPIPS (↓)	(d) Deraining	PSNR(↑)	SSIM(↑)
KinD++ [76]	17.752	0.760	0.198	AttnGAN [45]	31.59	0.9170
KinD++-SKF [68]	20.363	0.805	0.201	DuRN [34]	31.24	0.9259
DCC-Net [77]	22.72	0.81	-	RainAttn [46]	31.44	0.9263
SNR-Aware [66]	24.608	0.840	0.151	IDT [64]	31.87	0.9313
LLFlow [59]	<u>25.19</u>	<u>0.93</u>	0.11	RainDiff64 [82]	32.29	<u>0.9422</u>
LLFormer [58]	23.649	0.816	0.169	RainDiff128 [82]	32.43	0.9334
Ours (RDDM)	25.392	0.937	<u>0.116</u>	Ours (RDDM)	32.51	0.9563

Table 3. Quantitative comparison results of image generation on the CelebA (256×256) dataset [36], shadow removal on the ISTD dataset [57], low-light enhancement on the LOL [61] dataset, and deraining on the RainDrop [45] dataset. “S, NS, ALL” in (b) denote shadow area (S), non-shadow area (NS) and whole image (ALL). The sampling steps are 5 for shadow removal and deraining, 2 for low-light.

that removing residuals controls semantic transitions, while the second case (Fig. 6(d3)) shows that diversity is significantly reduced because noise is removed first. Fig. 6(d) validates our argument that residuals control directional semantic drift (certainty) and noise controls random perturbation (diversity). See Appendix B.4 for more details.

6. Experiments

Image Generation. We can convert a pre-trained⁶ DDIM [51] to RDDM by coefficient transformation using Eq. 17, and generate images by Eq. 12. Table 3(a) verifies that the quality of the generated images before and after the conversion is nearly the same⁷. We show the generated face images with 10 sampling steps in Fig. 7(a).

Image Restoration. We extensively evaluate our method on several image restoration tasks, including shadow removal, low-light enhancement, deraining, and deblurring on 5 datasets. Notably, our RDDM uses an identical UNet and is trained with a batch size of 1 for all these tasks. In contrast, SOAT methods often involve elaborate network architectures, such as multi-stage [13, 59, 80], multi-branch [10], Transformer [58], and GAN [27], or sophisticated loss functions like the chromaticity [20], texture similarity [74], and edge loss [70]. Table 3 and Fig. 7(b-c) show that our RDDM is competitive with the SOTA restora-

⁶<https://huggingface.co/google/ddpm-celebahq-256>

⁷The subtle differences in larger sampling steps may stem from errors introduced by numerical representation limitations during coefficient transformation, which may accumulate and amplify in larger sampling steps.



Figure 7. Application of our RDDM. (a) Image generation on the CelebA dataset [36]. (b) Shadow removal on the ISTD dataset [57]. (c) Low-light enhancement on the LOL dataset [61]. (d) Image inpainting (center and irregular mask). (e) The image translation process can be regarded as first translating the semantics and then generating the details. These images in (b) are magnified using MulingViewer [30].

tion methods. See Appendix B for more training details and comparison results.

We extend DDPM [17]/DDIM [51], initially uninterpretable for image restoration, into a unified and interpretable diffusion model for both image generation and restoration by introducing residuals. However, the residual diffusion process represents the directional diffusion from target images to conditional input images, which does not involve a priori information about the image restoration task, and therefore is not limited to it. Beyond image generation and restoration, we show examples of image inpainting and image translation to verify that our RDDM has the potential to be a unified and interpretable methodology for image-to-image distribution transformation. **We do not intend to achieve optimal performance on all tasks by tuning all hyperparameters.** The current experimental results show that RDDM 1) achieves consistent **image generation** performance with DDIM after coefficient transformation, 2) competes with state-of-the-art **image restoration** methods using a generic UNet with only an ℓ_1 loss, a batch size of 1, and fewer than 5 sampling steps, and 3) has satisfactory visual results of **image inpainting** and **image translation** (see Fig. 7(d-e), Fig. 14, or Fig. 15 in Appendix B.3), which

validates our RDDM.

7. Conclusions

We present a unified dual diffusion model called Residual Denoising Diffusion Models (RDDM) for image restoration and image generation. This is a three-term mixture framework beyond the previous denoising diffusion framework with two-term mixture. We demonstrate that our sampling process is consistent with that of DDPM and DDIM through coefficient schedule transformation, and propose a partially path-independent generation process. Our experimental results on four different image restoration tasks show that RDDM achieves SOTA performance in no more than five sampling steps. We believe that our model and framework hold the potential to provide a unified methodology for image-to-image distribution transformation and pave the way for the multi-dimensional diffusion process.

8. Acknowledgments

This work was supported by National Natural Science Foundation of China under Grants 61991413, 62273339, 62073205, 62306253.

References

- [1] Saeed Anwar and Nick Barnes. Densely residual laplacian super-resolution. *IEEE TPAMI*, 44(3):1192–1204, 2020. [2](#)
- [2] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022. [2](#), [5](#), [16](#)
- [3] Fan Bao, Chongxuan Li, Jiacheng Sun, Jun Zhu, and Bo Zhang. Estimating the optimal covariance with imperfect mean in diffusion probabilistic models. In *Proc. ICML*, 2022. [2](#)
- [4] Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. In *Proc. ICLR*, 2022. [2](#)
- [5] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006. [13](#)
- [6] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input / output image pairs. In *Proc. CVPR*, 2011. [21](#)
- [7] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proc. CVPR*, pages 3291–3300, 2018. [21](#)
- [8] Zipei Chen, Chengjiang Long, Ling Zhang, and Chunxia Xiao. CANet: A Context-Aware Network for Shadow Removal. In *Proc. ICCV*, pages 4723–4732, 2021. [19](#), [20](#)
- [9] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proc. CVPR*, pages 8188–8197, 2020. [23](#), [28](#)
- [10] Xiaodong Cun, Chi Man Pun, and Cheng Shi. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting GAN. In *Proc. AAAI*, pages 10680–10687, 2020. [7](#), [17](#), [19](#), [20](#)
- [11] Mauricio Delbracio and Peyman Milanfar. Inversion by direct iteration: An alternative to denoising diffusion for image restoration. *Transactions on Machine Learning Research*, 2023. [2](#), [4](#), [5](#), [16](#), [23](#), [26](#)
- [12] Huiyu Duan, Wei Shen, Xiongkuo Min, Yuan Tian, Jae-Hyun Jung, Xiaokang Yang, and Guangtao Zhai. Develop then rival: A human vision-inspired framework for superimposed image decomposition. *IEEE TMM*, 2022. [2](#)
- [13] Lan Fu, Changqing Zhou, Qing Guo, Felix Juefei-Xu, Hongkai Yu, Wei Feng, Yang Liu, and Song Wang. Auto-Exposure Fusion for Single-Image Shadow Removal. In *Proc. CVPR*, pages 10566–10575, 2021. [7](#), [17](#), [19](#), [20](#)
- [14] Lanqing Guo, Chong Wang, Wenhan Yang, Siyu Huang, Yufei Wang, Hanspeter Pfister, and Bihan Wen. Shadowdiffusion: When degradation prior meets diffusion model for shadow removal. In *Proc. CVPR*, 2023. [2](#), [16](#), [28](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016. [2](#), [3](#)
- [16] 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. In *Proc. NeurIPS*, 2017. [4](#)
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proc. NeurIPS*, pages 6840–6851, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#), [12](#), [13](#), [14](#), [15](#), [16](#), [19](#)
- [18] Weipeng Hu and Haifeng Hu. Domain discrepancy elimination and mean face representation learning for nir-vis face recognition. *IEEE Signal Processing Letters*, 28:2068–2072, 2021. [26](#)
- [19] Xiaowei Hu, Chi Wing Fu, Lei Zhu, Jing Qin, and Pheng Ann Heng. Direction-aware spatial context features for shadow detection and removal. *IEEE TPAMI*, 42(11): 2795–2808, 2020. [7](#), [19](#), [20](#)
- [20] Yeying Jin, Aashish Sharma, and Robby T. Tan. DC-ShadowNet: Single-Image Hard and Soft Shadow Removal Using Unsupervised Domain-Classifer Guided Network. In *Proc. ICCV*, pages 5007–5016, 2021. [7](#), [17](#)
- [21] Yeying Jin, Wenhan Yang, Wei Ye, Yuan Yuan, and Robby T Tan. Des3: Attention-driven self and soft shadow removal using vit similarity and color convergence. *arXiv preprint arXiv:2211.08089*, 2022. [25](#)
- [22] Yeying Jin, Wenhan Yang, Wei Ye, Yuan Yuan, and Robby T Tan. Shadowdiffusion: Diffusion-based shadow removal using classifier-driven attention and structure preservation. *arXiv preprint arXiv:2211.08089*, 2022. [1](#)
- [23] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *Proc. ICLR*, 2018. [23](#), [26](#), [28](#)
- [24] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In *Proc. NeurIPS*, pages 21696–21707, 2021. [2](#)
- [25] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *Proc. ICLR*, 2014. [3](#)
- [26] Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *Found. Trends Mach. Learn.*, 12(4):307–392, 2019. [3](#)
- [27] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *Proc. ICCV*, pages 8878–8887, 2019. [7](#), [17](#), [23](#)
- [28] Hieu Le and Dimitris Samaras. Shadow removal via shadow image decomposition. In *Proc. ICCV*, pages 8578–8587, 2019. [17](#)
- [29] Guan-Hong Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. I²sb: Image-to-image schrödinger bridge. In *Proc. ICML*, 2023. [2](#), [4](#), [5](#), [16](#), [23](#)
- [30] Jiawei Liu. MulimgViewer: A multi-image viewer for image comparison and image stitching. [8](#)
- [31] Jiawei Liu, Qiang Wang, Huijie Fan, Wentao Li, Liangqiong Qu, and Yandong Tang. A decoupled multi-task network for shadow removal. *IEEE TMM*, 2023. [3](#), [7](#), [19](#), [20](#)
- [32] Jiawei Liu, Qiang Wang, Huijie Fan, Jiandong Tian, and Yandong Tang. A shadow imaging bilinear model and three-branch residual network for shadow removal. *IEEE TNNLS*, pages 1–15, 2023. [2](#), [3](#)
- [33] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In *Proc. CVPR*, pages 10561–10570, 2021. [21](#)

- [34] Xing Liu, Masanori Suganuma, Zhun Sun, and Takayuki Okatani. Dual residual networks leveraging the potential of paired operations for image restoration. In *Proc. CVPR*, pages 7007–7016, 2019. [7](#)
- [35] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *Proc. ICLR*, 2023. [5](#), [16](#)
- [36] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proc. ICCV*, 2015. [5](#), [6](#), [7](#), [8](#), [15](#)
- [37] Zhihao Liu, Hui Yin, Yang Mi, Mengyang Pu, and Song Wang. Shadow removal by a lightness-guided network with training on unpaired data. *IEEE TIP*, 30:1853–1865, 2021. [19](#), [20](#)
- [38] Gunter Loffler, Grigori Yourganov, Frances Wilkinson, and Hugh R Wilson. fmri evidence for the neural representation of faces. *Nature neuroscience*, 8(10):1386–1391, 2005. [26](#)
- [39] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proc. CVPR*, pages 11461–11471, 2022. [1](#), [2](#)
- [40] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Image restoration with mean-reverting stochastic differential equations. In *Proc. ICML*, 2023. [4](#)
- [41] 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. In *Proc. ICLR*, 2021. [2](#)
- [42] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proc. CVPR*, pages 14225–14234, 2021. [26](#)
- [43] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proc. CVPR*, pages 3883–3891, 2017. [22](#), [23](#), [28](#)
- [44] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proc. ICML*, pages 8162–8171, 2021. [2](#), [5](#), [6](#), [16](#)
- [45] 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 *Proc. CVPR*, pages 2482–2491, 2018. [5](#), [7](#), [17](#), [22](#), [23](#), [28](#)
- [46] Yuhui Quan, Shijie Deng, Yixin Chen, and Hui Ji. Deep learning for seeing through window with raindrops. In *Proc. ICCV*, pages 2463–2471, 2019. [7](#)
- [47] K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*. Cambridge University Press, 3 edition, 2006. [2](#), [6](#)
- [48] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. CVPR*, pages 10684–10695, 2022. [1](#), [2](#), [5](#), [6](#)
- [49] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI*, 45(4):4713–4726, 2022. [1](#), [2](#), [16](#), [27](#), [28](#)
- [50] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. ICML*, pages 2256–2265, 2015. [2](#)
- [51] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proc. ICLR*, 2021. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#), [13](#), [14](#), [15](#), [16](#), [17](#), [24](#)
- [52] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *32*, 2019. [2](#), [5](#)
- [53] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. ICLR*, 2021. [2](#), [15](#), [16](#)
- [54] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *Proc. ICML*, 2023. [29](#)
- [55] Maitreya Suin, Kuldeep Purohit, and AN Rajagopalan. Spatially-attentive patch-hierarchical network for adaptive motion deblurring. In *Proc. CVPR*, pages 3606–3615, 2020. [23](#)
- [56] Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxim: Multi-axis mlp for image processing. In *Proc. CVPR*, pages 5769–5780, 2022. [2](#)
- [57] Jifeng Wang, Xiang Li, and Jian Yang. Stacked Conditional Generative Adversarial Networks for Jointly Learning Shadow Detection and Shadow Removal. In *Proc. CVPR*, pages 1788–1797, 2018. [5](#), [7](#), [8](#), [17](#), [18](#), [19](#), [20](#), [26](#), [27](#), [28](#)
- [58] Tao Wang, Kaihao Zhang, Tianrun Shen, Wenhan Luo, Bjorn Stenger, and Tong Lu. Ultra-high-definition low-light image enhancement: A benchmark and transformer-based method. In *Proc. AAAI*, pages 2654–2662, 2023. [7](#), [27](#)
- [59] Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex Kot. Low-light image enhancement with normalizing flow. In *Proc. AAAI*, pages 2604–2612, 2022. [7](#), [27](#)
- [60] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proc. CVPR*, pages 17683–17693, 2022. [23](#), [24](#), [28](#)
- [61] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *Proc. BMVC*, 2018. [4](#), [5](#), [7](#), [8](#), [20](#), [21](#), [28](#)
- [62] Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alexandros G Dimakis, and Peyman Milanfar. Deblurring via stochastic refinement. In *Proc. CVPR*, pages 16293–16303, 2022. [2](#), [16](#), [23](#)
- [63] Hugh R Wilson, Gunter Loffler, and Frances Wilkinson. Synthetic faces, face cubes, and the geometry of face space. *Vision research*, 42(27):2909–2923, 2002. [26](#)
- [64] Jie Xiao, Xueyang Fu, Aiping Liu, Feng Wu, and Zheng-Jun Zha. Image de-raining transformer. *IEEE TPAMI*, 2022. [7](#)
- [65] Ke Xu, Xin Yang, Baocai Yin, and Rynson WH Lau. Learning to restore low-light images via decomposition-and-enhancement. In *Proc. CVPR*, pages 2281–2290, 2020. [17](#), [21](#), [23](#), [28](#)
- [66] Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *Proc. CVPR*, pages 17714–17724, 2022. [7](#), [21](#), [22](#)

- [67] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. In *Proc. NeurIPS*, 2023. 4
- [68] Wu Yuhui, Pan Chen, Wang Guoqing, Yang Yang, Wei Jiwei, Li Chongyi, and Heng Tao Shen. Learning semantic-aware knowledge guidance for low-light image enhancement. In *Proc. CVPR*, 2023. 7, 21
- [69] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. In *Proc. ECCV*, pages 492–511, 2020. 21
- [70] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proc. CVPR*, pages 14821–14831, 2021. 2, 7, 17, 23
- [71] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for fast image restoration and enhancement. *IEEE TPAMI*, 45(2):1934–1948, 2022. 21
- [72] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE TIP*, 26(7):3142–3155, 2017. 2
- [73] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, pages 586–595, 2018. 17, 21
- [74] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *Proc. ACM MM*, pages 1632–1640, 2019. 7, 17
- [75] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image restoration. *IEEE TPAMI*, 43(7):2480–2495, 2020. 2
- [76] Yonghua Zhang, Xiaojie Guo, Jiayi Ma, Wei Liu, and Jiawan Zhang. Beyond brightening low-light images. *IJCV*, 129: 1013–1037, 2021. 7, 21
- [77] Zhao Zhang, Huan Zheng, Richang Hong, Mingliang Xu, Shuicheng Yan, and Meng Wang. Deep color consistent network for low-light image enhancement. In *Proc. CVPR*, pages 1899–1908, 2022. 7
- [78] Chuanjun Zheng, Daming Shi, and Wentian Shi. Adaptive unfolding total variation network for low-light image enhancement. In *Proc. ICCV*, pages 4439–4448, 2021. 21
- [79] Yurui Zhu, Jie Huang, Xueyang Fu, Feng Zhao, Qibin Sun, and Zheng-Jun Zha. Bijective mapping network for shadow removal. In *Proc. CVPR*, pages 5627–5636, 2022. 7, 17, 19, 20
- [80] Yurui Zhu, Zeyu Xiao, Yanchi Fang, Xueyang Fu, Zhiwei Xiong, and Zheng-Jun Zha. Efficient model-driven network for shadow removal. In *Proc. AAAI*, 2022. 7, 17, 19, 21
- [81] Zhengxia Zou, Sen Lei, Tianyang Shi, Zhenwei Shi, and Jieping Ye. Deep adversarial decomposition: A unified framework for separating superimposed images. In *Proc. CVPR*, pages 12806–12816, 2020. 2
- [82] Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE TPAMI*, pages 1–12, 2023. 1, 2, 7, 16, 23

Appendix

A. Derivations and Proofs

A.1. Perturbed Generation Process

For the reverse generation process from I_t to I_{t-1} , we can represent the transfer probabilities $q(I_{t-1}|I_t, I_0, I_{res})$ by Bayes' rule:

$$q(I_{t-1}|I_t, I_0, I_{res}) = q(I_t|I_{t-1}, I_0, I_{res}) \frac{q(I_{t-1}|I_0, I_{res})}{q(I_t|I_0, I_{res})}, \quad (22)$$

where $q(I_{t-1}|I_0, I_{res}) = \mathcal{N}(I_{t-1}; I_0 + \bar{\alpha}_{t-1}I_{res}, \bar{\beta}_{t-1}^2\mathbf{I})$ from Eq. 7, and $q(I_t|I_{t-1}, I_0, I_{res}) = q(I_t|I_{t-1}, I_{res})^8 = \mathcal{N}(I_t; I_{t-1} + \alpha_t I_{res}, \beta_t^2\mathbf{I})$ from Eq. 9. Thus, we have (considering only the exponential term)

$$\begin{aligned} q(I_{t-1}|I_t, I_0, I_{res}) &= \mathcal{N}(I_{t-1}; \mu_t(x_t, I_0, I_{res}), \Sigma_t(x_t, I_0, I_{res})\mathbf{I}) \\ &\propto \exp\left(-\frac{1}{2}\left(\frac{(I_t - I_{t-1} - \alpha_t I_{res})^2}{\beta_t^2} + \frac{(I_{t-1} - I_0 - \bar{\alpha}_{t-1}I_{res})^2}{\bar{\beta}_{t-1}^2}\right) - \frac{(I_t - I_0 - \bar{\alpha}_t I_{res})^2}{\bar{\beta}_t^2}\right) \\ &= \exp\left(-\frac{1}{2}\left(\frac{\bar{\beta}_t^2}{\beta_t^2 \bar{\beta}_{t-1}^2} I_{t-1}^2 - 2\left(\frac{I_t - \alpha_t I_{res}}{\beta_t^2} + \frac{\bar{\alpha}_{t-1} I_{res} + I_0}{\bar{\beta}_{t-1}^2}\right) I_{t-1} + C(I_t, I_0, I_{res})\right)\right), \end{aligned} \quad (23)$$

where the $C(I_t, I_0, I_{res})$ term is not related to I_{t-1} . From Eq. 23, $\mu_t(x_t, I_0, I_{res})$ and $\Sigma_t(x_t, I_0, I_{res})$ are represented as follows,

$$\mu_t(x_t, I_0, I_{res}) = \left(\frac{I_t - \alpha_t I_{res}}{\beta_t^2} + \frac{\bar{\alpha}_{t-1} I_{res} + I_0}{\bar{\beta}_{t-1}^2}\right) / \frac{\bar{\beta}_t^2}{\beta_t^2 \bar{\beta}_{t-1}^2} \quad (24)$$

$$= \frac{\bar{\beta}_{t-1}^2}{\bar{\beta}_t^2} I_t + \frac{\beta_t^2 \bar{\alpha}_{t-1} - \bar{\beta}_{t-1}^2 \alpha_t}{\bar{\beta}_t^2} I_{res} + \frac{\beta_t^2}{\bar{\beta}_t^2} I_0 \quad (25)$$

$$= I_t - \alpha_t I_{res} - \frac{\beta_t^2}{\bar{\beta}_t} \epsilon, \quad (26)$$

$$\Sigma_t(x_t, I_0, I_{res}) = \frac{\beta_t^2 \bar{\beta}_{t-1}^2}{\bar{\beta}_t^2}. \quad (27)$$

Eq. 7 is used for the derivation from Eq. 25 to Eq. 26. Then, we define the generation process to start from $p_\theta(I_T) \sim \mathcal{N}(I_T; \mathbf{0}, \mathbf{I})$,

$$p_\theta(I_{t-1}|I_t) = q(I_{t-1}|I_t, I_0^\theta, I_{res}^\theta), \quad (28)$$

where $I_0^\theta = I_t - \bar{\alpha}_t I_{res}^\theta - \bar{\beta}_t \epsilon_\theta$ by Eq. 7. Here we only consider L_{t-1} in [17],

$$L_{t-1} = D_{KL}(q(I_{t-1}|I_t, I_0, I_{res}) || p_\theta(I_{t-1}|I_t)) \quad (29)$$

$$= \mathbb{E} \left[\left\| I_t - \alpha_t I_{res} - \frac{\beta_t^2}{\bar{\beta}_t} \epsilon - (I_t - \alpha_t I_{res}^\theta - \frac{\beta_t^2}{\bar{\beta}_t} \epsilon_\theta) \right\|^2 \right], \quad (30)$$

where D_{KL} denotes KL divergence. The noise ϵ can be represented as a transformation of I_{res} using Eq. 16, and then Eq. 30 simplifies to:

$$L_{res}(\theta) := \mathbb{E} \left[\lambda_{res} \left\| I_{res} - I_{res}^\theta(I_t, t, I_{in}) \right\|^2 \right] \quad (31)$$

In addition, the residuals I_{res} can be represented as a transformation of ϵ using Eq. 16, and then Eq. 30 simplifies to:

$$L_\epsilon(\theta) := \mathbb{E} \left[\lambda_\epsilon \left\| \epsilon - \epsilon_\theta(I_t, t, I_{in}) \right\|^2 \right], \quad (32)$$

We reset the weights, i.e., $\lambda_{res}, \lambda_\epsilon \in \{0, 1\}$, similar to DDPM [17] discarding weights (or regarding them as reweighting) in simplified training objective.

⁸Each step in Eq. 7 adds a new random Gaussian noise in the random forward diffusion. Thus for simplicity, we assume $q(I_t|I_{t-1}, I_0, I_{res}) = q(I_t|I_{t-1}, I_{res})$, it follows that I_0 is not important for I_t when I_{t-1} presents as a condition.

A.2. Deterministic Implicit Sampling

If $q_\sigma(I_{t-1}|I_t, I_0, I_{res})$ is defined in Eq. 11, we have:

$$q(I_t|I_0, I_{res}) = \mathcal{N}(I_t; I_0 + \bar{\alpha}_t I_{res}, \bar{\beta}_t^2 \mathbf{I}). \quad (33)$$

Proof. Similar to the evolution from DDPM [17] to DDIM [51], we can prove the statement with an induction argument for t from T to 1. Assuming that Eq. 33 holds at T , we just need to verify $q(I_{t-1}|I_0, I_{res}) = \mathcal{N}(I_{t-1}; I_0 + \bar{\alpha}_{t-1} I_{res}, \bar{\beta}_{t-1}^2 \mathbf{I})$ at $t-1$ from $q(I_t|I_0, I_{res})$ at t using Eq. 33. Given:

$$q(I_t|I_0, I_{res}) = \mathcal{N}(I_t; I_0 + \bar{\alpha}_t I_{res}, \bar{\beta}_t^2 \mathbf{I}), \quad (34)$$

$$q_\sigma(I_{t-1}|I_t, I_0, I_{res}) = \mathcal{N}(I_{t-1}; I_0 + \bar{\alpha}_{t-1} I_{res} + \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2} \frac{I_t - (I_0 + \bar{\alpha}_t I_{res})}{\bar{\beta}_t}, \sigma_t^2 \mathbf{I}), \quad (35)$$

$$q(I_{t-1}|I_0, I_{res}) := \mathcal{N}(\tilde{\mu}_{t-1}, \tilde{\Sigma}_{t-1}) \quad (36)$$

Similar to obtaining $p(y)$ from $p(x)$ and $p(y|x)$ using Eq.2.113-Eq.2.115 in [5], the values of $\tilde{\mu}_{t-1}$ and $\tilde{\Sigma}_{t-1}$ are derived as following:

$$\tilde{\mu}_{t-1} = I_0 + \bar{\alpha}_{t-1} I_{res} + \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2} \frac{(I_0 + \bar{\alpha}_t I_{res}) - (I_0 + \bar{\alpha}_t I_{res})}{\bar{\beta}_t} = I_0 + \bar{\alpha}_{t-1} I_{res}, \quad (37)$$

$$\tilde{\Sigma}_{t-1} = \sigma_t^2 \mathbf{I} + \left(\frac{\sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}}{\bar{\beta}_t} \right)^2 \bar{\beta}_t^2 \mathbf{I} = \beta_{t-1}^2 \mathbf{I}. \quad (38)$$

Therefore, $q(I_{t-1}|I_0, I_{res}) = \mathcal{N}(I_{t-1}; I_0 + \bar{\alpha}_{t-1} I_{res}, \bar{\beta}_{t-1}^2 \mathbf{I})$. In fact, the case ($t = T$) already holds, thus Eq. 33 holds for all t .

Simplifying Eq. 11. Eq. 11 can also be written as:

$$I_{t-1} = I_0 + \bar{\alpha}_{t-1} I_{res} + \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2} \frac{I_t - (I_0 + \bar{\alpha}_t I_{res})}{\bar{\beta}_t} + \sigma_t \epsilon_t, \quad (39)$$

$$= \frac{\sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}}{\bar{\beta}_t} I_t + \left(1 - \frac{\sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}}{\bar{\beta}_t}\right) I_0 + \left(\bar{\alpha}_{t-1} - \frac{\sqrt{\bar{\alpha}_t \bar{\beta}_{t-1}^2 - \sigma_t^2}}{\bar{\beta}_t}\right) I_{res} + \sigma_t \epsilon_t \quad (40)$$

$$= I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{res} - (\bar{\beta}_t - \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}) \epsilon + \sigma_t \epsilon_t, \quad (41)$$

where $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Eq. 41 is consistent with Eq. 12, and Eq. 7 is used for the derivation from Eq. 40 to Eq. 41.

A.3. Coefficient Transformation

For image generation, $I_{in} = 0$, thus Eq. 16 can also be written as:

$$I_t = (\bar{\alpha}_t - 1) I_{res} + \bar{\beta}_t \epsilon \quad (42)$$

$$I_{res} = \frac{I_t - \bar{\beta}_t \epsilon}{\bar{\alpha}_t - 1}. \quad (43)$$

If the residuals I_{res}^θ are represented as a transformation of ϵ_θ using Eq. 43, Eq. 12 is simplified to

$$I_{t-1} = I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{res}^\theta - (\bar{\beta}_t - \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}) \epsilon_\theta + \sigma_t \epsilon_t \quad (44)$$

$$= I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) \frac{I_t - \bar{\beta}_t \epsilon_\theta}{\bar{\alpha}_t - 1} - (\bar{\beta}_t - \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2}) \epsilon_\theta + \sigma_t \epsilon_t \quad (45)$$

$$= \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} I_t - \left(\frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \bar{\beta}_t - \sqrt{\bar{\beta}_{t-1}^2 - \sigma_t^2} \right) \epsilon_\theta + \sigma_t \epsilon_t \quad (46)$$

$$= \frac{\sqrt{\bar{\alpha}_{DDIM}^{t-1}}}{\sqrt{\bar{\alpha}_{DDIM}^t}} I_t - \left(\frac{\sqrt{\bar{\alpha}_{DDIM}^{t-1}} \sqrt{1 - \bar{\alpha}_{DDIM}^t}}{\sqrt{\bar{\alpha}_{DDIM}^t}} - \sqrt{1 - \bar{\alpha}_{DDIM}^{t-1} - \sigma_t^2} \right) \epsilon_\theta + \sigma_t \epsilon_t \quad (47)$$

$$= \sqrt{\bar{\alpha}_{DDIM}^{t-1}} \left(\frac{I_t - \sqrt{1 - \bar{\alpha}_{DDIM}^t} \epsilon_\theta}{\sqrt{\bar{\alpha}_{DDIM}^t}} \right) + \sqrt{1 - \bar{\alpha}_{DDIM}^{t-1} - \sigma_t^2} \epsilon_\theta + \sigma_t \epsilon_t. \quad (48)$$

Algorithm 1: Coefficient initialization, transformation, and adjustment.

Input : The initial conditions $\bar{\alpha}_T = 1$, $\bar{\beta}_T^2 > 0$, $T = 1000$, and $t \in \{1, 2, \dots, T\}$. The hyperparameter $\eta = 1$ for the random generation process and $\eta = 0$ for deterministic implicit sampling. Variance modes have “DDIM” and “DDIM→RDDM”. The coefficient adjustment mode Adjust=“Alpha”, “Beta”, or “Alpha+Beta”.

Output: The adjusted coefficients $\bar{\alpha}_t^*$, $\bar{\beta}_t^*$ and σ_t^* .

```
// (1) Coefficient Initialization
1  $\beta_{DDIM}^t \leftarrow \text{Linspace}(0.0001, 0.02, T)$  ▷ linear schedule [51]
2  $\alpha_{DDIM}^t \leftarrow 1 - \beta_{DDIM}^t$ 
3  $\bar{\alpha}_{DDIM}^t \leftarrow \text{Cumprod}(\alpha_{DDIM}^t)$  ▷ cumulative multiplication
4  $\sigma_t^2(DDIM) \leftarrow \eta \frac{(1 - \alpha_{DDIM}^{t-1})}{1 - \alpha_{DDIM}^t} (1 - \frac{\bar{\alpha}_{DDIM}^t}{\bar{\alpha}_{DDIM}^{t-1}})$ 
// (2) Coefficient Transformation
5  $\bar{\alpha}_t \leftarrow 1 - \sqrt{\bar{\alpha}_{DDIM}^t}$  ▷ Eq. 17
6  $\bar{\beta}_t \leftarrow \sqrt{1 - \bar{\alpha}_{DDIM}^t}$  ▷ Eq. 17
7  $\sigma_t^2(RDDM) \leftarrow \eta \frac{(\bar{\beta}_t^2 - \bar{\beta}_{t-1}^2) \bar{\beta}_{t-1}^2}{\bar{\beta}_t^2}$ 
// (3) Consistent Sampling Process
8 if Variance==“DDIM” then
9 |  $\sigma_t^* \leftarrow \sqrt{\sigma_t^2(DDIM)}$  ▷ consistent sampling process with DDIM [51] and DDPM [17]
10 else if Variance==“DDIM→RDDM” then
11 |  $\sigma_t^* \leftarrow \sqrt{\sigma_t^2(RDDM)}$  ▷ sum-constrained variance schedule
12 end
// (4) Coefficient Adjustment
13  $\alpha_t \leftarrow \text{Linspace}(1 - t/T, 1)$  ▷ linearly decreasing by Eq. 18
14  $\beta_t^2 \leftarrow \text{Linspace}(t/T, 1) \cdot \bar{\beta}_T^2$  ▷ control the noise intensity in  $I_T$  by  $\bar{\beta}_T^2$ 
15 if Adjust==“Alpha” then
16 |  $\bar{\alpha}_t^* \leftarrow \text{Cumsum}(\alpha_t)$ ,  $\bar{\beta}_t^* \leftarrow \bar{\beta}_t$  ▷ cumulative sum
17 else if Adjust==“Beta” then
18 |  $\bar{\alpha}_t^* \leftarrow \bar{\alpha}_t$ ,  $\bar{\beta}_t^* \leftarrow \sqrt{\text{Cumsum}(\beta_t^2)}$  ▷ cumulative sum
19 else if Adjust==“Alpha+Beta” then
20 |  $\bar{\alpha}_t^* \leftarrow \text{Cumsum}(\alpha_t)$ ,  $\bar{\beta}_t^* \leftarrow \sqrt{\text{Cumsum}(\beta_t^2)}$  ▷ coefficient reinitialization
21 else
22 |  $\bar{\alpha}_t^* \leftarrow \bar{\alpha}_t$ ,  $\bar{\beta}_t^* \leftarrow \bar{\beta}_t$ 
23 end
24 return  $\bar{\alpha}_t^*$ ,  $\bar{\beta}_t^*$ ,  $\sigma_t^*$  ▷ sampling with adjusted coefficients by Eq. 12
```

Eq. 48 is consistent with Eq.12 in DDIM [51] by replacing σ_t^2 with $\sigma_t^2(DDIM)$, and Eq. 17 is used for the derivation from Eq. 46 to Eq. 47. Thus, our sampling process is consistent with that of DDPM [51] and DDIM [17] by transforming coefficient/variance schedules.

We present the pipeline of coefficient transformation in Algorithm 1. Fig. 4 (b)(c) show the result of coefficient transformation. In Eq. 17, in addition to the coefficient $\bar{\alpha}_t, \bar{\beta}_t^2$ being replaced by $\bar{\alpha}_{DDIM}^t$, the variance σ_t^2 is also replaced with $\sigma_t^2(DDIM)$ to be consistent with DDIM [51] ($\eta = 0$) and DDPM [17] ($\eta = 1$). In fact, for DDIM [51] ($\eta = 0$), the variance is equal to 0 and does not need to be converted. Therefore, we analyze the difference between the variance of our RDDM and the variance of DDPM [17] in Appendix A.4.

A.4. Perturbed Generation Process with Sum-constrained Variance

From Eq. 12, the variance of our RDDM ($\eta = 1$) is

$$\sigma_t^2(RDDM) = \eta \frac{\beta_t^2 \bar{\beta}_{t-1}^2}{\bar{\beta}_t^2}. \quad (49)$$




Figure 8. Perturbed generation process with sum-constrained variance on the CelebA 64×64 dataset [36]. When the coefficient schedules $(\bar{\alpha}_t, \bar{\beta}_t)$ are consistent, RDDM $(\eta = 0)$ is the same as DDIM $(\eta = 0)$, but RDDM $(\eta = 1)$ is different from DDPM $(\eta = 1)$ due to different variance schedules (σ_t^2) . At different sampling steps, our RDDM has the same total noise, while DDPM has a different total noise. Notably, all the results in Fig. 8 can be generated from the same pre-trained model via variance transformation in Appendix A.3. In other words, the RDDM provides a sum-constrained variance strategy, which can be used directly in the pre-trained DDPM without re-training the model.

We replace $\bar{\beta}_t^2$ by $\bar{\alpha}_{DDIM}^t$ using Eq. 17 and replace β_t^2 by $\bar{\beta}_t^2 - \bar{\beta}_{t-1}^2$,

$$\sigma_t^2(RDDM) = \eta \bar{\alpha}_{DDIM}^{t-1} \frac{(1 - \bar{\alpha}_{DDIM}^{t-1})}{1 - \bar{\alpha}_{DDIM}^t} (1 - \frac{\bar{\alpha}_{DDIM}^t}{\bar{\alpha}_{DDIM}^{t-1}}), \quad (50)$$

while the variance of DDPM [17] $(\eta = 1)$ is

$$\sigma_t^2(DDIM) = \eta \frac{(1 - \bar{\alpha}_{DDIM}^{t-1})}{1 - \bar{\alpha}_{DDIM}^t} (1 - \frac{\bar{\alpha}_{DDIM}^t}{\bar{\alpha}_{DDIM}^{t-1}}). \quad (51)$$

Our variance is much smaller than the variance of DDPM [17] because $0 < \bar{\alpha}_{DDIM}^{t-1} < \alpha_{DDIM}^1 < 1$ (e.g., $\alpha_{DDIM}^1 = 0.02$ in linear schedule [51]). Compared to $\sigma_t^2(DDIM) \approx 1$ [53], the variance of our RDDM is sum-constrained,

$$\sum_{i=1}^T \sigma_i^2(RDDM) = \sum_{i=1}^T \eta \beta_i^2 \frac{\bar{\beta}_{i-1}^2}{\bar{\beta}_i^2} \leq \sum_{i=1}^T \beta_i^2 \leq 1, \quad (52)$$

where $\sum_{i=1}^T \beta_t^2 = \bar{\beta}_T^2 = 1$ for image generation. This is also consistent with the directional residual diffusion process with perturbation defined in Eq. 7. A qualitative comparison of our RDDM ($\eta = 1$) with DDIM [51] ($\eta = 0$) and DDPM [17] ($\eta = 1$) is shown in Fig. 8. Notably, for $\eta = 0$, our RDDM is consistent with DDIM [51] (in Fig. 5(a)(b)).

A.5. Comparison With Other Methods

The main difference is that to adapt the denoising diffusion, score, flow, or Schrödinger’s bridge to image restoration, they choose the noise (Shadow Diffusion [14], SR3 [49], and WeatherDiffusion [82]), the residual (DvSR [62] and Rectified Flow [35]), the target image (ColdDiffusion [2] and InDI [11]), or its linear transformation term (I2SB [29]), which is similar to a special case of our RDDM when it only predicts noise (SM-N) or residuals (SM-Res), while we introduce residual estimation but also embrace noise both for generation and restoration (SM-Res-N). We highlight that residuals and noise are equally important, e.g., the residual prioritizes certainty while the noise emphasizes diversity.

Differences from DDPM [17]. 1) DDPM is not interpretable for image restoration, while our RDDM is a unified, interpretable diffusion model for both image generation and restoration. 2) Differences in the definition of the forward process lead to different variance strategies. Our RDDM has sum-constrained variance (much smaller than the variance of DDPM), while DDPM has preserving variance [53] (see Appendix A.4).

Differences from IDDPM [44]. In the original DDPM [17], for the transfer probabilities $p_\theta(I_{t-1}|I_t)$ in Eq. 5, the mean $\mu_\theta(I_t, t)$ is learnable, while the variance Σ_t is fixed. IDDPM [44] highlights the importance of **estimating both the mean and variance**, demonstrating that learning variances allow for fewer sampling steps with negligible differences in sample quality. However, IDDPM [44] still only involves denoising procedures, and crucially, IDDPM, like DDPM, is thus not interpretable for image restoration. In addition, IDDPM [44] proposes three alternative ways to parameterize $\mu_\theta(I_t, t)$, i.e., predict mean $\mu_\theta(I_t, t)$ directly with a neural network, predict noise ϵ , or predict clean image I_0 . IDDPM [44] does not predict the clean image I_0 and noise ϵ at the same time, while both the residuals and the noise are predicted for our SM-Res-N.

The **essential difference** is that, our RDDM contains a mixture of three terms (i.e., input images I_{in} , target images I_0 , and noise ϵ) beyond DDPM/IDDPM (a mixture of two terms, i.e. I_0 and ϵ). **We emphasize that residuals and noise are equally important:** the residual prioritizes certainty, while the noise emphasizes diversity. Furthermore, our RDDM preserves the original DDPM generation framework by coefficient/variance transformation (Eq. 17), enabling seamless transfer of improvement techniques from DDPM, such as variance optimization from IDDPM.

Differences from ColdDiffusion [2]. 1) ColdDiffusion aims to remove the random noise entirely from the diffusion model, and replace it with other transforms (e.g., blur, masking), while our RDDM still embraces noise diffusion. Notably, we argue that noise is necessary for generative tasks that emphasize diversity (see Table 1). In fact, since ColdDiffusion discards random noise, extra noise injection is required to improve generation diversity. 2) To simulate the degradation process for different restoration tasks, ColdDiffusion attempts to use a Gaussian blur operation for deblurring, a snowification transform for snow removal., etc. These explorations may lose generality and differ fundamentally from our residual learning. RDDM represents directional diffusion from target images to input images using residuals, without designing specific degradation operators for each task. Additionally, RDDM provides solid theoretical derivation, while ColdDiffusion lacks theoretical justification.

Differences from DvSR [62]. [62] indeed use residual. But they 1) predict the initial clean image from a blurring image via a traditional (non-diffusion) network, calculate the residuals between the ground truth of the clean image and the predicted clean image 2) use denoising-based diffusion models predict noise like DDPM [17] and use a linear transformation of the noise to represent the residuals. They treat the residual predictions as an image generation task, aiming to produce diverse and plausible outputs based on the initial predicted clean image. Beyond simply building a diffusion model on top of residuals, we redefine a new forward process that allows simultaneous diffusion of residuals and noise, wherein the target image progressively diffuses into a purely noise or a noise-carrying input image.

Differences from InDI [11] and I2SB [29]. We can conclude that the forward diffusion of InDI, I2SB, and our RDDM is consistent in the form of a mixture of three terms (i.e., input images I_{in} , target images I_0 , and noise ϵ) beyond the denoising-based diffusion (a mixture of two terms, i.e. I_0 and ϵ). Substituting $I_{res} = I_{in} - I_0$ into Eq. 16 results in $I_t = \bar{\alpha}_t I_{in} + (1 - \bar{\alpha}_t) I_0 + \bar{\beta}_t \epsilon$. This resulted I_t has the same format as Eq.8 in InD ($x_t = ty + (1 - t)x + \sqrt{t}\epsilon_t \eta_t$), and is the same format as Eq.11 in I2SB. Similar to Eq. 17 (from our RDDM to DDPM/DDIM), transforming coefficients leads to complete consistency. However, our RDDM can further extend DDPM/DDIM, InD, and I2SB to independent double diffusion processes, and holds the potential to pave the way for the multi-dimensional diffusion process. From the initial stages of constructing a new forward process, our RDDM uses independent coefficient schedules to control the diffusion of residuals and noise. This provides a more general, flexible, and scalable framework, and inspires our partially path-independent generation process, demonstrated in Fig. 6 and Fig. 16(b-f) with stable generation across various diffusion rates

Tasks	Image Restoration				Image	Image	Image
	Shadow Removal	Low-light	Deblurring	Deraining	Generation	Inpainting	Translation
Datasets	ISTD	LOL	GoPro	RainDrop	CelebA	CelebA-HQ	CelebA-HQ AFHQ
Batch size	1	1	1	1	128	64	64
Image size	256	256	256	256	64	64	64
$\bar{\beta}_T^2$	0.01	1	0.01	1	1	1	1
I_{in}	I_{in}	I_{in}	I_{in}	I_{in}	0	0	0
Sampling steps	5	2	2	5	10	10	10
Loss type	ℓ_1	ℓ_1	ℓ_1	ℓ_1	ℓ_2	ℓ_2	ℓ_2
Loss	$L_{res} + L_\epsilon$	L_{res}	$L_{res} + L_\epsilon$	$L_{res} + L_\epsilon$	L_ϵ	L_{res}, L_ϵ	L_{res}, L_ϵ
Sampling Method	SM-Res-N-2Net	SM-Res	SM-Res-N-2Net	SM-Res-N-2Net	SM-N	SM-Res-N-2Net	SM-Res-N-2Net
Optimizer	Adam	Adam	Adam	Adam	RAdam	RAdam	RAdam
Learning rate	8e-5	8e-5	8e-5	8e-5	2e-4	2e-4	2e-4
Training iterations	80k	80k	400k	120k	100k	100k	100k
Schedules	$\alpha_t : P(1-x, 1)$ $\beta_t^2 : P(x, 1)$	$\alpha_t : P(1-x, 1)$ $\beta_t^2 : P(x, 1)$	$\alpha_t : P(1-x, 1)$ $\beta_t^2 : P(x, 1)$	$\alpha_t : P(1-x, 1)$ $\beta_t^2 : P(x, 1)$	$\alpha_{DDIM}^t \rightarrow \alpha_t : P(1-x, 1)$ α_t, β_t^2	$\alpha_t : P(1-x, 1)$ $\beta_t^2 : P(x, 1)$	$\alpha_{DDIM}^t \rightarrow \alpha_t, \beta_t^2$

Table 4. Experimental settings for training our RDDM. “SM-Res-N-2Net” is described in Appendix B.2. Two optimizers can be implemented in L_{res}, L_ϵ . We use “SM-Res-N-2Net” and $L_{res} + L_\epsilon$ on the SID-RGB dataset [65].

and path variations.

B. Experiments

B.1. Training Details

We use a UNet architecture⁹ for both residual prediction and noise prediction in our RDDM. The UNet settings remain consistent across all tasks, including the channel size (64) and channel multiplier (1,2,4,8). Detailed experimental settings can be found in Table 4. Training and testing for all experiments in Table 4 can be conducted on a single Nvidia GTX 3090.

Image Generation. For comparison with DDIM [51], we convert the α_{DDIM}^t schedule of DDIM [51] into the α_t and β_t^2 schedules of our RDDM using Eq. 17 in Section 5.2 and Section 6. In fact, a better coefficient schedule can be used for training in our RDDM, e.g., α_t (linearly decreasing) and β_t^2 (linearly increasing) in Table 2. The quantitative results were evaluated by Frechet Inception Distance (FID) and Inception Score (IS).

Image Restoration. We extensively evaluate our method on several image restoration tasks, including shadow removal, low-light enhancement, image deraining, and image deblurring on 5 different datasets. For fair comparisons, the results of other SOTA methods are provided from the original papers whenever possible. For all image restoration tasks, the images are resized to 256, and the networks are trained with a batch size of 1. We use shadow masks and shadow images as conditions for shadow removal (similar to [28, 79]), while other image restoration tasks use the degraded image as condition inputs. For low-light enhancement, we use histogram equalization for pre-processing. To cope with the varying tasks and dataset sizes, we only modified the number of training iterations, $\bar{\beta}_T^2$ and sampling steps (5 steps for shadow removal and deraining, 2 steps for low-light and deblurring) as shown in Table 4. α_t is initialized using a linearly decreasing schedule (i.e., $P(1-x, 1)$ in Eq. 18), while β_t^2 is initialized using a linearly increasing schedule (i.e., $P(x, 1)$). The quantitative results were evaluated by Mean Absolute Error (MAE), Peak Signal to Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [73].

Notably, our RDDM uses an identical UNet architecture and is trained with a batch size of 1 for all these tasks. In contrast, SOAT methods often involve elaborate network architectures, such as multi-stage [13, 70, 80], multi-branch [10], and GAN [27, 45, 57], or sophisticated loss functions like the chromaticity [20], texture similarity [74], and edge loss [70].

Image Inpainting and Image Translation. To increase the diversity of the generated images, conditional input images were not fed into the deresidual and denoising network (see Fig. 19).

⁹Our RDDM is implemented by modifying <https://github.com/lucidrains/denoising-diffusion-pytorch> repository.

Algorithm 2: Training Pipeline Using AOSA.

Input : A degraded input image I_{in} and its corresponding ground truth image I_0 . Gaussian noise ϵ . Time condition t . Coefficient schedules $\bar{\alpha}$ and $\bar{\beta}_t$. The initial learnable parameters $\lambda_{res}^\theta = 0.5$. Network G with parameters θ . The initial learning rate l . n is the training iterations number. m is the iterations number of AOSA. The threshold of shifting training strategies, $\delta = 0.01$.

Output: Trained well parameters, $\theta, \lambda_{res}^\theta$.

```
1  $\theta \leftarrow \mathbb{I}\mathfrak{H}$  (G) ▷ initialize network parameters
2 for  $i \leftarrow 1$  to  $n + m$  do
3    $t \sim \bar{\mathfrak{H}}$  ( $\{1, 2, \dots, T\}$ ),  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $I_{res} \leftarrow I_{in} - I_0$ 
4    $I_t \leftarrow I_0 + \bar{\alpha}_t I_{res} + \bar{\beta}_t \epsilon$  ▷ synthesize  $I_t$  by Eq. 7
5    $I_{out} \leftarrow G(I_t, t, I_{in})$ 
6    $I_{res}^\theta \leftarrow \lambda_{res}^\theta \times I_{out} + (1 - \lambda_{res}^\theta) \times f_{\epsilon \rightarrow res}(I_{out})$  ▷  $f_{\epsilon \rightarrow res}(\cdot)$ : from  $\epsilon$  to  $I_{res}$  using Eq. 16
7    $\epsilon_\theta \leftarrow \lambda_{res}^\theta \times f_{res \rightarrow \epsilon}(I_{out}) + (1 - \lambda_{res}^\theta) \times I_{out}$  ▷  $f_{res \rightarrow \epsilon}(\cdot)$ : from  $I_{res}$  to  $\epsilon$  using Eq. 16
8    $L_{auto} \leftarrow \mathbb{L}\mathfrak{S}$  ( $I_{res}^\theta, I_{res}, \epsilon_\theta, \epsilon$ ) ▷ based on Eq. 53
9    $\theta, \lambda_{res}^\theta \stackrel{\pm}{\leftarrow} -\nabla_{\theta, \lambda_{res}^\theta} (\mathcal{L}_{auto}, l)$  ▷ updating gradient
10  if  $\mathfrak{B}$  ( $\lambda_{res}^\theta - 0.5$ )  $< \delta$  then
11    |  $\mathfrak{B}$  ▷ adversarial-like training
12  else
13    |  $\lambda_{res}^\theta \leftarrow \mathbb{D}\mathfrak{B}$  ( $\lambda_{res}^\theta$ ) ▷ halt the gradient updates
14    |  $\theta \leftarrow \mathbb{I}\mathfrak{H}$  (G) ▷ reinitialize network parameters if  $\lambda_{res}^\theta > 0.5$  then
15    | |  $\lambda_{res}^\theta \leftarrow 1$  ▷ SM-Res
16    | else
17    | |  $\lambda_{res}^\theta \leftarrow 0$  ▷ SM-N
18    | end
19  end
20 end
```

Sampling Method	Network	MAE(↓)	SSIM(↑)	PSNR(↑)
SM-Res	Residual network	4.76	0.959	30.72
SM-Res-N-2Net	Residual network+noise network	<u>4.67</u>	<u>0.962</u>	<u>30.91</u>
SM-Res-N-1Net	One network, only shared encoder	4.72	0.959	30.73
SM-Res-N-1Net	One network	4.57	0.963	31.10

Table 5. Ablation studies of sampling methods and network structures on the ISTD dataset [57]. “SM-Res-N-1Net+One network” denotes to output 6 channels using a network, where the 0-3-th channels are residual and the 3-6-th channels are noise.

B.2. SAMPLING Details

In Section 4.3, we described the empirical selection process for SM-Res or SM-N. Below, we give some additional methods in detail.

Automatic selection of SM-Res or SM-N by AOSA. We present the motivation, conceptualization, and implementation pipeline (Algorithm 2) of the Automatic Objective Selection Algorithm (AOSA) as follows:

Step 1. At the initial simultaneous training (similar to SM-Res-N), we do not know whether the network output (I_{out}) is residual or noise. Therefore, we set $\lambda_{res}^\theta = 0.5$ to denote the probability that the output is residual (I_{res}^θ), and $1 - \lambda_{res}^\theta$ is the probability that the output is noise (ϵ_θ).

Step 2. We then impose loss constraints on both residual and noise estimation weighted by the learned parameter (λ_{res}^θ), as follows:

$$L_{auto}(\theta) := \lambda_{res}^\theta E \left[\|I_{res} - I_{res}^\theta(I_t, t, I_{in})\|^2 \right] + (1 - \lambda_{res}^\theta) E \left[\|\epsilon - \epsilon_\theta(I_t, t, I_{in})\|^2 \right]. \quad (53)$$

The joint loss functions $L_{auto}(\theta)$ drive the network to gradually tend to output residuals or noise based on the input. For

Sampling Method	Shadow	Deraining	Deblurring	Generation	
	PSNR/SSIM	PSNR/SSIM	PSNR/SSIM	FID(↓)	IS(↑)
SM-Res	30.72/0.959	<u>31.96/0.9509</u>	<u>32.32/0.957</u>	31.47	1.73
SM-N	11.34/0.175	19.15/0.7179	9.49/0.087	23.25	2.05
SM-Res-N-2Net	<u>30.91/0.962</u>	32.51/0.9563	32.40/0.963	28.90	1.78
SM-Res-N-1Net	31.10/0.963	31.79/0.9504	31.69/0.951	<u>28.57</u>	<u>1.81</u>

Table 6. Generalization analysis of “SM-Res-N-1Net+One network”.

example, in the image restoration task with deterministic input, it should be simpler for the network to estimate a clear image than noise. In contrast, for the image generation task with random noise input, it is simpler for the network to estimate the noise than a clear image.

Step 3. To enable learning of λ_{res}^θ , we then include it in the network computation, allowing gradient transmission. Since λ_{res}^θ denotes the probability that the output is residual, the estimated residual I_{res}^θ can be represented as $\lambda_{res}^\theta \times I_{out} + (1 - \lambda_{res}^\theta) \times f_{\epsilon \rightarrow res}(I_{out})$. $f_{\epsilon \rightarrow res}(\cdot)$ represents the transformation from ϵ to I_{res} using Eq. 16. Similarly, ϵ_θ can be represented as $\lambda_{res}^\theta \times f_{res \rightarrow \epsilon}(I_{out}) + (1 - \lambda_{res}^\theta) \times I_{out}$. This is very similar to the cross-entropy loss function.

Step 4. As the training process is completed, our objective should be to estimate only noise (SM-N) or residuals (SM-Res). By utilizing the learned λ_{res}^θ , we can determine when to switch from an adversarial-like process (residuals vs. noise in Step 2) to a single prediction (residuals or noise). This transition can be controlled, for instance, by setting a condition such as $\text{abs}(\lambda_{res}^\theta - 0.5) \geq 0.01$. When the network’s tendency to estimate residuals surpasses 51% probability, we set λ_{res}^θ to 1 and halt the gradient updates for λ_{res}^θ .

The experimental results were consistent with the empirical analysis in Section 4.3 and verified the effectiveness of AOSA. For instance, the initial simultaneous training switches to residual learning (SM-Res) for shadow removal and low-light enhancement in approximately 300 iterations, and to denoising learning (SM-N) for image generation in approximately 1000 iterations. To summarize, AOSA achieves the same inference cost as the current denoising-based diffusion methods [17] with the plug-and-play training strategy.

SM-Res-N. Both the residuals and the noise are predicted, which can be implemented with two or one networks. **SM-Res-N-2Net.** If computational resources are sufficient, two separate networks can be trained for noise and residual predictions, and the optimal sampling method can be determined during testing. This setting easily obtains a well-suited network for the target task, and facilitates the exploration of the decoupled diffusion process and the partially path-independent generation process in Section 5. **SM-Res-N-1Net.** To avoid training two separate networks, another solution is to simply use a joint network (i.e., a shared encoder and decoder) to output 6 channels where the 0-3-th channels are residual and the 3-6-th channels are noise. This setting loses the decoupling property of RDDM, but can achieve dual prediction with a slight cost. Table 5 shows that the joint network (i.e., SM-Res-N-1Net+One network) achieves the best shadow removal results (MAE 4.57), even better than two independent networks (4.67). A network with the shared encoder (MAE 4.72) has a slight performance degradation compared to the independent two networks (4.67). We conduct further generalization experiments on Table 6, indicating

Method	MAE(↓)			SSIM(↑)			PSNR(↑)			LPIPS(↓)
	S	NS	ALL	S	NS	ALL	S	NS	ALL	
ST-CGAN [57]	10.33	6.93	7.47	0.981	0.958	0.929	33.74	29.51	27.44	-
DSC [19] ¶	9.48	6.14	6.67	0.967	-	-	33.45	-	-	-
DHAN [10]	8.14	6.04	6.37	0.983	-	-	34.50	-	-	-
CANet [8]	8.86	6.07	6.15	-	-	-	-	-	-	-
LG-ShadowNet [37]	10.23	5.38	6.18	0.979	0.967	0.936	31.53	29.47	26.62	-
FusionNet [13]	7.77	5.56	5.92	0.975	0.880	0.945	34.71	28.61	27.19	0.1204
UnfoldingNet [80]	7.87	4.72	5.22	0.987	<u>0.978</u>	0.960	36.95	31.54	29.85	-
BMNet [79]	7.60	4.59	5.02	<u>0.988</u>	0.976	0.959	35.61	32.80	30.28	0.0377
DMTN [31]	<u>7.00</u>	<u>4.28</u>	<u>4.72</u>	0.990	0.979	0.965	35.83	<u>33.01</u>	<u>30.42</u>	<u>0.0368</u>
Ours (RDDM)	6.67	4.27	4.67	<u>0.988</u>	0.979	<u>0.962</u>	<u>36.74</u>	33.18	30.91	0.0305

Table 7. Shadow removal results on the ISTD dataset [57]. We report the MAE, SSIM and PSNR in the shadow area (S), non-shadow area (NS), and whole image (ALL).




Figure 9. More visual comparison results for shadow removal on the ISTD dataset [57].




Figure 10. More visual comparison results for Low-light enhancement on the LOL dataset [61].

“SM-Res-N-1Net+One network” isn’t bad, but not always the best.

B.3. More Results

Shadow removal. We compare RDDM with DSC [19], FusionNet [13], BMNet [79] and DMTN [31] on the ISTD dataset [57]. The ISTD dataset [57] contains shadow images, shadow masks, and shadow-free image triplets (1,330 for training; 540 for testing). Table 3(b), Fig. 7(b), and Fig. 9 demonstrate the superiority of our method. In addition, we compare RDDM with more shadow removal methods (e.g., ST-CGAN [57], DHAN [10], CANet [8], LG-ShadowNet [37],




Figure 11. More visual comparison results for Low-light enhancement on the SID-RGB dataset [65].

Low-light (SID-RGB)	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
SID [7]	21.16	0.6398	0.4026
D&E [65]	22.13	0.7172	0.3794
MIR-Net[69, 71]	22.34	0.7031	0.3562
UTVNet [78]	22.69	0.7179	0.3417
SNR-Aware [66]	22.87	0.625	-
Our RDDM (2 step)	23.97	0.8392	0.2433
Our RDDM (5 step)	<u>23.80</u>	<u>0.8378</u>	0.2289

Table 8. Quantitative comparison results of Low-light enhancement on the SID-RGB dataset [65]. The results of MIR-Net are reported by [78].

UnfoldingNet [80]) in Table 7.

Low-light enhancement. We evaluate our RDDM on the LOL datasets [61] (500 images) and SID-RGB [65] dataset (5,094 images), and compare our method with the current SOTA methods [33, 66, 68, 71, 76–78]. To unify and simplify the data loading pipeline for training, we only evaluate the RGB low-light image dataset [61, 65], not the RAW datasets (e.g., FiveK [6]). Table 3(c), Fig. 7(c), and Fig. 10 show that our RDDM achieves the best SSIM and LPIPS [73] and can recover better visual quality on the LOL [61] dataset. Table 8 shows the low-light enhancement results on the SID-RGB [65]



Figure 12. More visual comparison results for deraining on the RainDrop [45] dataset.



Figure 13. More visual comparison results for deblurring on the GoPro [43] dataset.

dataset of different methods. Our RDDM outperforms the state-of-the-art SNR-Aware [66] by a 4.8% PSNR and a 34.2%

Deblurring (GoPro)	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
Deblurgan-v2 [27]	29.55	0.934	0.117
Suin <i>et al.</i> [55]	31.85	0.948	-
MPRNet [70]	32.66	0.959	0.089
DvSR [62]	31.66	0.948	0.059
Uformer-B [60]	32.97	0.967	0.0089
I2SB [29]	29.31	0.906	0.0961
InDI [11]	31.49	0.946	0.058
Our RDDM (2 step)	32.40	<u>0.963</u>	0.0415
Our RDDM (10 step)	31.67	0.950	<u>0.0379</u>

Table 9. Quantitative comparison results of deblurring on the GoPro dataset [43].

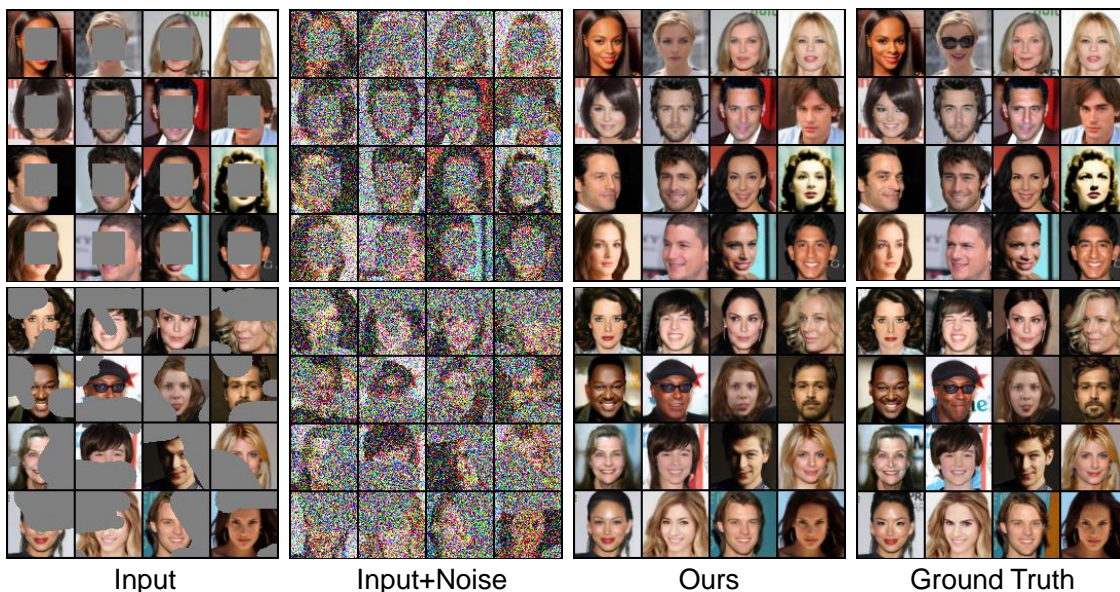


Figure 14. More visual results for image inpainting on the CelebA-HQ [23] dataset.

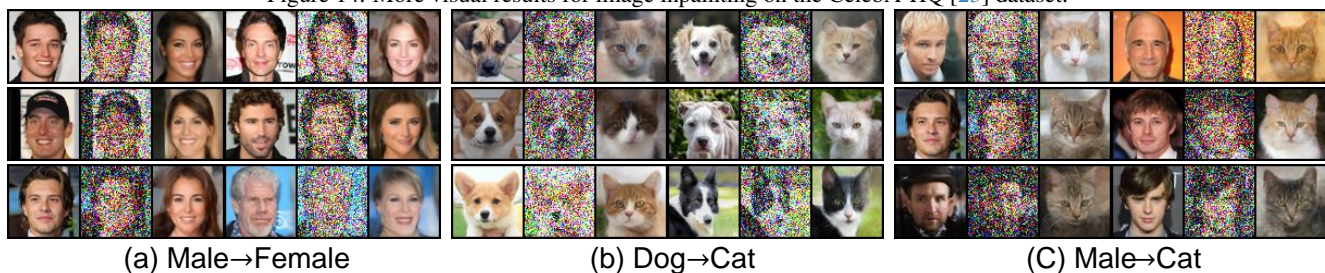


Figure 15. More visual results for image translation on the CelebA-HQ [23] and AFHQ [9] datasets.

SSIM improvement on the SID-RGB [65] dataset. Fig. 11 shows that our RDDM outperforms competitors in detail recovery (sharper text of 1st row), and color vibrancy (2nd & 3rd rows), avoiding issues like gray shading and detail blurring.

Image deraining. We make a fair comparison with the current SOTA diffusion-based image restoration method - Rain-Diff128 [82] (“128” denotes the 128×128 patch size for training) on the RainDrop dataset [45] (1119 images). Rain-Diff128 [82] feeds the degraded input image as a condition to the denoising network, which requires 50 sampling steps to generate a clear image from the noise, while our RDDM requires only 5 sampling steps to recover the degraded image from the noise-carrying input image and outperforms RainDiff128 [82], as shown in Table 3(d) and Fig. 12.

Image deblurring. We evaluate our method on the widely used deblurring dataset - GoPro [43] (3,214 images). Table 9




Figure 16. More visual results for the partially path-independent generation process. Two networks are used to estimate residuals and noise separately, i.e., $I_{res}^\theta(I_t, \tilde{\alpha}_t \cdot T)$ and $\epsilon_\theta(I_t, \tilde{\beta}_t \cdot T)$ ($\eta = 0$).

shows that our RDDM is the second best LPIPS. Fig. 13 shows that our method is competitive with the SOTA deblurring methods (e.g, Uformer-B [60]) in terms of visual quality.

Image Inpainting and Image Translation. We show more qualitative results of image inpainting (Fig. 14) and translation (Fig. 15).

B.4. Partially Path-independent Generation Process

Fig. 16(b-f) provides evidence supporting the partially path-independent generation process, demonstrating the robustness of the generative process within a certain range of diffusion rates (step size per step) and path variation, e.g., converting DDIM [51] to a uniform diffusion speed in Fig. 16(c). However, excessive disturbances can result in visual inconsistencies, as depicted in Fig. 16(h)(i). Furthermore, Fig. 16(c) and Fig. 16(g) illustrate that even when the paths are the same, the variation in diffusion speed significantly impacts the quality of the generated images. This highlights the importance of carefully considering and controlling the diffusion speed and generation path during the generation process.

We also investigated two reverse paths to gain insight into the implications of the proposed partial path independence. In the first case, the residuals are removed first, followed by the noise: $I(T) \xrightarrow{-I_{res}} I(0) + \tilde{\beta}_T \epsilon \xrightarrow{-\tilde{\beta}_T \epsilon} I(0)$, as shown in Fig. 17(b1)(b2). The second case involves removing the noise first and then the residuals: $I(T) \xrightarrow{-\tilde{\beta}_T \epsilon} I_{in} \xrightarrow{-I_{res}} I(0)$. In the first case, images are successfully generated (as shown in Fig. 17(b)) which exhibit a striking similarity to the default images in Fig. 17(a). However, the second case shown in Fig. 17(c) fails to go from I_{in} to $I(0)$ due to $I_{in} = 0$ in the generation task. Figure 17(d) shows the intermediate visualization results of removing the noise first.



Figure 17. Special paths of the partially path-independent generation process. Two networks are used to estimate residuals and noise separately, i.e., $I_{res}^\theta(I_t, \bar{\alpha}_t \cdot T)$, $\epsilon_\theta(I_t, \bar{\beta}_t \cdot T)$ ($\eta = 0$).

B.5. Ablation Studies

We have analyzed the sampling method in Table 1, the coefficient schedule in Table 2, and the network structure for SM-Res-N in Table 5.

Sampling Methods. We present the results for noise predictions only (SM-N) in Fig. 18. Fig. 18 (b) and (c) illustrate that estimating only the noise poses challenges as colors are distorted, and it becomes difficult to retain information from the input shadow image. We found that increasing sampling steps does not lead to improved results from Fig. 18 (b) to Fig. 18 (c), which may be an inherent limitation when estimating only the noise for image restoration. Actually, this is also reflected in DeS3 [21] (a shadow removal method based on a denoising diffusion model), where DeS3 [21] specifically designs the loss against color bias. Additionally, training with batch size 1 may contribute to poor results of only predicting noise. However, estimating only the residuals (SM-Res) with batch size 1 does not exhibit such problems for image restoration, as demonstrated in Fig. 18 (d)&(e) and Table 1, further demonstrating the merits of our RDDM. For image inpainting, SM-Res-N-2Net can generate more realistic face images compared to SM-N and SM-Res, as shown in Fig. 19(d-f). If computational resources are sufficient, to obtain better image quality for an unknown task, we suggest that two separate networks can be trained for noise and residual predictions, and the optimal sampling method can be determined during




Figure 18. Visualizing ablation studies of sampling methods.

testing. If computational resources are limited, the sampling method can be determined empirically (see Section 4.3).

RDDM (SM-Res-N-2Net)	metric	1 step	2 step	5 step	10 step	100 step
$\bar{\beta}_T^2 = 0.01$	MAE-ALL (\downarrow)	4.83	4.69	4.67	4.72	4.90
	PSNR-S (\uparrow)	36.83	36.98	36.74	36.59	36.41
	LPIPS (\downarrow)	0.0344	0.0308	0.0305	0.0314	0.0334
$\bar{\beta}_T^2 = 1$	MAE-ALL (\downarrow)	5.07	4.94	4.90	4.87	4.99
	PSNR-S (\uparrow)	36.93	37.20	37.07	37.01	36.62
	LPIPS (\downarrow)	0.0346	0.0314	0.0298	0.0300	0.0319

Table 10. Ablation studies of sampling steps for shadow removal on the ISTD dataset [57].

Sampling Steps. Table 10 shows that our RDDM performance improves as the number of sampling steps increases for shadow removal. Unlike the findings in InDI [11] where one step yields the best MSE reconstruction, RDDM achieved its best MAE at 5 steps for shadow removal, which may be due to differences in coefficient schedules and sampling methods.




Figure 19. Visualizing ablation studies of sampling methods for image inpainting on the CelebA-HQ [23] dataset. (g-i) The conditional input image (a) is not used as an input to the deresidual and denoising network in the generation process from (b) to (f). Compared to (g-i), the diversity of the generated images in (j-l) decreases.

Certainty and diversity. Indeed, feeding conditional input images (Fig. 19(a)) into the deresidual and denoising network enhances the certainty of the generated images, while diminishing diversity, as shown in Fig. 19(j-l). Generating a clear target image directly from a noisy-carrying degraded image (Fig. 19(b)) without any conditions increases diversity, but changes non-missing regions (Fig. 19(g-i)).

Noise Perturbation Intensity. Table 10 shows that noise might play a beneficial role in restoring image details and enhancing perceptual quality. For image generation, the diversity of the generated images decreases significantly as $\bar{\beta}_T^2$ decreases from 1 in Fig. 20(c) to 0.01 in Fig. 20(b). The experiment is related to the mean face [18, 38, 42, 63] and

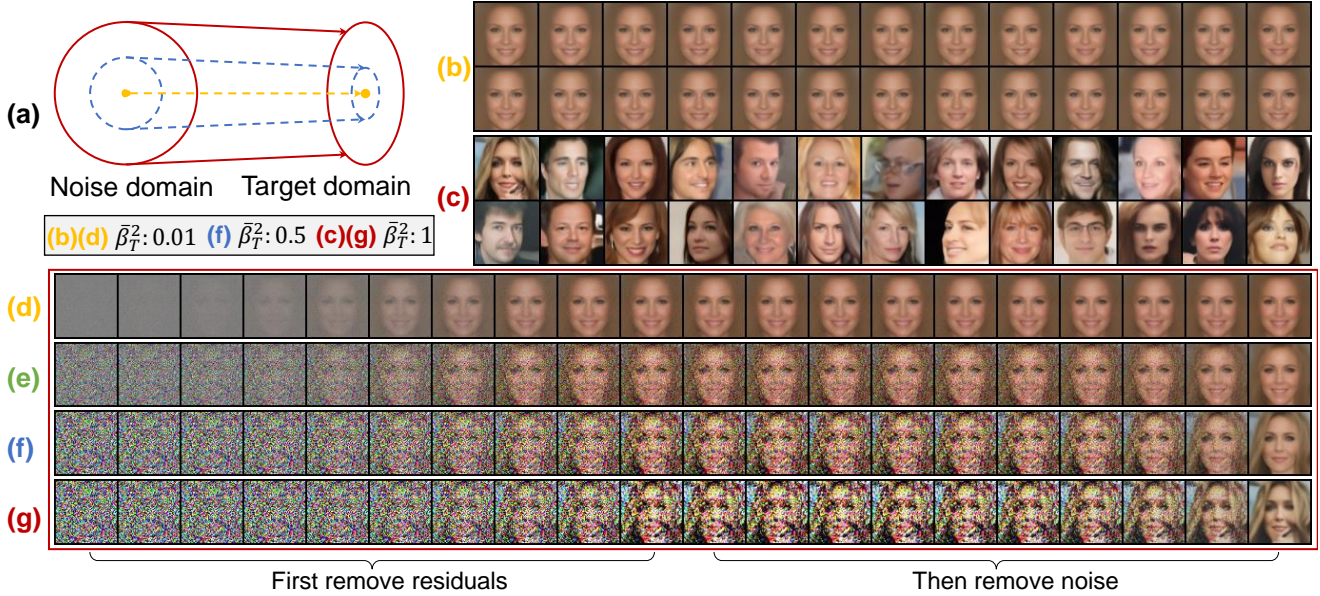


Figure 20. Visualizing ablation studies of noise perturbation intensity ($\bar{\beta}_T^2$). (a) We change the variance ($\bar{\beta}_T^2$) during testing, specifically by variance transformation via Algorithm 1. (b-c) When $\bar{\beta}_T^2$ decreases from 1 in (c) to 0.01 in (b), the diversity of the generated images decreases significantly. (d-g) We visualize each step in the generation process. $\bar{\beta}_T^2 = 0.01$ in (d), $\bar{\beta}_T^2 = 0.1$ in (e), $\bar{\beta}_T^2 = 0.5$ in (f), and $\bar{\beta}_T^2 = 1$ in (g). The sampling method is SM-Res-N-2Net with 10 sampling steps.

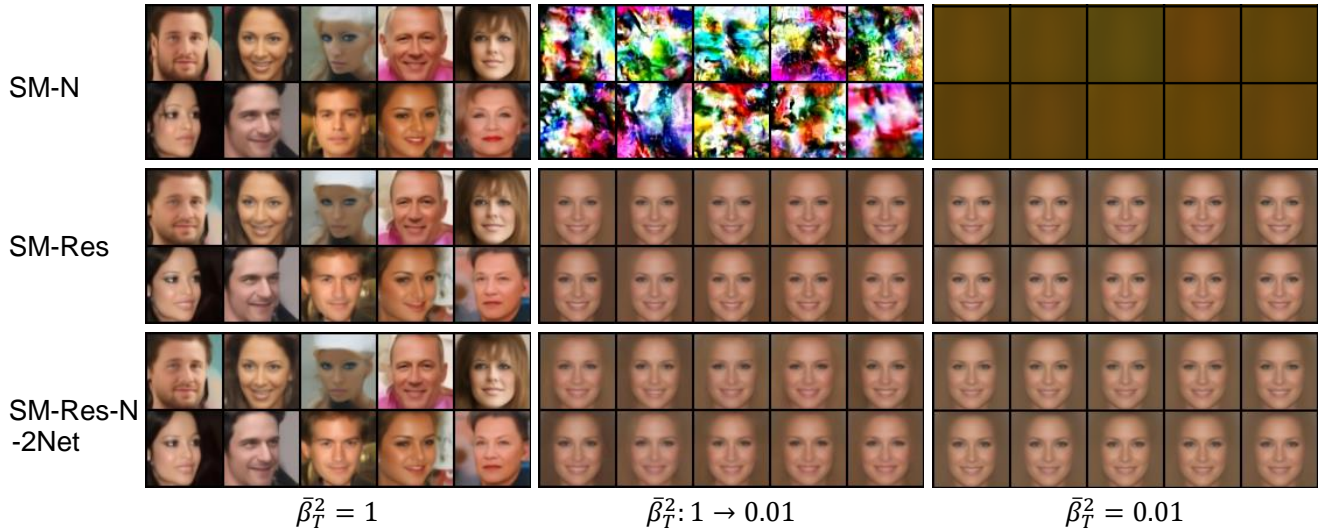


Figure 21. Visualizing ablation studies of sampling methods with different intensities of noise perturbation ($\bar{\beta}_T^2$). “ $\bar{\beta}_T^2: 1 \rightarrow 0.01$ ” denotes that the variance ($\bar{\beta}_T^2$) is changed during testing by variance transformation via Algorithm 1. The sampling steps are 10.

could provide useful insights to better understand the generative process. Fig. 21 shows that modifying $\bar{\beta}_T^2$ during testing ($\bar{\beta}_T^2: 1 \rightarrow 0.01$) causes SM-N to fail to generate meaningful faces. SM-Res-N-2Net including deresidual and denoising networks can generate meaningful face images like SM-Res, indicating that the denoising network can perform denoising when modifying $\bar{\beta}_T^2$, but cannot obtain robust residuals (I_{res}^θ) in the sampling process by Eq. 16. In summary, the deresidual network is relatively robust to noise variations compared to the denoising network.

Resource efficiency. Due to fewer sampling steps, our RDDM inference time and performance is comparable to Illflow [59], and LLFormer [58] (not diffusion-based). Compared to SR3 [49], our RDDM (only res in Table 11(b)) has 10x fewer training iterations, 10x fewer parameters, 10x faster inference time, and 10% improvement in PSNR and SSIM on the ISTD [57] dataset (shadow removal). For a fair comparison, priori shadow masks are used in SR3 [49] with a batch

(a) Low-light	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)	Params(M)	MACs(G) \times Steps	Inference Time(s)
LLformer	23.649	0.816	0.169	24.51	22.0 \times 1 = 22.0	$0.09 \times 1 = 0.09$
LLFlow	<u>25.19</u>	<u>0.93</u>	0.11	<u>17.42</u>	$286.33 \times 1 = 286.3$	$0.18 \times 1 = 0.18$
Ours(RDDM)	25.392	0.937	<u>0.116</u>	7.73	$32.9 \times 2 = 65.8$	0.03 \times 2 = 0.06

(b) Shadow Removal	MAE(\downarrow)	PSNR(\uparrow)	SSIM(\uparrow)	Params(M)	MACs(G) \times Steps	Inference Time(s)
Shadow Diffusion [14]	4.12	32.33	0.969	-	-	-
SR3 [49] (80k)	14.22	25.33	0.780	155.29	$155.3 \times 100 = 15530.0$	$0.02 \times 100 = 2.00$
SR3 [49] (500K)	13.38	26.03	0.820	155.29	$155.3 \times 100 = 15530.0$	$0.02 \times 100 = 2.00$
SR3 [49] (1000K)	11.61	27.49	0.871	155.29	$155.3 \times 100 = 15530.0$	$0.02 \times 100 = 2.00$
Ours (only res, 80k)	4.76	30.72	0.959	7.74	33.5 \times 5 = 167.7	0.03 \times 5 = 0.16
Ours (80k)	<u>4.67</u>	<u>30.91</u>	<u>0.962</u>	<u>15.49</u>	$67.1 \times 5 = 335.5$	$0.06 \times 5 = 0.32$

(c) Deraining	PSNR(\uparrow)	SSIM(\uparrow)	Params(M)	MACs(G) \times Steps	Inference Time(s)
RainDiff64[28]	32.29	0.9422	109.68	$252.4 \times 10 = 2524.2$	$0.03 \times 10 = 0.38$
RainDiff128[28]	<u>32.43</u>	0.9334	109.68	$248.4 \times 50 = 12420.0$	$0.038 \times 50 = 1.91$
Ours (only res)	31.96	<u>0.9509</u>	7.73	32.9 \times 5 = 164.7	0.032 \times 5 = 0.16
Ours	32.51	0.9563	<u>15.47</u>	$65.8 \times 5 = 329.3$	$0.07 \times 5 = 0.35$

Table 11. Resource efficiency and performance analysis by THOP. “MAC” means multiply-accumulate operation. (a) Low-light enhancement on the LoL dataset [61]. (b) Shadow removal on the ISTD dataset [57]. For a fair comparison, a priori shadow mask are used in SR3 with a batch size of 1. (c) Deraining on the RainDrop dataset [45].




Figure 22. Failure cases. (a1-a4) Shadow removal on the ISTD dataset [57]. (b1-b4) Low-light enhancement on the SID-RGB dataset [65]. (c1-c3) Deblurring on the GoPro [43] dataset. (d) Image translation (male/dog \rightarrow cat) on the CelebA-HQ [23] and AFHQ [9] datasets.

size of 1. ShadowDiffusion [14] uses SR3 [49] and Uformer [60], which has a higher PSNR but is also expected to be more computationally expensive. **Our RDDM with SM-Res requires only 4.8G of GPU memory for training.** Experiments in low-light enhancement, shadow removal, and deraining demonstrate the effectiveness of RDDM, enabling computationally-constrained researchers to utilize diffusion modeling for image restoration tasks.

Accelerating Convergence. The residual prediction in our RDDM helps the diffusion process to be more certain, which can accelerate the convergence process, e.g., fewer training iterations and higher performance in Table 11(b).

Failure case. We present some failure cases in Fig. 22.

C. Discussions, Limitations, and Further Work

Limitations. Our primary focus has been on developing a unified prototype model for image restoration and generation, which may result in certain performance limitations when compared to task-specific state-of-the-art methods. To further improve the performance of a specific task, potential avenues for exploration include using a UNet with a larger number of parameters, increasing the batch size, conducting more iterations, and implementing more effective training strategies, such as learning rate adjustments customized for different tasks. For the image generation task, although Table 2 showcases the development of an improved coefficient schedule, attaining state-of-the-art performance in image generation necessitates further investigation and experimentation. In summary, while we recognize the existing performance limitations for specific tasks, we are confident that our unified prototype model serves as a robust foundation for image restoration and generation.

Further Work. Here are some interesting ways to extend our RDDM.

1. In-depth analysis of the relationship between RDDM and curve/multivariate integration.
2. Development of a diffusion model trained with one set of pre-trained parameters to handle several different tasks.
3. Implementing adaptive learning coefficient schedules to reduce the sampling steps while improving the quality of the generated images.
4. Constructing interpretable multi-dimensional latent diffusion models for multimodal fusion, e.g., generating images using text and images as conditions.
5. Adaptive learning noise intensity (β_T^2) for an unknown new task.
6. Exploring residuals in distillation (e.g., introducing dual diffusion into consistency models [54]).