

Learning Efficient and Effective Trajectories for Differential Equation-based Image Restoration

Zhiyu Zhu, Jinhui Hou, Hui Liu, Huanqiang Zeng, and Junhui Hou, *Senior Member, IEEE*

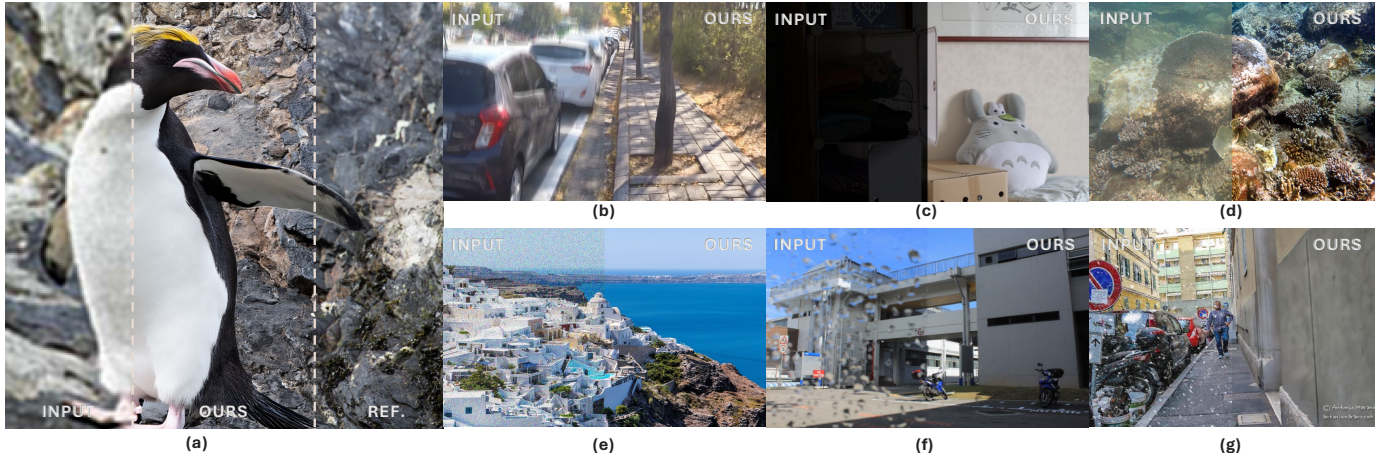


Fig. 1. Visual demonstration of the reconstruction performance for the proposed method on different image restoration tasks, where (a)–(e) indicate the tasks of single image super-resolution, deblurring, low-light enhancement, underwater enhancement, denoising, raindrop removal, and desnowing, respectively. Impressively, benefiting from the strong capacity of diffusion models, the proposed method can even generate more clear content than **Reference** images in (a).

Abstract—The differential equation-based image restoration approach aims to establish learnable trajectories connecting high-quality images to a tractable distribution, e.g., low-quality images or a Gaussian distribution. In this paper, we reformulate the trajectory optimization of this kind of method, focusing on enhancing both reconstruction quality and efficiency. Initially, we navigate effective restoration paths through a reinforcement learning process, gradually steering potential trajectories toward the most precise options. Additionally, to mitigate the considerable computational burden associated with iterative sampling, we propose cost-aware trajectory distillation to streamline complex paths into several manageable steps with adaptable sizes. Moreover, we fine-tune a foundational diffusion model (FLUX) with 12B parameters by using our algorithms, producing a unified framework for handling 7 kinds of image restoration tasks. Extensive experiments showcase the *significant* superiority of the proposed method, achieving a maximum PSNR improvement of 2.1 dB over state-of-the-art methods, while also greatly enhancing visual perceptual quality. Project page: <https://zhu-zhiyu.github.io/FLUX-IR/>.

This work was supported in part by the NSFC Excellent Young Scientists Fund 62422118, in part by the Hong Kong Research Grants Council under Grant 11218121, and in part by Hong Kong Innovation and Technology Fund MHP/117/21. Zhiyu Zhu and Jinhui Hou contributed to this paper equally. *Corresponding author: Junhui Hou.*

Zhiyu Zhu, Jinhui Hou, and Junhui Hou are with the Department of Computer Science, City University of Hong Kong, Hong Kong SAR (e-mail: zhiyuzhu2@my.cityu.edu.hk; jhhou3-c@my.cityu.edu.hk; jh.hou@cityu.edu.hk).

Hui Liu is with the School of Computing and Information Sciences, Saint Francis University, Hong Kong SAR (e-mail: h2liu@sfu.edu.hk).

Huanqiang Zeng is with the School of Engineering, Huaqiao University, Quanzhou 362021, China, and also with the School of Information Science and Engineering, Huaqiao University, Xiamen 361021, China (e-mail: zeng0043@hqu.edu.cn).

Index Terms—Image Restoration, Diffusion Models, Reinforcement Learning.

I. INTRODUCTION

IMAGE restoration involves the enhancement of low-quality images afflicted by various degradations like underwater and low-light conditions, raindrops, low-resolution, and noise to achieve high-quality outputs. It serves as a fundamental processing unit for visual recognition [1], [2], communication [3] and virtual reality [4]. Traditional image restoration methods typically rely on optimization procedures incorporating human priors such as sparsity, low rankness, and self-similarity [5], [6]. The emergence of deep learning techniques [7] has significantly transformed this domain. Initially delving into neural network architectures [8], image restoration has progressed beyond simple regression networks [8]–[11], exploring avenues like adversarial training [12], [13], algorithm unrolling [14]–[16] and flow-based methods [17]–[19].

Recently, a new category of generative models, namely diffusion models, has shown their strong potential for image synthesis and restoration [20], [21]. Generally, diffusion models construct probabilistic flow (PF) between a tractable distribution and the target distribution. The forward process typically involves incrementally introducing noise until reaching a manageable distribution, often a Gaussian. On the other hand, the reverse process can be obtained by maximizing the posterior of the forward Markov chain [22] or by sampling the reverse stochastic differential equation (SDE) [23], [24] or ordinary differential equation (ODE) [24].

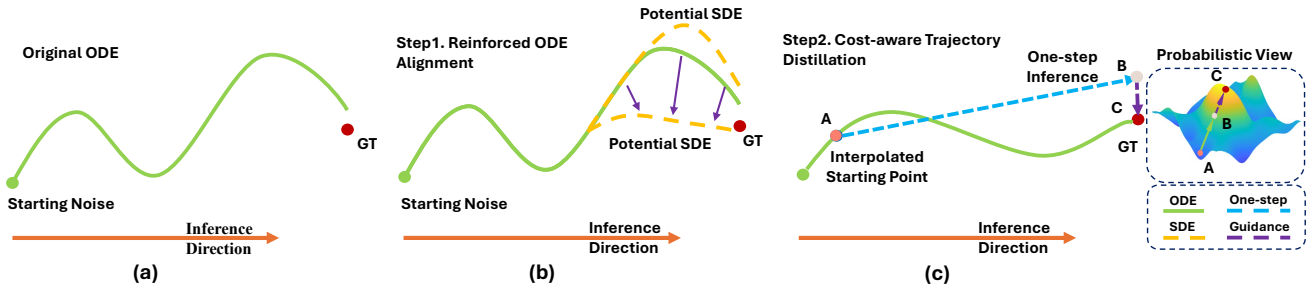


Fig. 2. Illustration of the workflow of the proposed method. Given a pre-trained diffusion model for image restoration, our trajectory optimization process contains the following two stages. (1) **Reinforced ODE alignment**, which aligns the deterministic ODE trajectory shown in (a) to the most effective modulated SDE trajectory, as shown in (b). (2) **Distillation cost-aware ODE acceleration** in (c), which achieves high-quality one-step inference via delicate designs based on the task, with the knowledge of the original pre-trained model preserved. Note that to preserve the original knowledge of the pre-trained diffusion model, we aim to find a trajectory with less modification of gradient $\frac{dX}{dt}$. Through theoretical analyses and experimental validations, we find that for image restoration tasks, degraded measurements usually lie in the low probability region from the probabilistic space of high-quality samples. Thus, we also utilize the input measurements as negative guidance to rectify the gradient of log-density. As shown in the sub-figure of the probabilistic view, **A**, **B**, and **C** correspond to the low-quality measurement, reconstructed sample, and reconstruction by a low-quality image as negative guidance, respectively. We refer the readers to Fig. 9, which illustrates the trajectories directly from the diffusion data points.

Three categories of methods stand out for harnessing the potent generative capabilities of diffusion models for image restoration. The first category [25]–[27] leverages the progressive integration nature of the differential equations to maximize the diffusion posterior on the scene of a low-quality image, then progressively reversing to the high-quality samples. Although sampling-based methods can directly take advantage of large, pre-trained models and get rid of network training, the posterior optimization process may take more time and computational resources [25], [28], and their performance is noncompetitive compared with supervised training diffusion model [29]. The second category incorporates the reconstruction outcomes from a fixed pre-trained diffusion model as a prior [30], subsequently refining these outcomes through trainable neural networks, akin to algorithm unrolling techniques. Lastly, the third set of methods [20], [29] trains a diffusion trajectory conditioned on low-quality samples [29] or establishes a direct linkage between low-quality and high-quality image distributions [31]. Since the diffusion trajectory is explicitly refined by training on the paired dataset, this kind of approach has the most potential and effectiveness.

Due to the probabilistic nature and separate training approach of diffusion models, the reverse generation trajectory might exhibit instability or chaos [32]. To tackle such an issue, some work tends to rectify the generation trajectories to be straight [32] or directly train a consistency model [33]. Although, this manner allows us to achieve adversarial training during the diffusion process, excessive regularization could significantly impair diffusion performance. As shown in Fig. 2, we aim to realign the diffusion trajectory towards the most effective path using a reinforcement learning approach. Furthermore, considering that the resulting trajectories may be complex and require extensive steps for sampling, we then propose a novel trajectory distillation process to alleviate this issue, which analyzes and lessens the cost of a diffusion model distillation process. Extensive experiments demonstrate the significant advantages of the proposed trajectory augmentation strategies on various image restoration tasks over state-of-the-art methods.

In summary, we make the following key contributions in this paper:

- we propose a novel trajectory optimization paradigm for boosting both the efficiency and effectiveness of differential equation-based image restoration methods;
- we theoretically examine the accumulated score estimation error in diffusion models for image restoration and introduce a reinforcement learning-based ODE trajectory augmentation algorithm that leverages the modulated SDE to generate potential high-quality trajectories;
- we improve the inference efficiency by introducing an acceleration distillation process, where we preserve the model’s original capacity via investigating the distillation cost and utilizing the low-quality images for initial state interpolation and diffusion guidance;
- we establish new records on various image restoration tasks, including de-raining, low-light, under-water enhancement, image super-resolution, image de-blurring, and image de-noising; and
- we calibrate a unified diffusion model for various image restoration tasks, based on the recent foundational diffusion model named *FLUX-DEV* with 12B parameters.

The remainder of this paper is organized as follows. Sec. II briefly reviews related work concerning diffusion models. Additionally, Sec. III offers essential mathematical formulations serving as the foundational backdrop for the proposed approach. Sec. IV details the proposed method, followed by comprehensive experiments in Sec. V. Finally, Sec. VI concludes this paper.

II. RELATED WORK

The differential equation-based deep generative model [24] represents a kind of learning strategy inspired by physical non-equilibrium diffusion processes [34]. It involves a forward diffusion process that progressively introduces noise into the data until a tractable distribution is reached, followed by a reversal to establish the data generation process. Ho *et al.* [22] firstly explored an effective diffusion formulation by parameterizing the reverse process as maximizing the posterior

of reverse steps. Song *et al.* [24], [35]–[37] generalized such discrete diffusion steps into a continuous formulation by stochastic differential equation. Thus, the diffusion process can then be treated as an integral of the corresponding differential equation. Besides, the ordinary differential equation introduced in [36] removes the additional noise in the reverse process, which enables more inference acceleration designs [38]–[40]. Xu *et al.* [41] introduce a Poisson noise-based diffusion model. Robin *et al.* [42] convert the diffusion process into the latent domain to achieve high fidelity high-resolution diffusion. Blattmann *et al.* [43], [44] extended the image diffusion model into a high-dimensional video diffusion process. Karras *et al.* [45] examined different formulations of the diffusion model and proposed a concise formulation. Dockhorn *et al.* [46] augmented the diffusion space by introducing an auxiliary velocity variable and constructing a diffusion process running in the joint space. Chen *et al.* [47] conducted theoretical convergence analysis on the score-based diffusion model.

In addition to the above, some diffusion-based methods bridge the distributions between different types of images. For example, Liu *et al.* [48] constructed a diffusion bridge by applying maximum likelihood estimation of latent trajectories with input from an auxiliary distribution. Li *et al.* [49] proposed an image-to-image translation diffusion model based on a Brownian Bridge diffusion process. Zhou *et al.* [50] utilized the h -transform to make a constraint of the endpoint of forward diffusion with formulations of reverse process derived by reformulating Kolmogorov equations.

A. Sampling Strategies for Accelerating Diffusion

Given the progressive noise reduction process, hundreds of steps are usually required to sample high-quality images. One solution to expedite inference involves crafting a refined sampling strategy. For example, Song *et al.* [51] pioneered the use of denoising diffusion implicit models (DDIM) to hasten diffusion sampling by disentangling each step in a non-Markov Chain fashion. Leveraging the semi-linear attributes of ODE and SDE formulations in diffusion models, Lu *et al.* [38]–[40] introduced a series of integral solvers with analytic solutions for the associated ODEs or SDEs. Zhang *et al.* [52] delved into the significant variance in distribution shifts and isolated an exponential variance component from the score estimation model, thereby mitigating discretization errors. Zhou *et al.* [53] introduced learnable solvers grounded in the mean value theorem for integrals. Xue *et al.* [54], [55] proposed accelerating diffusion sampling through an improved stochastic Adams method and precise ODE steps. Dockhorn *et al.* [56] advocated for higher-order denoising diffusion solvers based on truncated Taylor methods.

B. Trajectory Distillation-based Diffusion Acceleration

An alternative solution for accelerating generation involves directly adjusting diffusion trajectories. Liu *et al.* [32] introduced a method for straight flow regularization, which certifies the diffusion generation trajectory to be linear. Song *et al.* [33], [57] presented the consistency model, aligning each point on the trajectory directly with noise-free endpoints.

Kim *et al.* [58] introduced trajectory consistency distillation to regularize gradients that can consistently map to corresponding points on the trajectory. Moreover, Zhou *et al.* [59], [60] proposed to make distillation of a pre-trained diffusion model to a student one-step generator via measuring discrepancy by additional learned score function.

Moving beyond training the diffusion model into a single-step generator, segmenting the process into multiple sections, with each being linear, presents a viable solution for rapidly generating high-quality outcomes in diffusion models [61]–[63].

C. Diffusion-based Image Restoration

In the realm of diffusion-based image restoration, we outline works that fall into two main categories: those employing a training-free diffusion sampling strategy [25]–[27], [64], [65] and those relying on model training [20], [29], [66]–[68].

Within the first category, Kawar *et al.* [25] introduced an unsupervised posterior sampling method for image restoration using a pre-trained diffusion model. Chung *et al.* [26] proposed to regularize the intermediate derivative from the reconstruction process to improve image restoration. Wang *et al.* [27] decoupled the image restoration into range-null spaces and focused on the reconstruction of null space, which contains the degraded information. Zhu *et al.* [64] combined the traditional plug-and-play image restoration method into the diffusion process. Regarding the training-based methods, Luo *et al.* [29] proposed a mean-reverting SDE with its reverse formulation to boost diffusion-based image restoration. Jiang *et al.* [69] proposed a wavelet-based diffusion model for low-light enhancement. Yi *et al.* [70] introduced a dual-branches diffusion framework combining reflectance and illumination reconstruction process. Wang *et al.* [71] proposed a DDIM-inspired diffusion-based framework for the distillation of an image restoration model. Tang *et al.* [72] introduced a transformer-based model for underwater enhancement using diffusion techniques.

Based on the aforementioned analysis for related work, although there are many works introducing the diffusion model to the field of image restoration, there is a limited number of works considering the accumulated score-estimation error by PF characteristics of the diffusion model. Moreover, acceleration of the sampling speed is also an essential research topic for differential equation-based image restoration frameworks.

III. PRELIMINARY

In this section, we provide a succinct overview of the diffusion model techniques, acceleration strategies (including trajectory solvers and distillation), and boosting strategies through alignment, laying the foundation for the subsequent sections.

A. Diffusion Model

Here, we consider image restoration as a case study to briefly elucidate the diffusion model.

Given a low-quality measurement denoted as $\mathbf{Y} \in \mathbb{R}^{H \times W \times 3}$, image restoration methods strive to reconstruct corresponding high-quality outputs, represented as $\mathbf{X} \in \mathbb{R}^{H' \times W' \times 3}$. The probabilistic nature of this restoration process involves maximizing a posterior $\mathbf{P}_\theta(\mathbf{X}|\mathbf{Y})$, where θ the parameter set of the learnable neural module. The differential equation-based methods [25] generally learn to construct a probabilistic flow (PF), e.g., $\mathbf{P}(\mathbf{X}_0|\mathbf{X}_1, \mathbf{Y})$, $\mathbf{P}(\mathbf{X}_1|\mathbf{X}_2, \mathbf{Y})$, \dots , $\mathbf{P}(\mathbf{X}_{N-1}|\mathbf{X}_N, \mathbf{Y})$, to bridge the marginal distributions $\mathbf{P}(\mathbf{X}_0)$ (same as $\mathbf{P}(\mathbf{X})$) and $\mathbf{P}(\mathbf{X}_N)$, usually as $\mathcal{N}(\mathbf{0}, \sigma_N^2 \mathbf{I})$ with σ being the standard deviation of noise.

In particular, the formulation based on DDPM [21], [22] constructs the PF connecting $\mathbf{P}(\mathbf{X})$ and a standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, where the forward process is formulated as a diffusion process gradually substituting the data component with Gaussian noise, expressed as $\mathbf{X}_t = \bar{\alpha}_t \mathbf{X}_0 + \sqrt{(1 - \bar{\alpha}_t^2)} \epsilon_t$, $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Conversely, the reverse process is deduced by computing the posterior $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t+1}, \mathbf{X}_0) = \frac{\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{X}_t)\mathbf{P}(\mathbf{X}_t|\mathbf{X}_0)}{\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{X}_0)}$ with its resolution derived through the following ancestral sampling process:

$$\mathbf{X}_t \sim \mathcal{N}\left(\frac{1}{\sqrt{\alpha_{t+1}}}(\mathbf{X}_{t+1} - \frac{1 - \alpha_{t+1}}{\sqrt{1 - \bar{\alpha}_{t+1}}} \epsilon_{t+1}), \frac{1 - \bar{\alpha}_t}{1 - \bar{\alpha}_{t+1}}(1 - \alpha_{t+1})\mathbf{I}\right). \quad (1)$$

As the noise component ϵ_{t+1} is typically challenging to handle, it is represented by a neural network parameterization $\epsilon_\theta(\mathbf{X}_{t+1}, t + 1)$, trained using a loss term $\mathcal{L}_{DDPM} = \|\epsilon - \epsilon_\theta(\bar{\alpha}_t \mathbf{X}_0 + \sqrt{(1 - \bar{\alpha}_t^2)} \epsilon_t)\|_2^2$. By iteratively sampling Eq. (1), we can deduce \mathbf{X}_0 from the sample \mathbf{X}_N of a tractable distribution. Essentially, these sampling procedures serve to connect samples from two distributions along a high-dimensional trajectory, when minimizing the step size from the discrete to continuous spaces, i.e., $\Delta \bar{\alpha} \rightarrow d\alpha$.

Score-based models [24] achieve diffusion model generalization by reformulating the forward process via the following SDE:

$$d\mathbf{X} = f(\mathbf{X}, t)dt + g(t)d\omega, \quad (2)$$

where $t \in [\delta, T]$ denotes the timestamp of the diffusion process, with δ serving as a small number for numerical stability. The solution to the reverse SDE can be expressed as

$$d^{(s)}\mathbf{X} = [f(\mathbf{X}, t) - g^2(t)\nabla_x \log \mathbf{P}(\mathbf{X})]dt + g(t)d\omega. \quad (3)$$

Furthermore, by redefining Kolmogorov's forward equation [24], an equivalent reverse ODE formulation emerges:

$$d^{(o)}\mathbf{X} = \left[f(\mathbf{X}, t) - \frac{1}{2}g^2(t)\nabla_x \log \mathbf{P}(\mathbf{X}) \right] dt. \quad (4)$$

Based on the diffusion strategies for $f(\mathbf{X}, t)$ and $g^2(t)$, the score-based diffusion models can be categorized into two variants, namely Variance Preserving (VP) or Variance Exploding (VE). Specifically, VP entails $f(\mathbf{X}, t) = \frac{d \log \alpha_t}{dt} \mathbf{X}$ and $g^2(t) = \frac{d\sigma_t^2}{dt} - 2 \frac{d \log \alpha_t}{dt} \sigma_t^2$, where $\sigma_t = \sqrt{1 - \alpha_t^2}$, while VE involves $f(\mathbf{X}, t) = \mathbf{0}$ and $g^2(t) = \frac{d\sigma_t^2}{dt}$.

B. Integral-Solver

To derive the reconstruction sample, we can calculate the integral of the reverse ODE trajectory in Eq. (4) as $\hat{\mathbf{X}}_0 = \mathbf{X}_T + \int_T^0 [f(\mathbf{X}, t) - \frac{1}{2}g^2(t)\nabla_x \log \mathbf{P}(\mathbf{X})] dt$. Given the deterministic nature of the entire process, an analytical formulation can be leveraged. DPM-Solver [38], [39] first introduces an exact solution based on the semi-linear property of the diffusion model, expressed as

$$\mathbf{X}_{t-\Delta t} = \frac{\alpha_{t-\Delta t}}{\alpha_t} \mathbf{X}_t - \alpha_{t-\Delta t} \int_{\lambda_t}^{\lambda_{t-\Delta t}} e^{-\lambda} \hat{\epsilon}_\theta(\hat{\mathbf{X}}_\lambda, \tau) d\lambda, \quad (5)$$

where λ denotes the log-SNR (Signal-to-Noise Ratio), i.e., $\lambda := \log(\frac{\alpha_t}{\sigma_t})$. By computing this integral using various Order Taylor series for the non-linear term $\hat{\epsilon}_\theta(\hat{\mathbf{X}}_\lambda, \lambda)$, we can finally derive the result.

C. Rectified Flow & Trajectory Distillation

As a kind of naturally easily sampled model, the rectified flow-based method learns a straight flow between the random noise and target image domain, which can be formulated as

$$d\mathbf{X}_t = (\mathbf{X}_0 - \mathbf{X}_T)dt. \quad (6)$$

Similarly, distillation-based enhancement [33], [73] of diffusion models aims to regularize the inherent trajectory pattern to derive the diffusion models, which can be easily sampled. The nature of such kind of methods is to regularize

$$\mathcal{L}_{slope} = \mathcal{D}(\text{Solver}(\mathbf{X}_t, t, \theta) | \text{Solver}(\mathbf{X}_s, s, \phi)), \quad (7)$$

where $\text{Solver}(\cdot)$ denotes the ODE/SDE solvers, solving the integral with outputting $\hat{\mathbf{X}}_0$ from an arbitrary intermediate state \mathbf{X}_s (resp. \mathbf{X}_t) with timestamp s (resp. t); θ (resp. ϕ) represents the model weights of the student (resp. teacher) network; ϕ can be exponential moving average between the student net and original pre-trained model; and $\mathcal{D}(\cdot | \cdot)$ indicates the divergence metric for two inputs, such as L1/L2 norm and LPIPS [74].

D. Reinforcement Tuning of Diffusion Models

Given that the diffusion model is trained using discrete trajectory points, the accrued errors during the inference process are typically overlooked. Promising solutions to address these issues are found in reinforcement learning-based methods [75], [76], which optimize entire trajectories in a decoupled fashion. This approach involves separating the inference Markov Chain and enhancing each step towards a direction of high reward, expressed as

$$\nabla_\theta \mathcal{J}(\mathbf{X}_t) = - \int_{P_{\theta'}} \frac{\nabla_\theta P_\theta(\mathbf{X}_t)}{P_{\theta'}(\mathbf{X}_t)} \mathcal{R}(\hat{\mathbf{X}}_0) + \kappa \nabla_\theta \mathcal{D}(P_{\theta'} | P_\theta), \quad (8)$$

where $\mathcal{R}(\cdot)$ denotes the reward function assessing the reconstruction quality or trajectory performance $\mathcal{J}(\cdot)$ the optimization objective, and κ the weight assigned to the regularization term.

Algorithm 1 Reinforced ODE Trajectory Learning

-
- 1: **Repeat**
 - 2: $\mathbf{X}_0 \sim q(\mathbf{X}_0)$, $\tau \sim \text{Uniform}(\{1, 2, \dots, T\})$ and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$;
 - 3: For score (resp. rectified flow)-based method, we generate N reverse trajectories $\{\mathcal{T}_i\}_{i=1}^N$ with the τ^{th} step as M-SDE via Eqs. (11) or (12) (resp. Eq. (13)), other steps as ODE by Eq. (5) (resp. Eq. (6))
 - 4: Find the best trajectory \mathcal{T}_i from N M-SDE trajectories, whose performance is measured by the reward function $\mathcal{R}(\cdot)$.
 - 5: Calculate $\text{Solver}(d^{(\circ)}\mathbf{X}, t_\tau, t_{\tau-1})$ via Eq. (5).
 - 6: Take gradient descent step on $\nabla_{\theta}\mathcal{L}_A(t_i)$ by Eq. (16);
 - 7: **Until** converged
-

IV. PROPOSED METHOD

A. Overview

Learning effective and efficient trajectories is critical for differential equation-based image restoration. The inherent Markov Chain property of the diffusion model complicates the precise regularization of the entire trajectory.

In this work, we propose a reinforcement learning-inspired alignment process in Sec. IV-B for improving effectiveness. Specifically, by projecting the accumulated error back to different steps, we theoretically reason the necessity of adaptively modulating the noise intensity of differential equations. Based on that, we align the ODE trajectory with the most effective alternative drawn from multiple candidate trajectories that are sampled by solving different modulated-SDEs (M-SDEs).

Subsequently, in Sec. IV-C, we propose a *cost-aware* trajectory distillation strategy for boosting efficiency. This strategy leverages intrinsic characteristics of image restoration tasks to lessen the distillation burden. We utilize the low-quality image as a coarse estimation and negative guidance and give corresponding theoretical analysis. Note that the proposed strategy can be adapted to both score-based [24] and rectified flow-based [32] diffusion models. Due to the page limits, we mainly utilize the formulation of a score-based diffusion model in this paper, and we also refer readers to the *appendixes* for extensive theoretical elaborations.

B. Reinforcing ODE with Modulated Differential Equations

Diffusion models are usually trained through individual steps originating from the decoupled PF. However, during the inference phase, they usually operate in a progressively noise-removing manner. Due to inherent score function errors, these models often accumulate inaccuracies. While optimizing the entire trajectory could mitigate this issue, traditional diffusion models, even with ODE-solvers, struggle to yield satisfactory outcomes within a limited number of steps. Differently, we propose to align the learned trajectories with the most effective alternatives through reinforcement learning. Generally, our reinforcement learning-inspired approach aims to maximize the expectation of a reward function as

$$\nabla_{\theta}\mathcal{L} = -\nabla_{\theta}\mathbb{E}_{x \sim \mathbf{P}_{\theta}(\mathbf{X})}\mathcal{R}(\mathbf{X}). \quad (9)$$

Given the deterministic nature of ODE sampling, optimizing the ODE towards the optimal trajectory could lead to maximizing the likelihood of obtaining the optimal sample ($\mathbf{P}_{\theta}(\mathbf{X}_{\text{optim}}) \rightarrow 1$). Unfortunately, the deterministic property of ODE trajectory also makes it struggle to generate diverse trajectories given a randomly initialized starting noise point. It cannot fulfill the reinforcement learning needs, which requires diverse alternatives for measuring and selecting a better optimization direction. Thus, we argue a **potential solution** for reinforcement training-based ODE trajectory augmentation should involve leveraging SDE to produce diverse restoration trajectories and aligning the deterministic ODE trajectory with the most effective SDE trajectory. However, the SDE is formulated with a fixed noise intensity level, which is too rigid and inflexible to adapt to different conditions. As shown in *Appendix C*, we give theoretical proof that for an image restoration trajectory ended with \mathbf{X}_0 , we need to adjust the intensity of injected noise, conditioned on the reconstruction error $\|\mathbf{X}_0 - \mathbf{X}_0^*\|_2$ and corresponding timestamp t . Thus, a more flexible and controllable SDE formulation is necessary for the ODE trajectory correction.

Building upon the preceding analyses, we utilize a *modulated SDE*, adjusted by a tunable factor $\gamma_{\phi}(\|\mathbf{X}_0 - \mathbf{X}_0^*\|_2, t) > 0$, abbreviated as γ_{ϕ} , which is a small MLP parameterized with ϕ . We also refer readers to *Appendix B* for the proof of Modulated-SDE (M-SDE) and its solver. Serving as a versatile sampling trajectory, M-SDE encompasses ODE, SDE, and DDIM-like sampling, each tailored through distinct parameterizations [52], [77]. In what follows, we give the detailed formulations of M-SDE with different diffusion models.

Score-based Diffusion. The inverse M-SDE of the diffusion forward process, as described in Eq. (2), can be explicitly formulated as

$$d^{(\gamma)}\mathbf{X} = [f(x, t) - \frac{1 + \gamma_{\phi}^2}{2}g^2(\mathbf{X}, t)\nabla_x \log p(\mathbf{X}|\mathbf{Y})]dt + \gamma_{\phi}g(\mathbf{X}, t)d\hat{\omega}. \quad (10)$$

Moreover, we give the analytical formulation of integral solvers (see *Appendix B-B* for the proof) for the VP-M-SDE and VE-M-SDE as

$$\begin{aligned} \mathbf{X}_{t-\Delta t} &= \frac{\alpha_{t-\Delta t}}{\alpha_t}\mathbf{X}_t - [1 + \gamma_{\phi}^2]\epsilon_{\theta}\left(\frac{\alpha_{t-\Delta t}}{\alpha_t}\sigma_t - \sigma_{t-\Delta t}\right) \\ &\quad - \sqrt{2}\gamma_{\phi}\epsilon\alpha_{t-\Delta t}\sqrt{\log\frac{\alpha_{t-\Delta t}}{\alpha_t}}, \end{aligned} \quad (11)$$

$$\begin{aligned} \mathbf{X}_{t-\Delta t} &= \mathbf{X}_t - [1 + \gamma_{\phi}^2]\epsilon_{\theta}(\sigma_t - \sigma_{t-\Delta t}) \\ &\quad - \sqrt{2}\gamma_{\phi}\epsilon\sqrt{\sigma_t^2 - \sigma_{t-\Delta t}^2}. \end{aligned} \quad (12)$$

Rectified Flow. We have proven that the following integral formulation of M-SDE has the same PF with the original ODE formulation as Eq. (6) for the rectified flow-based model,

$$\mathbf{X}_{t-\Delta t} = \frac{[\mathbf{X}_t - \alpha_t\Delta t\frac{d\mathbf{X}_t}{dt} - \beta_k\epsilon]}{(1 + \alpha_t\Delta t - t) + \sqrt{(t - \alpha_t\Delta t)^2 + \beta_k^2}}, \quad (13)$$




Fig. 3. Visualization of distillation costs, where (a) illustrates the total distillation cost (summation from different steps) across trajectories with varying numbers of inference steps, and (b) provides a detailed illustration of the distillation cost for each step from two different inference step setting, with the step-index ranging from the high-noise to the low-noise regions.

where α_t is a scalar ($\alpha_t > 1$), and β_k is formulated as

$$\beta_k = \sqrt{\frac{(t - \Delta_t)^2 [1 - (t - \alpha\Delta_t)]^2}{[1 - (t - \Delta_t)]^2} - (t - \alpha\Delta_t)^2}. \quad (14)$$

The M-SDE of Eq. (13) can be formulated as

$$d\mathbf{X} = \mathbf{X}_0 dt + \frac{t - \alpha_t}{1 - t} \mathbf{X}_T dt + \sqrt{2(\alpha_t - 1)} d\omega. \quad (15)$$

We refer readers to *Appendix A* for the detailed proof. Then, to close a certain step of M-SDE and ODE, we introduce the following alignment loss as

$$\mathcal{L}_A(t_i) = \mathcal{D} \left(\text{Solver} \left(d^{(\gamma)} \mathbf{X}, t_i, t_{i-1} \right), \text{Solver} \left(d^{(o)} \mathbf{X}, t_i, t_{i-1} \right) \right), \quad (16)$$

where $\text{Solver}(d\mathbf{X}, t_i, t_{i-1})$ indicates the integral solver for a the derivative expression $d\mathbf{X}$ from timestamp t_i to t_{i-1} , and we adopt L2 norm to achieve the divergence measurement $\mathcal{D}(\cdot, \cdot)$. Algorithm 1 summarizes the entire optimization process. Through such a reinforcement learning process, the proposed method can even trained with some in-differentiable metrics, e.g., NIQE.

C. Distillation Cost-aware Diffusion Sampling Acceleration

In this section, we first explicitly model the cost-value of the diffusion model distillation process. Then, based on both empirical and theoretical results of distillation cost, we propose a novel trajectory distillation pipeline to manage high-quality few-step inference, which consists of a multi-step distillation strategy and a negative guidance policy from low-quality images.

Distillation Cost Analysis. To accelerate diffusion sampling, model distillation [32], [33], [78] condenses the knowledge from precise and multi-step sampling outcomes into shorter procedures, like the direct regularization in Eq. (7). However, this condensed distillation process may require adjustments to the initial neural parameter distributions, potentially decreasing network performance. To derive efficient and effective reconstruction, we argue that a good distillation training method for diffusion models should not only enable precise integration with fewer inference steps but also leverage the original Neuron-ODEs while minimizing alterations to neural network parameters. To this end, we propose a distillation cost-aware diffusion acceleration strategy that leverages the special characteristics of image restoration tasks to lessen the learning burden of the diffusion network.

Algorithm 2 Diffusion Acceleration Distillation

- 1: **Repeat**
- 2: $\mathbf{X}_0 \sim q(\mathbf{X}_0)$ and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$.
- 3: Interpolate the diffusion starting state $\mathbf{X}_{T-\delta}$ via Eq. (18).
- 4: Calculate $\text{Solver}(d^{(o)}\mathbf{X}, T - \delta, 0)$ via Eq. (5).
- 5: Take gradient descent step on $\nabla_{\theta} \mathcal{L}_D$ by Eq. (19).
- 6: **Until** converged

Algorithm 3 Inference Process of the Augmented Image Restoration Diffusion Models

- 1: $\epsilon \sim \mathcal{N}(0, \mathbf{I})$
- 2: Interpolate the starting state $\mathbf{X}_{T-\delta}$ via Eq. (18);
- 3: **for** $t \leftarrow T - \delta$ to 0 **do**
- 4: Derive the noise $\hat{\epsilon}_\theta$ with low-quality guidance via Eq. (20).
- 5: Calculate the reverse diffusion result $\mathbf{X}_{t-\Delta t}$ by substituting $\hat{\epsilon}_\theta$ and \mathbf{X}_t into Eq. (5);
- 6: **end for**

Specifically, to quantify the extent of neural network adjustments, we introduce the trajectory distillation cost defined as

$$\mathcal{C} = \sum_{i=1}^k \left\| \tilde{\epsilon} \left(\frac{\mathbf{X}_i^{i-1}}{t_i^{i-1}} \left| \frac{d\mathbf{X}_\epsilon}{dt} \right. \right) - \epsilon_\theta(\mathbf{X}_{t_i}, t_i) \right\|_2, \quad (17)$$

where k refers to the total number of steps in the student model, which is also known as the number of function evaluations (NFE), function $\tilde{\epsilon}(\mathbf{A}|\mathbf{B})$ denotes the inverse of $\epsilon(\cdot)$, calculating the corresponding ϵ value via making the \mathbf{B} term identical to the \mathbf{A} term, \mathbf{X}_i^{i-1} symbols $\mathbf{X}_{i-1} - \mathbf{X}_i$. We refer readers to *Appendix D-A* for the detailed formulation process of $\tilde{\epsilon}$ and \mathcal{C} .

To investigate the characteristics of the distillation cost of a typical one-step and two-step diffusion model, we calculate the corresponding value of Eq. (28) for k from 1 to 16 in Fig. 3-(a) and detailed illustration of distillation cost of each step in Fig. 3-(b). With acknowledgment of the aforementioned distillation-based prior, we can draw the following observations:

- the distillation cost exhibits a negative correlation with the total number of steps (Fig. 3-(a));
- initial steps contribute significantly to the overall distillation cost (Fig. 3-(b)); and
- the gradients of log-distribution of low-quality images can serve as guidance directions for the reverse diffusion process, as illustrated in Fig. 2.

The first observation supports the superiority of recent multi-step distillation models [58], [61], [79] over their single-step counterparts. Nonetheless, for efficient inference of the diffusion model, the inclusion of few or even single-step models remains essential. Furthermore, in *Appendix D*, we delve into the theoretical exploration of the existence of cost-effective multi-step distillations. Moreover, drawing on the second and third observations, we can outline the subsequent steps to alleviate the substantial learning burden associated with the distillation process.

Interpolation of the Initial State. Regarding the second observation, during the initial inference stages, diffusion models must synthesize data from pure noise—a challenging yet crucial aspect of the generation process. According to the characteristics of image restoration tasks, we have the low-quality image, which contains the same content as the target image, which can serve as the coarse estimation of low-quality image. Thus, to alleviate this burden, we propose synthesizing the noised latent representation as

$$\mathbf{X}_{T-\delta} = \alpha_{T-\delta} \mathbf{Y} + \sigma_{T-\delta} \epsilon, \quad (18)$$

where $\delta \geq 0$ is chosen sufficiently small to ensure SNR $(\frac{\alpha_{T-\delta}}{\sigma_{T-\delta}})$ to be sufficiently small. Additionally, we theoretically analyze this noised latent interpolation method in *Appendix E*. Subsequently, we can train our acceleration distillation neural network via the following loss term:

$$\mathcal{L}_D = \mathcal{D} \left(\text{SOLVER} \left(d^{(o)} \mathbf{X}, T - \delta, 0 \right), \bar{\mathbf{X}} \right), \quad (19)$$

where $\bar{\mathbf{X}}$ indicates the reference high-quality images.

Low-quality Images as Sampling Guidance. In the realm of image restoration, our objective is to reconstruct a high-quality image \mathbf{X} from a low-quality measurement \mathbf{Y} . Moreover, based on the previous observation, we propose leveraging low-quality images as sampling guidance, which can amplify the positive restoration components from the diffusion model, to ease the learning burden of the diffusion model. From a probabilistic view, the training of diffusion-based image restoration models strives to improve the alignment of final reconstruction \mathbf{X}_0 with reference image \mathbf{X} under the given condition \mathbf{Y} , i.e., improving $\mathbf{P}(\mathbf{X}_0 = \mathbf{X} | \mathbf{Y})$. Let $\mathbf{Y} = \mathcal{H}(\mathbf{X})$ with $\mathcal{H}(\cdot)$ being the degradation function. Considering that the degradation Jacobian matrix $\frac{\partial \mathcal{H}(\mathbf{X})}{\partial \mathbf{X}}$ often deviates from the identity matrix, there exists a notable discrepancy between the $\mathbf{P}(\mathbf{X}_0)$ and $\mathbf{P}(\mathbf{Y}_0)$. To address this, we propose a parameterized score function guided by the following principles:

$$\hat{\epsilon}_\theta = (1 + w) \epsilon_\theta - w \tilde{\epsilon}, \quad (20)$$

where $\tilde{\epsilon}$ indicates predicted noise by maximizing the posterior likelihood on low-quality images, serving as guidance for the diffusion process. w denotes a scalar for guidance strength. Specifically, it can be calculated by inverting the integral process as

$$\tilde{\epsilon} = \frac{\frac{\alpha_0}{\alpha_t} \mathbf{X}_t - \mathbf{Y}}{\sqrt{1 - \alpha_t^2 \alpha_0} - \sqrt{1 - \alpha_0^2}}. \quad (21)$$

Moreover, we explain the training and inference process in Algorithms 2 and 3 in detail, respectively.

Remark. During the distillation phase, we harness the intrinsic characteristics of the image restoration task to mitigate the substantial challenges associated with few-step inference. Specifically, we implement interpolation of the initial state and negative guided sampling to address two critical issues: the significant estimation errors encountered during the high-noise initial stage and the complexities associated with modeling the probabilistic data space. These features distinguish our

approach from existing techniques. Besides, our framework aims to improve the overall efficiency and effectiveness of inference processes in complex environments, thereby paving the way for future research and applications.

V. EXPERIMENTS

In this section, we thoroughly assess the proposed trajectory optimization strategies across various image restoration tasks. Initially, we confirm the task-specific enhancement capabilities by training individual smaller networks for tasks such as de-raining, low-light enhancement, and underwater enhancement in Sec. V-B. Furthermore, in Sec. V-C, we produce a unified perceptual image restoration network by fine-tuning the state-of-the-art T2I foundational diffusion framework *FLUX-DEV* [103], which has 12B parameters.

A. Experimental Settings

1) *Datasets:* We employ the following commonly used benchmark datasets to conduct experiments:

- **Image De-raining.** Rain-drop [104] and heavy Rain datasets [105] are utilized. Rain-drop consists of 1,119 image pairs. Each pair includes the same background scene, with one image degraded by raindrops and the other image free from raindrops. The images were captured using two identical glass panels, one sprayed with water and the other kept clean. The dataset encompasses diverse background scenes and raindrop patterns and was obtained using a Sony A6000 and a Canon EOS 60. The heavy rain dataset contains 9,000 and 1,800 synthetic images from [104], respectively. We utilized a subset of 8250 and 750 images for training and testing, respectively.
- **Low-light Image Enhancement.** LOLv1 [95] contains 485 low/normal-light image pairs for training and 15 pairs for testing, captured at various exposure times from the real scene. LOLv2 [106] is split into two subsets: LOLv2-real and LOLv2-synthetic. LOLv2-real comprises 689 pairs of low-/normal-light images for training and 100 pairs for testing, collected by adjusting the exposure time and ISO. LOLv2-synthetic was generated by analyzing the illumination distribution of low-light images, consisting of 900 paired images for training and 100 pairs for testing.
- **Underwater Image Enhancement.** The UIEB dataset [81] consists of 950 real-world underwater images and includes two subsets: 890 raw underwater images with the corresponding high-quality reference images and 60 challenging underwater images.
- **Image Desnowing.** The Snow100K-L [107] dataset includes 100k synthesized snowy images with corresponding snow-free reference images and snow masks. We randomly selected 1872 (resp. 601) images forming the training (resp. testing) dataset for image de-snowing.
- **Image Super-Resolution.** DIV2K [108] is a popular single-image super-resolution dataset that contains 1,000 images with different scenes and is split into 800 for training, 100 for validation, and 100 for testing. It




Fig. 4. Visual comparison of underwater image enhancement on U45 [80] and UIEB [81] datasets. U45 (**top**): (a) low-quality input, (b) CycleGAN [82], (c) MLE [83], (d) HCLR [84], (e) SemiUIR [85], (f) Ours($k = 1$) and (g) Ours($k = 10$). UIEB (**bottom**): except (b) reference image, the remaining columns are the same as those of U45. k is the number of inference steps, commonly referred to as the number of function evaluations (NFE).

TABLE I

QUANTITATIVE COMPARISONS OF DIFFERENT METHODS ON UNDERWATER IMAGE ENHANCEMENT. THE BEST AND SECOND-BEST RESULTS ARE HIGHLIGHTED IN **BOLD** AND UNDERLINED, RESPECTIVELY. “ \uparrow ” (RESP. “ \downarrow ”) MEANS THE LARGER (RESP. SMALLER), THE BETTER. “NFE” DENOTES THE NUMBER OF FUNCTION EVALUATIONS, WHICH CAN BE INTERPRETED AS THE INFERENCE STEPS.

Method	NFE	UIEBD				C60		U45	
		PSNR \uparrow	SSIM \uparrow	UCIQE \uparrow	UIQM \uparrow	UCIQE \uparrow	UIQM \uparrow	UCIQE \uparrow	UIQM \uparrow
Water-Net [81] TIP'19	1	16.31	0.797	0.606	2.857	0.597	2.382	0.599	2.993
Ucolor [86] TIP'21	1	21.09	0.872	0.580	3.048	0.553	2.482	0.573	3.159
MLLE [83] TIP'22	1	19.56	0.845	0.588	2.646	0.569	2.208	0.595	2.485
NAFNet [87] ECCV'22	1	22.69	0.870	0.592	3.044	0.559	2.751	0.594	3.087
Restormer [10] CVPR'22	1	23.70	0.907	0.599	3.015	0.570	2.688	0.600	3.097
SemiUIR [85] CVPR'23	1	24.31	0.901	0.605	3.032	0.583	2.663	0.606	3.185
HCLR-net [84] IJCV'24	1	25.00	0.925	0.607	3.033	<u>0.587</u>	2.695	0.610	3.103
Proposed Method	10	25.08	0.913	0.615	3.142	0.571	3.663	0.612	4.282
	2	<u>25.94</u>	<u>0.937</u>	<u>0.618</u>	<u>3.136</u>	0.576	<u>3.774</u>	<u>0.613</u>	<u>4.291</u>
	1	26.25	0.938	0.623	3.135	0.582	3.814	0.617	4.413

was collected for NTIRE2017 and NTIRE2018 Super-Resolution Challenges in order to encourage research on image super-resolution with more realistic degradation. Meanwhile, Flickr2K [109] consists of 2650 pairs with high-quality 2K images and corresponding degraded images.

- **Image Deblurring.** The GoPro [110] dataset contains 3,214 blurred images with a size of 1280×720 . The

images are divided into 2,103 training images and 1,111 test images. The dataset consists of pairs of a realistic blurry image and the corresponding ground truth sharp images that are obtained by a high-speed camera.

- **Image Denoising.** The noisy images were derived by randomly corrupting the aforementioned high-quality SR datasets with Gaussian noise by a standard deviation of 50.

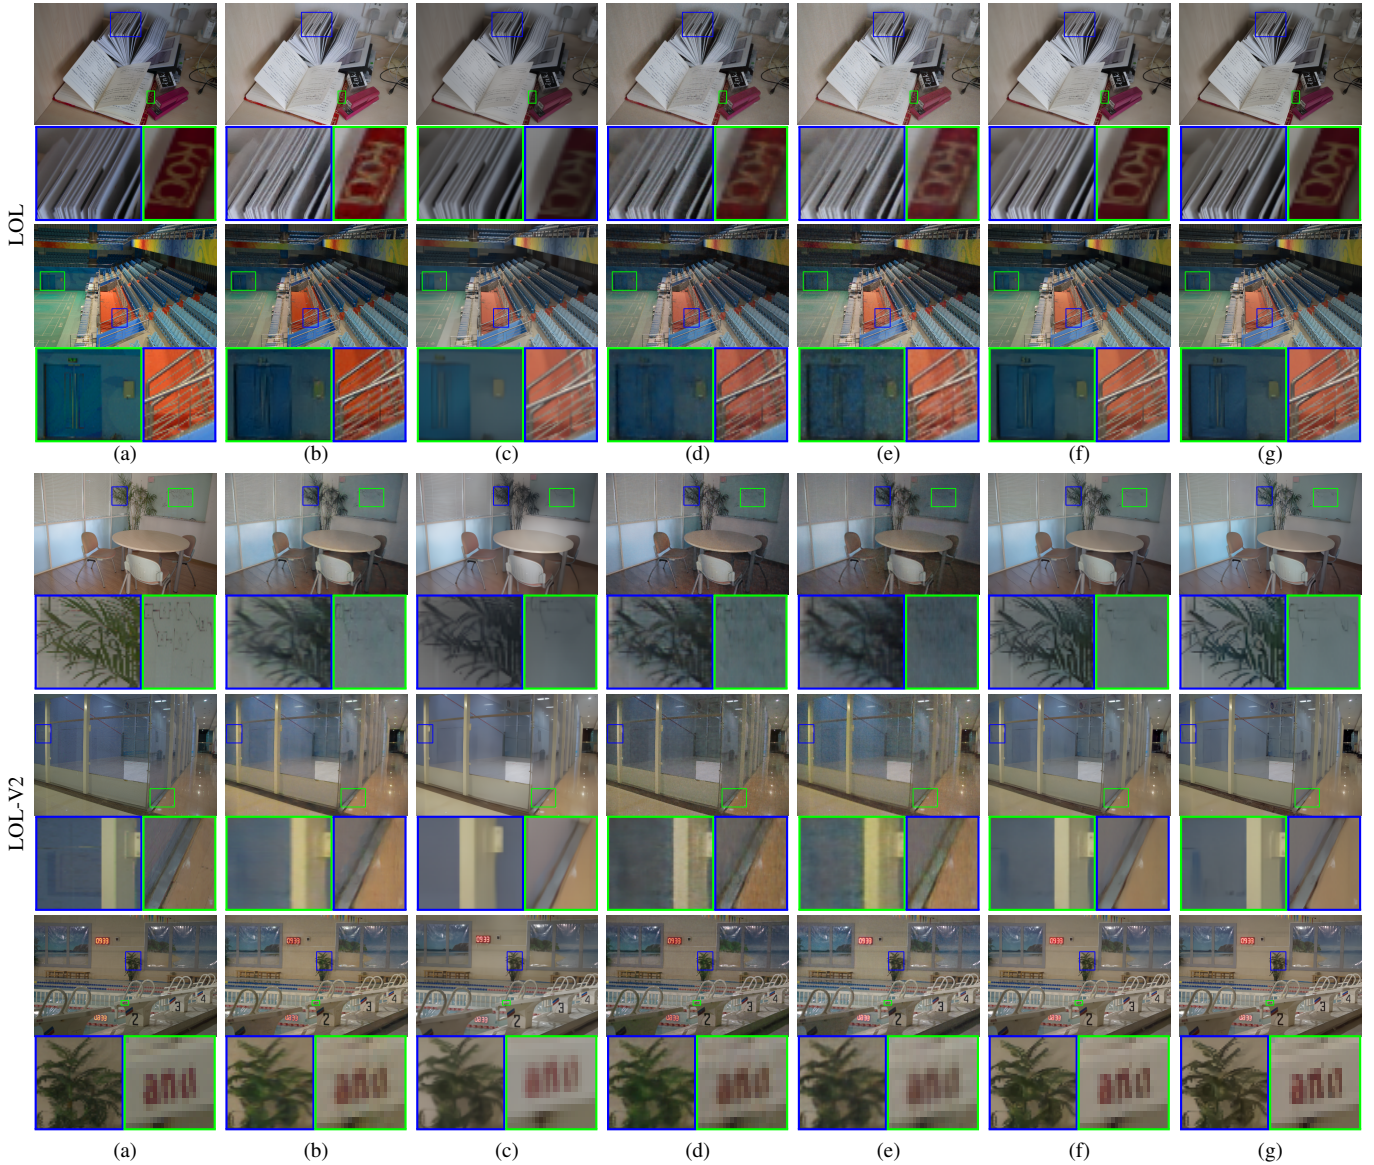


Fig. 5. Visual comparison of low-light enhancement on LOL and LOLV2 datasets. LOL (**top**): (a) reference images, (b) CID [88], (c) LLFlow [89], (d) RetinexFormer [90], (e) LLFormer [91], (f) Ours ($k = 1$), (g) Ours ($k = 10$). LOLV2 (**bottom**): except (b) SNR-Aware [92], the remaining columns are the same as those of LOL. Below each figure, we also visualize zoom-in regions marked by the blue and green boxes. k is the number of inference steps, commonly referred to as the number of function evaluations (NFE).

2) *Implementation details*: We conducted image restoration experiments to assess the efficacy of our method under two scenarios:

- **Task-specific restoration**, where a diffusion model is trained for each particular degradation. Specifically, we trained three diffusion networks based on GSAD [102] for low-light enhancement, de-raining, and underwater enhancement tasks. During this process, we utilized PSNR as the sole reward for the ODE alignment step. The model was trained on an RTX 3090 for 30,000 iterations for both alignment and acceleration, employing Adam optimizer with a learning rate of $5e^{-5}$, a training patch size of 256^2 , and a batch size of 2.
- **Unified restoration**, where a single diffusion model is trained to handle various types of degradation. Here,

we constructed the unified image restoration network based on FLUX-DEV [103]. Specifically, we first trained a low-quality U-Net encoder to make its feature map consistent with those from high-quality images. Then, we further trained a Control-Net by *XFLUX* [111]. We began by training a low-quality U-Net encoder to ensure its feature map aligned with those from high-quality images. Subsequently, we trained a control network using *XFLUX* [111]. During this training, we integrated features from *XFLUX* and the pre-trained encoder into the DiT structure of FLUX to enable rapid adaptation of our FLUX-IR framework. The model was trained on five NVIDIA H800 GPUs for 20,000 iterations, utilizing the Adam optimizer, a learning rate of $5e^{-5}$, a training patch size of 1024^2 , and a batch size of 128 (with gradient accu-

TABLE II

QUANTITATIVE COMPARISONS OF DIFFERENT METHODS ON **LOW-LIGHT ENHANCEMENT**. THE BEST AND SECOND-BEST RESULTS ARE HIGHLIGHTED IN **BOLD** AND UNDERLINED, RESPECTIVELY. “↑” (RESP. “↓”) MEANS THE LARGER (RESP. SMALLER), THE BETTER. “NFE” DENOTES THE NUMBER OF FUNCTION EVALUATIONS, WHICH CAN BE INTERPRETED AS THE INFERENCE STEPS.

Method	LOL				LOL-v2			
	NFE	PSNR↑	SSIM↑	LPIPS↓	NFE	PSNR↑	SSIM↑	LPIPS↓
Zero-DCE [93] CVPR’20	1	14.861	0.562	0.335	1	18.059	0.580	0.313
EnlightenGAN [94] TIP’21	1	17.483	0.652	0.322	1	18.640	0.677	0.309
RetinexNet [95] BMVC’18	1	16.770	0.462	0.474	1	18.371	0.723	0.365
DRBN [96] CVPR’20	1	19.860	0.834	0.155	1	20.130	0.830	0.147
KinD [97] MM’19	1	20.870	0.799	0.207	1	17.544	0.669	0.375
KinD++ [98] IJCV’20	1	21.300	0.823	0.175	1	19.087	0.817	0.180
MIRNet [99] TPAMI’22	1	24.140	0.842	0.131	1	20.357	0.782	0.317
LLFlow [89] AAAI’22	1	25.132	0.872	0.117	1	26.200	0.888	0.137
Retinexformer [90] ICCV’23	1	27.180	0.850	-	1	27.710	0.856	-
PyDiff [100] IJCAI’23	4	27.090	0.879	0.100	-	-	-	-
LLFormer [91] AAAI’23	1	25.758	0.823	0.167	1	26.197	0.819	0.209
SNR-Aware [92] CVPR’22	1	26.716	0.851	0.152	1	27.209	0.871	0.157
LLFlow-L-SKF++ [101] TPAMI’24	1	26.894	0.879	0.095	1	28.453	0.909	0.117
GSAD [102] NeurIPS’23	20	27.839	0.877	0.091	10	28.818	0.895	0.095
Proposed Method	10	28.581	0.883	0.084	10	29.535	0.898	0.086
	2	<u>28.360</u>	0.886	0.088	2	<u>29.794</u>	0.898	<u>0.088</u>
	1	28.184	<u>0.885</u>	<u>0.086</u>	1	29.911	<u>0.904</u>	0.101

mulation). After this pre-training phase, we enhanced the FLUX-IR model using our proposed strategy. Given the significant size and training costs associated with FLUX, we combined reinforcement learning and distillation into a single training phase, setting the diffusion timestamp to 9 and interpolating the initial state using the low-quality latent. We then applied reinforcement learning with guidance to further improve performance over a few inference steps. During this phase, the reward was calculated as the average of the following metrics: LPIPS, NIQE, MUSIQ, and CLIPQA, with normalization factors of -1 , -20 , 70 , and 1 , respectively. During training, we sampled the data from the task of super-resolution by a frequency of $5/10$ and others are the same of $5/60$.

3) *Methods under comparison & Evaluation metrics*: On the experiments of specific image restoration tasks, we compared the proposed method with the state-of-the-art methods in the fields of underwater enhancement, low-light enhancement, and deraining. For unified image restoration, we mainly compared the proposed method with the unified method for fairness. Moreover, according to the ill-posed nature of the image restoration inverse problems and to preserve the strong image synthesis capacity of the pre-trained FLUX model, we did not enforce the proposed method to fully approach the reference image and utilized more unpaired and perceptual scores, e.g., NIQE [112], MUSIQ [113], and CLIPQA [114], to validate the performance of our FLUX-IR.

B. Task-Specific Image Restoration Diffusion Models

Underwater Image Enhancement. The quantitative comparisons are presented in Table I, demonstrating that the proposed method significantly outperforms state-of-the-art techniques, such as HCLR-net [84] and SemiUIR [85], on the UIEBD dataset by **1.3** dB. The single step model even beats the multi-

step counterparts. We visualized the enhanced results in Fig. 4. For the U45 dataset, due to the fact that there is no ground truth available, we only provided the low-quality input with the corresponding reconstruction. Our method reconstructs more clear details with visually pleasing color, especially for the 1^{st} and 3^{rd} rows on the U45 dataset. Moreover, on the UIEB dataset, our method may even generate more visually pleasing results than the reference image, shown as the first example. The enhanced image has more soft light, making it easier to distinguish the foreground object, e.g., the shark, and the background scene, e.g., coral.

Low-light Image Enhancement. Table II presents the performance of various methods in low-light enhancement. It is important to note that *LOL-V1* serves as a highly competitive benchmark. Nonetheless, the proposed method demonstrates a notable improvement of **0.7** dB. Furthermore, it achieves an enhancement exceeding **1.1** dB on *LOL-V2*. Fig. 5 visually compares the results of different methods. The proposed method generates more clear details than other compared methods, e.g., the leaves in the 1^{st} and 3^{rd} examples from the *LOL-V2* dataset, and the small word “ano” in 3_{rd} examples marked by the green rectangle. Moreover, even for the region with extremely low-light conditions, our method can also accurately reconstruct it, shown as the local zoom in the region with green rectangle in the first example of the *LOL* dataset.

Deraining. Substantial improvements of **2.1** dB and **0.9** dB are evidenced in Tables III and IV, respectively. Notably, the proposed method, employing both dual-step and single-step approaches, demonstrates superior performance compared to specialized deraining diffusion models. This underscores the necessity and effectiveness of our strategies for ODE trajectory augmentation and simplification. This highlights the necessity and effectiveness of our strategies for ODE trajectory augmentation and simplification. Fig. 6 provides visual results



Fig. 6. Visual comparison on the tasks of raindrop removal and image deraining. Raindrop removal (**top**): (a) low-quality input, (b) reference samples, (c) IDT [115], (d) GridFormer [116], (e) RainDropDiff [65], (f) Ours ($k = 1$), (g) Ours ($k = 10$). Deraining (**bottom**): (a) low-quality input, (b) reference samples, (c) GridFormer [116], (d) WeatherDiff64 [65], (e) WeatherDiff128 [65], (f) Ours ($k = 1$), (g) Ours ($k = 10$). k is the number of inference steps, commonly referred to as the number of function evaluations (NFE).

TABLE III

QUANTITATIVE COMPARISONS OF DIFFERENT METHODS ON **IMAGE DERAINING**. THE BEST AND SECOND-BEST RESULTS ARE HIGHLIGHTED IN **BOLD** AND UNDERLINED, RESPECTIVELY. “ \uparrow ” (RESP. “ \downarrow ”) MEANS THE LARGER (RESP. SMALLER), THE BETTER. “NFE” DENOTES THE NUMBER OF FUNCTION EVALUATIONS, WHICH CAN BE INTERPRETED AS THE INFERENCE STEPS.

Method	NFE	Outdoor-Rain		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
pix2pix [117] CVPR’17	1	19.09	0.7100	-
HRGAN [118] CVPR’19	1	21.56	0.8550	0.154
PCNet [119] TIP’21	1	26.19	0.9015	0.132
MPRNet [120] CVPR’21	1	28.03	0.9192	0.089
Restormer [10] CVPR’22	1	29.97	0.9215	0.074
WeatherDiff ₆₄ [65] TPAMI’23	25	29.41	0.9312	0.059
WeatherDiff ₁₂₈ [65] TPAMI’23	25	29.28	0.9216	0.061
DTPM [121] CVPR’24	50	30.48	0.9210	0.054
DTPM [121] CVPR’24	10	30.92	0.9320	0.062
DTPM [121] CVPR’24	4	30.99	0.9340	0.064
Proposed Method	10	32.08	<u>0.9424</u>	0.065
	2	33.10	0.9439	<u>0.058</u>
	1	<u>32.61</u>	0.9419	0.064

TABLE IV

QUANTITATIVE COMPARISONS OF DIFFERENT METHODS ON **RAINDROP REMOVING**. THE BEST AND SECOND-BEST RESULTS ARE HIGHLIGHTED IN **BOLD** AND UNDERLINED, RESPECTIVELY. “ \uparrow ” (RESP. “ \downarrow ”) MEANS THE LARGER (RESP. SMALLER), THE BETTER. “NFE” DENOTES THE NUMBER OF FUNCTION EVALUATIONS, WHICH CAN BE INTERPRETED AS THE INFERENCE STEPS.

Method	NFE	Raindrop		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DuRN [122] CVPR’19	1	31.24	0.9259	-
RaindropAttn [123] ICCV’19	1	31.44	0.9263	0.068
AttentiveGAN [104] CVPR’18	1	31.59	0.9170	0.055
IDT [115] TPAMI’22	1	31.87	0.9313	0.059
RainDropDiff ₆₄ [65] TPAMI’23	25	32.29	0.9422	0.058
RainDropDiff ₁₂₈ [65] TPAMI’23	25	32.43	0.9334	0.058
AST-B [124] CVPR’24	1	32.38	0.9350	0.066
DTPM [121] CVPR’24	50	31.44	0.9320	0.044
DTPM [121] CVPR’24	10	31.87	0.9370	<u>0.048</u>
DTPM [121] CVPR’24	4	32.72	0.9440	0.058
Proposed Method	10	33.32	0.9388	0.044
	2	<u>33.53</u>	<u>0.9444</u>	0.050
	1	33.63	0.9459	0.052

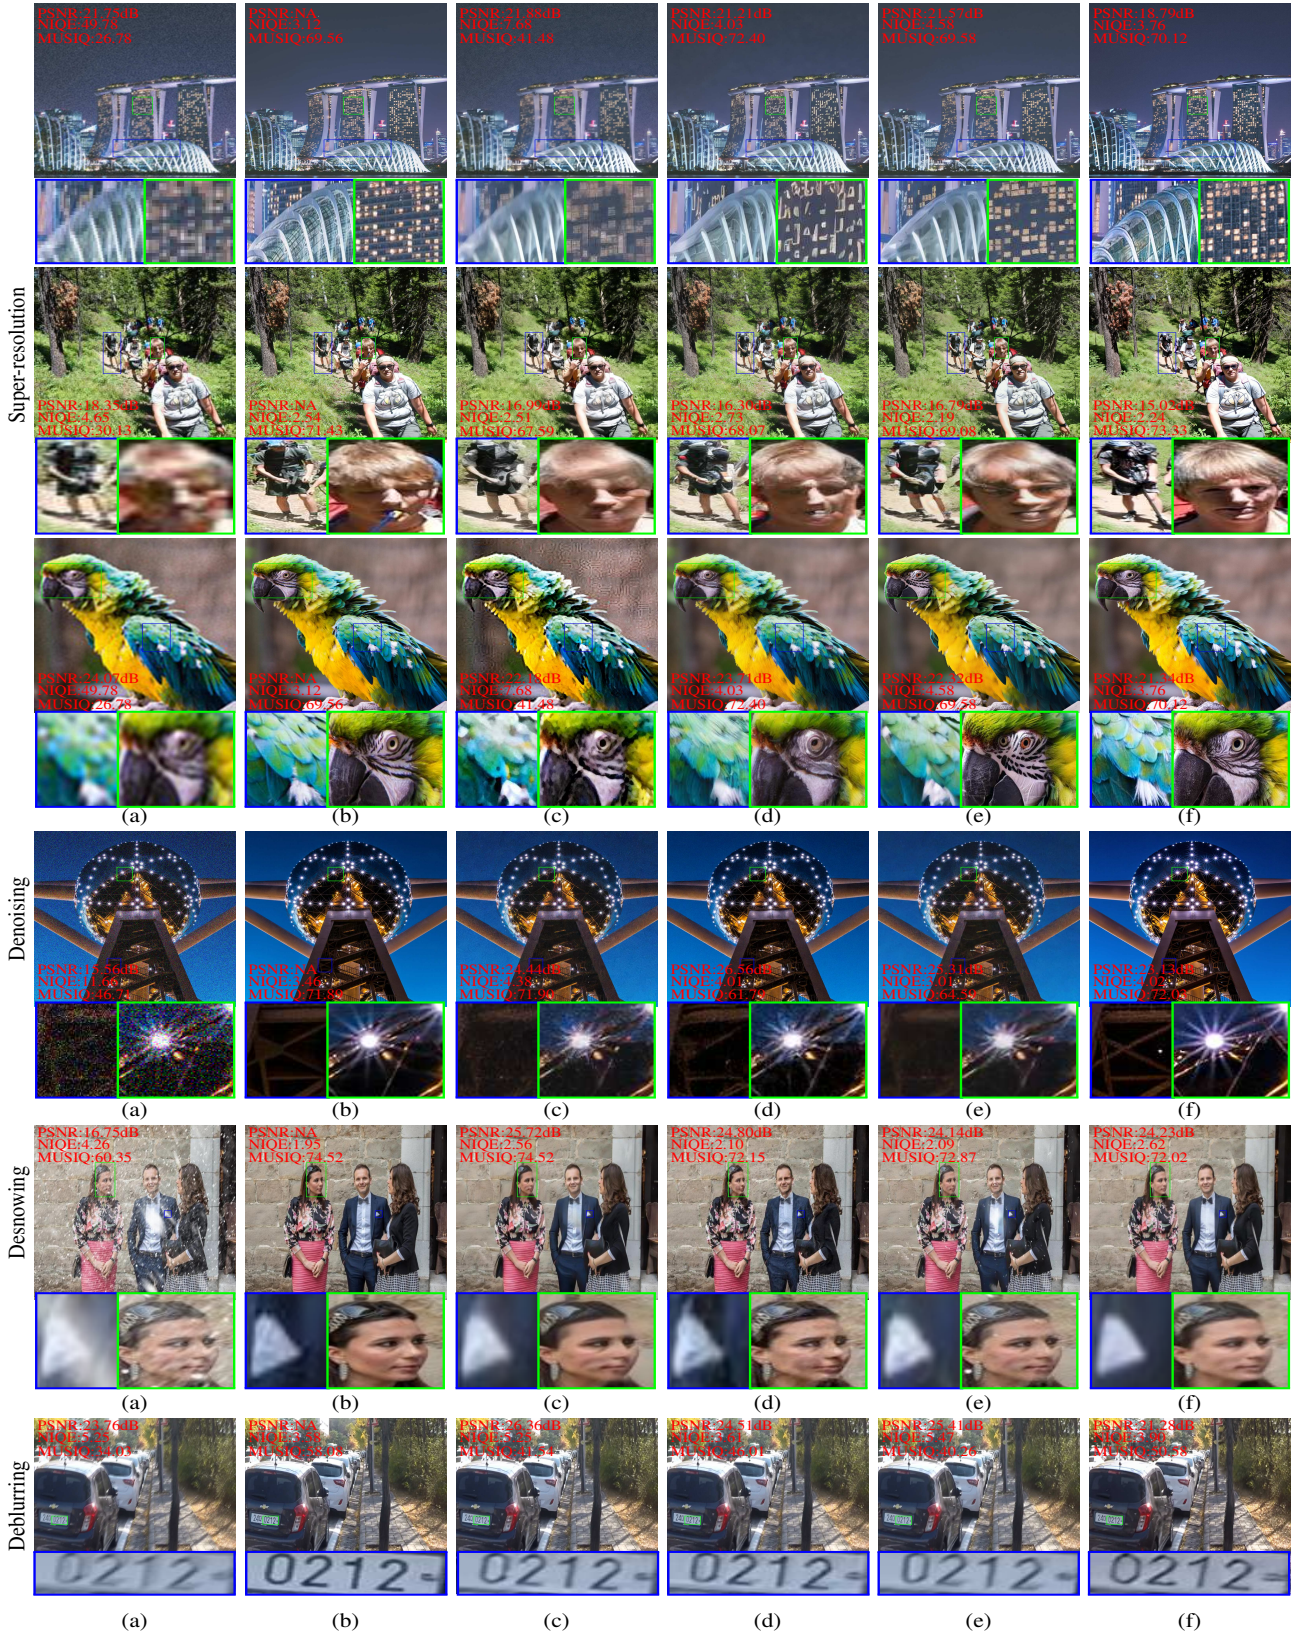


Fig. 7. Visual comparison of unified image restoration, where four different tasks are compared in the figure. In the super-resolution task, (a)-(f) indicates the input low-quality measurement, reference image, PASD [125], SeeSR [126], and FLUX-IR (Ours), respectively. For denoising, (c)-(e) represents AdaIR [127], DA-CLIP [68], and PromptIR [128], respectively. Moreover, for desnowing, (c)-(e) shows WGWS-Net [129], DA-CLIP [68], and DiffUIR [130], respectively. Finally for deblurring, (c)-(e) shows AdaIR [127], DA-CLIP [68], and DiffUIR [130], respectively. We annotated evaluation metrics of corresponding images by PSNR \uparrow , NIQE \downarrow , and MUSIQ \uparrow , respectively.

TABLE V

QUANTITATIVE COMPARISONS OF DIFFERENT METHODS ON UNIFIED IR TASKS. “↑” (RESP. “↓”) MEANS THE LARGER (RESP. SMALLER), THE BETTER. † MEANS THE COMPARED METHODS WERE DESIGNED AS UNIFIED IMAGE RESTORATION FRAMEWORKS.

Super-resolution	DIV2K-Val [108]					Desnowing	Snow100K-L [107]				
	NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑		NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑
SeeSR [126] CVPR’24	50	4.073	72.437	0.722	23.329	WGWS-Net [129] CVPR’23	1	3.345	70.167	0.507	28.933
PASD [125] ECCV’24	20	3.676	71.738	0.695	23.143	GridFormer [116] IJCV’24	1	2.918	70.854	0.488	30.792
DiffBIR† [131] ArXiv’23	50	4.333	64.672	0.651	22.548	DiffUIR† [130] CVPR’24	3	3.163	70.265	0.456	28.879
AutoDIR† [132] ECCV’24	50	4.865	55.196	0.458	23.939	DA-CLIP† [68] ICLR’24	100	2.775	69.394	0.375	28.641
FLUX-IR	20	3.491	73.188	0.721	20.248	FLUX-IR	20	2.769	65.985	0.440	23.442
	10	4.269	72.069	0.673	20.895		10	2.862	68.117	0.482	24.840
De-blurring	GoPro [110]					Denoising	DIV2K-Val [108]				
	NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑		NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑
DA-CLIP† [68] ICLR’24	100	3.937	39.925	0.212	28.619	PromptIR† [128] NeurIPS’23	1	2.856	64.025	0.571	26.927
DiffUIR† [130] CVPR’24	3	5.834	34.061	0.203	27.815	DA-CLIP† [68] ICLR’24	100	5.091	58.098	0.584	26.979
AdaIR† [127] ArXiv’24	1	5.514	33.263	0.198	28.464	AdaIR† [127] ArXiv’24	1	4.458	59.714	0.611	25.953
AutoDIR† [132] ECCV’24	50	6.164	33.354	0.197	28.444	AutoDIR† [132] ECCV’24	50	5.095	58.399	0.515	28.081
FLUX-IR	20	4.578	46.728	0.233	23.884	FLUX-IR	20	3.631	72.657	0.685	24.531
	10	4.521	42.045	0.199	24.868		10	4.199	71.519	0.631	25.373
Low-Light Enhancement	LOL [95]					Raindrop removal	Raindrop [104]				
	NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑		NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑
DA-CLIP† [68] ICLR’24	100	5.208	70.500	0.633	26.768	WeatherDiff [65] TPAMI’23	25	3.517	71.631	0.471	30.713
DiffUIR† CVPR’24	3	4.904	71.389	0.378	25.269	WGWS-Net [129] CVPR’23	1	3.479	71.731	0.435	33.430
AdaIR† [127] ArXiv’24	1	4.713	70.859	0.404	22.409	AST-B [124] CVPR’24	1	3.272	69.750	0.427	32.380
AutoDIR† [132] ECCV’24	50	4.200	71.095	0.398	22.896	DA-CLIP† [68] ICLR’24	100	4.817	67.592	0.488	31.207
FLUX-IR	20	4.045	73.233	0.484	25.029	FLUX-IR	20	3.201	70.105	0.507	25.846
	10	4.163	72.728	0.460	24.914		10	3.320	69.661	0.515	27.918
Deraining&Dehazing	Outdoor-Rain [105]					Underwater	UIEBD [81]				
	NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑		NFE	NIQE↓	MUSIQ↑	CLIPQA↑	PSNR↑
WGWS-Net [129] CVPR’23	1	3.968	70.835	0.459	30.609	NU ² Net [133] AAAI’23	1	4.717	47.810	0.541	25.221
GridFormer [116] IJCV’24	1	3.620	70.299	0.511	31.874	HCLR-net [84] IJCV’24	1	4.803	48.623	0.593	24.998
Ours-specific	10	3.874	70.738	0.490	32.082	Ours-specific	10	4.902	49.157	0.578	25.081
Ours-specific	1	3.882	70.848	0.493	32.612	Ours-specific	1	4.809	48.851	0.590	26.253
FLUX-IR	20	2.804	70.631	0.514	24.734	FLUX-IR	20	4.515	49.867	0.574	23.275
	10	3.023	69.862	0.521	25.856		10	4.596	49.958	0.574	23.409

of different methods. Specifically, for the task of raindrop removal, we visualized the regions seriously deteriorated by raindrops, e.g., the roof of a car and shelves on the playground, in zoom-in sub-figures, where the input measurement, i.e., Fig. 6-(a), indicates the object structure has been greatly damaged. However, even with this kind of degradation, the proposed method can accurately reconstruct the original structure, validating the superiority of the proposed method. The first sample of the Raindrop dataset also validates that the proposed method can correct the color of texture since the other methods show more red components compared with the proposed method. Meanwhile, we also visualize the deraining experimental results of *Outdoor-Rain* dataset in Fig. 6. Note that our method is not specifically trained to reconstruct words. However, to our surprise, the strong restoration capacity of the proposed method enables accurate reconstruction of words and numbers, e.g., “2.8” in the first sample and “U” in the second sample of the *Outdoor-Rain* dataset.

Based on the aforementioned analysis, we conclude that the proposed method can accurately reconstruct both natural texture and cultural markers, and greatly outperform the state-of-the-art method by a large extent, validating the effectiveness of the proposed image restoration diffusion augmentation strategy.

C. Unified Perceptual Image Restoration with FLUX-IR

We further validated the effectiveness of the proposed trajectory augmentation techniques in the unified image restoration task, which contains 7 distinct image restoration tasks. Experimental results are shown in Table V. Benefiting from strong capacity and our delicately designed learning scheme, FLUX-IR achieves extraordinary performance on the perceptual metrics, which even beats the task-specific methods, e.g., Stable SR [30] and PASD [125] from the task of super-resolution. We want to note that accurately quantifying the perceptual performance is a difficult issue. The outstanding performance of the proposed method, which may generate more reasonable and clear structures even than reference images, may make evaluation more difficult.

We visually compare various methods in Fig. 7 across four tasks: super-resolution, denoising, deblurring, and desnowing, which were not previously illustrated. Corresponding metrics are annotated in the corners of the images. While reference-aware metrics such as PSNR are critical for traditional image restoration tasks, relying solely on these metrics is inadequate for assessing quality in real-world scenarios. For instance, in the super-resolution task, manual interpolation, as depicted in Fig. 7-SR-(a), achieves superior PSNR scores compared to all other methods, yet its visual quality is inferior. Furthermore, evaluating quality without reference poses significant chal-

TABLE VI

RESULTS OF THE ABLATIVE STUDY ON THE EFFECT OF THE PROPOSED TWO TRAINING AUGMENTATION TECHNIQUES. “RL” DENOTES THE REINFORCEMENT LEARNING-BASED ALIGNMENT. “DISTILL” REPRESENTS THE FEW-STEP DISTILLATION FOR INFERENCE ACCELERATION. “INTERP” INDICATES THE INTERPOLATION OF THE STARTING POINT. “NGS” REPRESENTS UTILIZING THE LOW-QUALITY IMAGE AS NEGATIVE GUIDANCE.

IDX.	RL	DISTILL	INTERP	NGS	NFE	UIEBD			Raindrop			LOL-v2		
						PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
1	✗	✗	✗	✗	10	24.31	0.916	0.151	32.86	0.942	0.059	28.78	0.895	<u>0.094</u>
					1	19.99	0.802	0.284	20.92	0.357	0.674	20.35	0.717	0.255
2	✓	✗	✗	✗	10	25.08	0.913	0.162	33.32	0.938	0.044	29.54	<u>0.898</u>	0.086
					1	21.06	0.715	0.457	24.54	0.807	0.266	21.35	0.690	0.263
3	✓	✓	✗	✗	10	24.04	0.856	0.225	30.26	0.920	0.057	28.19	0.863	0.137
					1	25.80	0.923	0.153	33.35	0.942	<u>0.049</u>	29.75	0.904	0.096
4	✓	✓	✓	✗	10	24.20	0.923	0.136	29.80	0.918	0.056	28.30	0.864	0.135
					1	<u>25.94</u>	<u>0.937</u>	0.127	<u>33.53</u>	0.946	0.052	<u>29.86</u>	0.904	0.102
5	✓	✓	✓	✓	10	21.87	0.845	0.261	29.25	0.877	0.129	28.33	0.863	0.130
					1	26.25	0.938	<u>0.128</u>	33.63	0.946	0.052	29.91	0.904	0.101

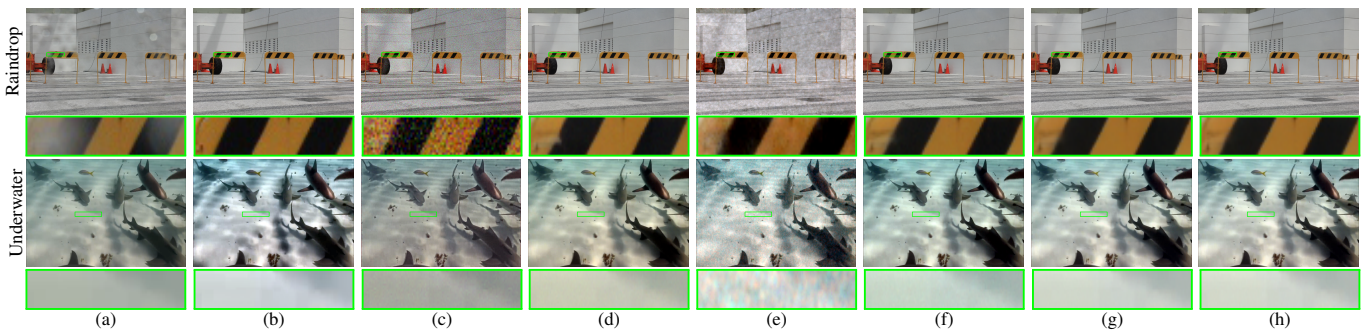


Fig. 8. Visual demonstration of ablation studies. (a) low-quality image, (b) reference image, (c) pretrained model ($k = 1$, Table VI-1), (d) pretrained model ($k = 10$, Table VI-1), (e) RL ($k = 1$, Table VI-2), (f) RL w/ DISTILL ($k = 1$, Table VI-3), (g) (f) w/ latent INTER ($k = 1$, Table VI-4), (h) (g) w/ NGS ($k = 1$, Table VI-5). k indicates the number of inference steps.




Fig. 9. Trajectories of different models. (a) Visualization of ODE and M-SDE trajectories for the pre-trained model. It can be seen that some M-SDE trajectories are more effective, being closer to the ground truth (GT) than the ODE counterpart. The optimal M-SDE trajectory is then selected as guidance for our reinforcement learning-based alignment process. (b) Visualization of ODE trajectories for the pre-trained model, our reinforced model, and our distilled model. “ t ” indicates the diffusion timestamp. Both reinforced and distilled models generate more effective ODE trajectories compared to the pre-trained model.

lenges, and a single perceptual metric may exhibit instability under certain conditions. Therefore, our approach of integrating multiple perceptual metrics for reinforcement learning is both practical and effective.

The following is a detailed analysis of the visual comparison results. Specifically, for the task of *super-resolution*, FLUX-IR significantly outperforms the compared methods, particularly evident in the depiction of building details and the feathers of the parrot in the first and third rows, respectively. These results not only match but sometimes surpass the perceptual quality of the reference images. The strong capabilities of FLUX-

IR are especially highlighted in human-related restoration, an area highly sensitive to our perception. As illustrated in the second row, FLUX-IR effectively generates realistic human faces and bodies, while the other methods struggle to produce plausible outcomes. Similar trends are observed in the tasks of denoising and desnowing. Lastly, for deblurring, this task differs from previous challenges involving noise, snowflakes, and downsampling, as the deblurring process merely mixes without compromising the original images. Consequently, all methods produce plausible images, as shown in the last row of Fig. 7.

D. Ablation Studies

Reconstruction Performance. We have applied the detailed ablation studies of each trajectory augmentation technique in Table VI. Specifically, for reinforcement learning-based ODE alignment (RL), it improves both single and multiple steps image restoration. Notably, models utilizing single-step alignment exhibit substantial improvements, such as an increase from 20.92 dB to 24.54 dB on the Raindrop dataset. This enhancement is attributed to the significant simplification of the trajectory post-alignment, which reduces integral estimation errors and benefits single-step inference. Conversely, few-step distillation (DISTILL) improves single-step performance but diminishes network performance with multiple steps. Furthermore, interpolation (INTERP) and guided sampling techniques (NGS) generally enhance single-step diffusion models by

alleviating the burden on neural networks to directly match Gaussian noise distributions with real-world data. However, for multi-step diffusion models, the use of interpolated noisy latent and negative guidance from low-quality images may present limitations. Then, the gradual performance improvement is also observed within visual demonstration of Fig 8.

Visualization of Trajectories. As our design is primarily centered on modeling trajectories of neural differential equations, we also analyze the variations in these trajectories to gain insights into the effects of our proposed method. As shown in Fig. 9-(a), we illustrate the ODE and M-SDE trajectories of the pre-trained model. The M-SDE trajectories exhibit more variance and potential compared to the ODE counterpart, which enables our reinforced ODE trajectory learning. Fig. 9-(b) depicts the ODE trajectories of the pre-trained, reinforced, and distilled models. Here, the reinforced ODE trajectory demonstrates a more direct path toward the target distribution, indicating that our reinforcement learning approach effectively optimizes the restoration path. Moreover, the distilled ODE trajectory closely approximates the reinforced trajectory, showcasing the efficacy of our trajectory distillation in preserving the optimized path while reducing computational complexity.

VI. CONCLUSION

We have presented an efficient yet effective trajectory optimization paradigm for image restoration-based diffusion models. Through reinforcement learning-based trajectory augmentation techniques, we boost the effectiveness of image restoration diffusion network. By employing different reward functions, we can flexibly guide the learning of the diffusion model toward either more objective or perceptual restoration. Moreover, based on distillation cost analysis, we introduced a diffusion acceleration distillation pipeline with several techniques to perverse the original knowledge of diffusion models and achieve single-step distillation. We have carried out extensive experiments on both task-specific image restoration diffusion and unified image restoration diffusion networks over more than 7 different image restoration tasks to validate the effectiveness of the proposed method. Moreover, we have also calibrated a *12B* rectified flow-based model for the image restoration task. Experimental results demonstrate the effectiveness of the proposed method, which generates clear and meaningful results compared with state-of-the-art methods. We believe our insights and findings would push the frontier of the field of image restoration.

REFERENCES

- [1] J. Noh, W. Bae, W. Lee, J. Seo, and G. Kim, "Better to follow, follow to be better: Towards precise supervision of feature super-resolution for small object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9725–9734.
- [2] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Proceedings of the European Conference on Computer Vision*, 2014, pp. 184–199.
- [3] M. Lu, T. Chen, H. Liu, and Z. Ma, "Learned image restoration for vvc intra coding," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, vol. 4, 2019.
- [4] M. V. Conde, U.-J. Choi, M. Burchi, and R. Timofte, "Swin2sr: Swinv2 transformer for compressed image super-resolution and restoration," in *European Conference on Computer Vision*, 2022, pp. 669–687.

- [5] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Transactions on Image Processing*, vol. 19, no. 11, pp. 2861–2873, 2010.
- [6] J. Yang, J. Wright, T. Huang, and Y. Ma, "Image super-resolution as sparse representation of raw image patches," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
- [7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [8] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 295–307, 2015.
- [9] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, "Residual dense network for image super-resolution," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2472–2481.
- [10] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, "Restormer: Efficient transformer for high-resolution image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5728–5739.
- [11] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, "Swinir: Image restoration using swin transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1833–1844.
- [12] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4681–4690.
- [13] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. Change Loy, "Esrgan: Enhanced super-resolution generative adversarial networks," in *Proceedings of the European Conference on Computer Vision Workshops*, 2018.
- [14] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.
- [15] K. Zhang, L. V. Gool, and R. Timofte, "Deep unfolding network for image super-resolution," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3217–3226.
- [16] I. Marivani, E. Tsiligianis, B. Cornelis, and N. Deligiannis, "Multimodal deep unfolding for guided image super-resolution," *IEEE Transactions on Image Processing*, vol. 29, pp. 8443–8456, 2020.
- [17] A. Lugmayr, M. Danelljan, L. Van Gool, and R. Timofte, "Srrflow: Learning the super-resolution space with normalizing flow," in *Proceedings of the European Conference on Computer Vision*, 2020, pp. 715–732.
- [18] Y. Kim and D. Son, "Noise conditional flow model for learning the super-resolution space," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 424–432.
- [19] J. Liang, A. Lugmayr, K. Zhang, M. Danelljan, L. Van Gool, and R. Timofte, "Hierarchical conditional flow: A unified framework for image super-resolution and image rescaling," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4076–4085.
- [20] B. Xia, Y. Zhang, S. Wang, Y. Wang, X. Wu, Y. Tian, W. Yang, and L. Van Gool, "Diffir: Efficient diffusion model for image restoration," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 13 095–13 105.
- [21] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 4, pp. 4713–4726, 2022.
- [22] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [23] B. D. Anderson, "Reverse-time diffusion equation models," *Stochastic Processes and their Applications*, vol. 12, no. 3, pp. 313–326, 1982.
- [24] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *Proceedings of the International Conference on Learning Representations*, 2021.
- [25] B. Kavar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 23 593–23 606, 2022.
- [26] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," in

- Proceedings of the International Conference on Learning Representations*, 2023.
- [27] Y. Wang, J. Yu, and J. Zhang, “Zero-shot image restoration using denoising diffusion null-space model,” in *Proceedings of the International Conference on Learning Representations*, 2023.
- [28] B. Fei, Z. Lyu, L. Pan, J. Zhang, W. Yang, T. Luo, B. Zhang, and B. Dai, “Generative diffusion prior for unified image restoration and enhancement,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9935–9946.
- [29] Z. Luo, F. K. Gustafsson, Z. Zhao, J. Sjölund, and T. B. Schön, “Image restoration with mean-reverting stochastic differential equations,” *Proceedings of the International Conference on Machine Learning*, 2023.
- [30] J. Wang, Z. Yue, S. Zhou, K. C. Chan, and C. C. Loy, “Exploiting diffusion prior for real-world image super-resolution,” *International Journal of Computer Vision*, 2024.
- [31] H. Chung, J. Kim, and J. C. Ye, “Direct diffusion bridge using data consistency for inverse problems,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [32] X. Liu, C. Gong, and qiang liu, “Flow straight and fast: Learning to generate and transfer data with rectified flow,” in *Proceedings of the International Conference on Learning Representations*, 2023.
- [33] Y. Song, P. Dhariwal, M. Chen, and I. Sutskever, “Consistency models,” in *Proceedings of the International Conference on Machine Learning*, vol. 202, 2023, pp. 32 211–32 252.
- [34] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *Proceedings of the International conference on machine learning*, 2015, pp. 2256–2265.
- [35] Y. Song and S. Ermon, “Generative modeling by estimating gradients of the data distribution,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [36] Y. Song and S. Ermon, “Improved techniques for training score-based generative models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 438–12 448, 2020.
- [37] Y. Song, C. Durkan, I. Murray, and S. Ermon, “Maximum likelihood training of score-based diffusion models,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 1415–1428, 2021.
- [38] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, “Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 5775–5787, 2022.
- [39] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, “Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models,” *arXiv preprint arXiv:2211.01095*, 2022.
- [40] K. Zheng, C. Lu, J. Chen, and J. Zhu, “Dpm-solver-v3: Improved diffusion ode solver with empirical model statistics,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [41] Y. Xu, Z. Liu, M. Tegmark, and T. Jaakkola, “Poisson flow generative models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 16 782–16 795, 2022.
- [42] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 684–10 695.
- [43] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts *et al.*, “Stable video diffusion: Scaling latent video diffusion models to large datasets,” *arXiv preprint arXiv:2311.15127*, 2023.
- [44] A. Blattmann, R. Rombach, H. Ling, T. Dockhorn, S. W. Kim, S. Fidler, and K. Kreis, “Align your latents: High-resolution video synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 22 563–22 575.
- [45] T. Karras, M. Aittala, T. Aila, and S. Laine, “Elucidating the design space of diffusion-based generative models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 26 565–26 577, 2022.
- [46] T. Dockhorn, A. Vahdat, and K. Kreis, “Score-based generative modeling with critically-damped langevin diffusion,” in *Proceedings of the International Conference on Learning Representations*, 2022.
- [47] S. Chen, S. Chewi, J. Li, Y. Li, A. Salim, and A. Zhang, “Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions,” in *Proceedings of the International Conference on Learning Representations*, 2023.
- [48] X. Liu, L. Wu, M. Ye, and Q. Liu, “Let us build bridges: Understanding and extending diffusion generative models,” *arXiv preprint arXiv:2208.14699*, 2022.
- [49] B. Li, K. Xue, B. Liu, and Y.-K. Lai, “Bbdm: Image-to-image translation with brownian bridge diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1952–1961.
- [50] L. Zhou, A. Lou, S. Khanna, and S. Ermon, “Denoising diffusion bridge models,” in *Proceedings of the International Conference on Learning Representations*.
- [51] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” in *Proceedings of the International Conference on Learning Representations*.
- [52] Q. Zhang and Y. Chen, “Fast sampling of diffusion models with exponential integrator,” *arXiv preprint arXiv:2204.13902*, 2022.
- [53] Z. Zhou, D. Chen, C. Wang, and C. Chen, “Fast ode-based sampling for diffusion models in around 5 steps,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 7777–7786.
- [54] S. Xue, M. Yi, W. Luo, S. Zhang, J. Sun, Z. Li, and Z.-M. Ma, “Sasolver: Stochastic adams solver for fast sampling of diffusion models,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [55] S. Xue, Z. Liu, F. Chen, S. Zhang, T. Hu, E. Xie, and Z. Li, “Accelerating diffusion sampling with optimized time steps,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 8292–8301.
- [56] T. Dockhorn, A. Vahdat, and K. Kreis, “Genie: Higher-order denoising diffusion solvers,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 30 150–30 166, 2022.
- [57] Y. Song and P. Dhariwal, “Improved techniques for training consistency models,” in *Proceedings of the International Conference on Learning Representations*, 2023.
- [58] D. Kim, C.-H. Lai, W.-H. Liao, N. Murata, Y. Takida, T. Uesaka, Y. He, Y. Mitsufuji, and S. Ermon, “Consistency trajectory models: Learning probability flow ode trajectory of diffusion,” in *Proceedings of the International Conference on Learning Representations*.
- [59] M. Zhou, H. Zheng, Z. Wang, M. Yin, and H. Huang, “Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation,” in *Proceedings of the International Conference on Machine Learning*, 2024.
- [60] M. Zhou, Z. Wang, H. Zheng, and H. Huang, “Long and short guidance in score identity distillation for one-step text-to-image generation,” *arXiv 2406.01561*, 2024.
- [61] Y. Ren, X. Xia, Y. Lu, J. Zhang, J. Wu, P. Xie, X. Wang, and X. Xiao, “Hyper-sd: Trajectory segmented consistency model for efficient image synthesis,” *arXiv preprint arXiv:2404.13686*, 2024.
- [62] F.-Y. Wang, Z. Huang, A. W. Bergman, D. Shen, P. Gao, M. Lingelbach, K. Sun, W. Bian, G. Song, Y. Liu *et al.*, “Phased consistency model,” *arXiv preprint arXiv:2405.18407*, 2024.
- [63] Q. Xie, Z. Liao, Z. Deng, S. Tang, H. Lu *et al.*, “Mlcm: Multi-step consistency distillation of latent diffusion model,” *arXiv preprint arXiv:2406.05768*, 2024.
- [64] Y. Zhu, K. Zhang, J. Liang, J. Cao, B. Wen, R. Timofte, and L. Van Gool, “Denoising diffusion models for plug-and-play image restoration,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1219–1229.
- [65] O. Özdenizci and R. Legenstein, “Restoring vision in adverse weather conditions with patch-based denoising diffusion models,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 8, pp. 10 346–10 357, 2023.
- [66] Y. Zhang, X. Shi, D. Li, X. Wang, J. Wang, and H. Li, “A unified conditional framework for diffusion-based image restoration,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [67] Z. Luo, F. K. Gustafsson, Z. Zhao, J. Sjölund, and T. B. Schön, “Refusion: Enabling large-size realistic image restoration with latent-space diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1680–1691.
- [68] Z. Luo, F. K. Gustafsson, Z. Zhao, J. Sjölund, and T. B. Schön, “Controlling vision-language models for universal image restoration,” *arXiv preprint arXiv:2310.01018*, 2023.
- [69] H. Jiang, A. Luo, H. Fan, S. Han, and S. Liu, “Low-light image enhancement with wavelet-based diffusion models,” *ACM Transactions on Graphics (TOG)*, vol. 42, no. 6, pp. 1–14, 2023.
- [70] X. Yi, H. Xu, H. Zhang, L. Tang, and J. Ma, “Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 12 302–12 311.
- [71] Y. Wang, W. Yang, X. Chen, Y. Wang, L. Guo, L.-P. Chau, Z. Liu, Y. Qiao, A. C. Kot, and B. Wen, “Sinsr: diffusion-based image super-resolution in a single step,” in *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, 2024, pp. 25 796–25 805.
- [72] Y. Tang, H. Kawasaki, and T. Iwaguchi, “Underwater image enhancement by transformer-based diffusion model with non-uniform sampling for skip strategy,” in *Proceedings of the ACM International Conference on Multimedia*, 2023, pp. 5419–5427.
- [73] C. Meng, R. Rombach, R. Gao, D. Kingma, S. Ermon, J. Ho, and T. Salimans, “On distillation of guided diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 14 297–14 306.
- [74] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [75] Y. Fan, O. Watkins, Y. Du, H. Liu, M. Ryu, C. Boutilier, P. Abbeel, M. Ghavamzadeh, K. Lee, and K. Lee, “Reinforcement learning for fine-tuning text-to-image diffusion models,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [76] B. Wallace, M. Dang, R. Rafailov, L. Zhou, A. Lou, S. Purushwalkam, S. Ermon, C. Xiong, S. Joty, and N. Naik, “Diffusion model alignment using direct preference optimization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 8228–8238.
- [77] Q. Zhang and Y. Chen, “Diffusion normalizing flow,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 16 280–16 291, 2021.
- [78] X. Liu, X. Zhang, J. Ma, J. Peng *et al.*, “Instaflow: One step is enough for high-quality diffusion-based text-to-image generation,” in *The Twelfth International Conference on Learning Representations*, 2023.
- [79] A. Sauer, D. Lorenz, A. Blattmann, and R. Rombach, “Adversarial diffusion distillation,” *arXiv preprint arXiv:2311.17042*, 2023.
- [80] C. Ancuti, C. O. Ancuti, T. Haber, and P. Bekaert, “Enhancing underwater images and videos by fusion,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 81–88.
- [81] C. Li, C. Guo, W. Ren, R. Cong, J. Hou, S. Kwong, and D. Tao, “An underwater image enhancement benchmark dataset and beyond,” *IEEE Transactions on Image Processing*, vol. 29, pp. 4376–4389, 2019.
- [82] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2017, pp. 2223–2232.
- [83] W. Zhang, P. Zhuang, H.-H. Sun, G. Li, S. Kwong, and C. Li, “Underwater image enhancement via minimal color loss and locally adaptive contrast enhancement,” *IEEE Transactions on Image Processing*, vol. 31, pp. 3997–4010, 2022.
- [84] J. Zhou, J. Sun, C. Li, Q. Jiang, M. Zhou, K.-M. Lam, W. Zhang, and X. Fu, “Hclr-net: Hybrid contrastive learning regularization with locally randomized perturbation for underwater image enhancement,” *International Journal of Computer Vision*, pp. 1–25, 2024.
- [85] S. Huang, K. Wang, H. Liu, J. Chen, and Y. Li, “Contrastive semi-supervised learning for underwater image restoration via reliable bank,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 18 145–18 155.
- [86] C. Li, S. Anwar, J. Hou, R. Cong, C. Guo, and W. Ren, “Underwater image enhancement via medium transmission-guided multi-color space embedding,” *IEEE Transactions on Image Processing*, vol. 30, pp. 4985–5000, 2021.
- [87] L. Chen, X. Chu, X. Zhang, and J. Sun, “Simple baselines for image restoration,” in *Proceedings of the European Conference on Computer Vision*, 2022, pp. 17–33.
- [88] Y. Feng, C. Zhang, P. Wang, P. Wu, Q. Yan, and Y. Zhang, “You only need one color space: An efficient network for low-light image enhancement,” *arXiv preprint arXiv:2402.05809*, 2024.
- [89] Y. Wang, R. Wan, W. Yang, H. Li, L.-P. Chau, and A. Kot, “Low-light image enhancement with normalizing flow,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 3, 2022, pp. 2604–2612.
- [90] Y. Cai, H. Bian, J. Lin, H. Wang, R. Timofte, and Y. Zhang, “Retinex-former: One-stage retinex-based transformer for low-light image enhancement,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 12 504–12 513.
- [91] T. Wang, K. Zhang, T. Shen, W. Luo, B. Stenger, and T. Lu, “Ultra-high-definition low-light image enhancement: A benchmark and transformer-based method,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 3, 2023, pp. 2654–2662.
- [92] X. Xu, R. Wang, C.-W. Fu, and J. Jia, “Snr-aware low-light image enhancement,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 714–17 724.
- [93] C. Guo, C. Li, J. Guo, C. C. Loy, J. Hou, S. Kwong, and R. Cong, “Zero-reference deep curve estimation for low-light image enhancement,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1780–1789.
- [94] Y. Jiang, X. Gong, D. Liu, Y. Cheng, C. Fang, X. Shen, J. Yang, P. Zhou, and Z. Wang, “Enlightengan: Deep light enhancement without paired supervision,” *IEEE Transactions on Image Processing*, vol. 30, pp. 2340–2349, 2021.
- [95] W. Y. J. L. Chen Wei, Wenjing Wang, “Deep retinex decomposition for low-light enhancement,” in *Proceedings of the British Machine Vision Conference*, 2018.
- [96] W. Yang, S. Wang, Y. Fang, Y. Wang, and J. Liu, “From fidelity to perceptual quality: A semi-supervised approach for low-light image enhancement,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3063–3072.
- [97] Y. Zhang, J. Zhang, and X. Guo, “Kindling the darkness: A practical low-light image enhancer,” in *Proceedings of the ACM International Conference on Multimedia*, 2019, pp. 1632–1640.
- [98] Y. Zhang, X. Guo, J. Ma, W. Liu, and J. Zhang, “Beyond brightening low-light images,” *International Journal of Computer Vision*, vol. 129, pp. 1013–1037, 2021.
- [99] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, “Learning enriched features for fast image restoration and enhancement,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 2, pp. 1934–1948, 2022.
- [100] D. Zhou, Z. Yang, and Y. Yang, “Pyramid diffusion models for low-light image enhancement,” in *Proceedings of the International Joint Conference on Artificial Intelligence*, 2023, pp. 1795–1803.
- [101] Y. Wu, G. Wang, S. Liu, Y. Yang, W. Li, X. Tang, S. Gu, C. Li, and H. T. Shen, “Towards a flexible semantic guided model for single image enhancement and restoration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [102] J. Hou, Z. Zhu, J. Hou, H. Liu, H. Zeng, and H. Yuan, “Global structure-aware diffusion process for low-light image enhancement,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [103] B. Andreas, S. Axel, L. Dominik, P. Dustin, B. Frederic, S. Harry, M. Jonas, L. Kyle, E. Patrick, R. Robin, K. Sumith, D. Tim, L. Yam, and E. Zion, “Flux,” 2024. [Online]. Available: <https://blackforestlabs.ai/>
- [104] R. Qian, R. T. Tan, W. Yang, J. Su, and J. Liu, “Attentive generative adversarial network for raindrop removal from a single image,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2482–2491.
- [105] K. Jiang, Z. Wang, P. Yi, C. Chen, B. Huang, Y. Luo, J. Ma, and J. Jiang, “Multi-scale progressive fusion network for single image deraining,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8346–8355.
- [106] W. Yang, W. Wang, H. Huang, S. Wang, and J. Liu, “Sparse gradient regularized deep retinex network for robust low-light image enhancement,” *IEEE Transactions on Image Processing*, vol. 30, pp. 2072–2086, 2021.
- [107] Y.-F. Liu, D.-W. Jaw, S.-C. Huang, and J.-N. Hwang, “Desnownet: Context-aware deep network for snow removal,” *IEEE Transactions on Image Processing*, vol. 27, no. 6, pp. 3064–3073, 2018.
- [108] E. Agustsson and R. Timofte, “Ntire 2017 challenge on single image super-resolution: Dataset and study,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2017.
- [109] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, “Enhanced deep residual networks for single image super-resolution,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition workshops*, 2017, pp. 136–144.
- [110] S. Nah, T. Hyun Kim, and K. Mu Lee, “Deep multi-scale convolutional neural network for dynamic scene deblurring,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3883–3891.
- [111] “Xflux,” 2024. [Online]. Available: <https://github.com/XLabs-AI/x-flux>
- [112] A. Mittal, R. Soundararajan, and A. C. Bovik, “Making a “completely blind” image quality analyzer,” *IEEE Signal processing letters*, vol. 20, no. 3, pp. 209–212, 2012.
- [113] J. Ke, Q. Wang, Y. Wang, P. Milanfar, and F. Yang, “Musiq: Multi-scale image quality transformer,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5148–5157.

- [114] J. Wang, K. C. Chan, and C. C. Loy, "Exploring clip for assessing the look and feel of images," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 2, 2023, pp. 2555–2563.
- [115] J. Xiao, X. Fu, A. Liu, F. Wu, and Z.-J. Zha, "Image de-raining transformer," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [116] T. Wang, K. Zhang, Z. Shao, W. Luo, B. Stenger, T. Lu, T.-K. Kim, W. Liu, and H. Li, "Gridformer: Residual dense transformer with grid structure for image restoration in adverse weather conditions," *International Journal of Computer Vision*, pp. 1–23, 2024.
- [117] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.
- [118] R. Li, L.-F. Cheong, and R. T. Tan, "Heavy rain image restoration: Integrating physics model and conditional adversarial learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1633–1642.
- [119] K. Jiang, Z. Wang, P. Yi, C. Chen, Z. Wang, X. Wang, J. Jiang, and C.-W. Lin, "Rain-free and residue hand-in-hand: A progressive coupled network for real-time image deraining," *IEEE Transactions on Image Processing*, vol. 30, pp. 7404–7418, 2021.
- [120] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, "Multi-stage progressive image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 821–14 831.
- [121] T. Ye, S. Chen, W. Chai, Z. Xing, J. Qin, G. Lin, and L. Zhu, "Learning diffusion texture priors for image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2524–2534.
- [122] X. Liu, M. Suganuma, Z. Sun, and T. Okatani, "Dual residual networks leveraging the potential of paired operations for image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7007–7016.
- [123] Y. Quan, S. Deng, Y. Chen, and H. Ji, "Deep learning for seeing through window with raindrops," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2463–2471.
- [124] S. Zhou, D. Chen, J. Pan, J. Shi, and J. Yang, "Adapt or perish: Adaptive sparse transformer with attentive feature refinement for image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2952–2963.
- [125] P. R. X. X. Tao Yang, Rongyuan Wu and L. Zhang, "Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization," in *Proceedings of the European Conference on Computer Vision*, 2023.
- [126] R. Wu, T. Yang, L. Sun, Z. Zhang, S. Li, and L. Zhang, "Seesr: Towards semantics-aware real-world image super-resolution," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 25 456–25 467.
- [127] Y. Cui, S. W. Zamir, S. Khan, A. Knoll, M. Shah, and F. S. Khan, "Adair: Adaptive all-in-one image restoration via frequency mining and modulation," *arXiv preprint arXiv:2403.14614*, 2024.
- [128] V. Potlapalli, S. W. Zamir, S. Khan, and F. Khan, "Promptir: Prompting for all-in-one image restoration," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [129] Y. Zhu, T. Wang, X. Fu, X. Yang, X. Guo, J. Dai, Y. Qiao, and X. Hu, "Learning weather-general and weather-specific features for image restoration under multiple adverse weather conditions," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 747–21 758.
- [130] D. Zheng, X.-M. Wu, S. Yang, J. Zhang, J.-F. Hu, and W.-S. Zheng, "Selective hourglass mapping for universal image restoration based on diffusion model," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 25 445–25 455.
- [131] X. Lin, J. He, Z. Chen, Z. Lyu, B. Dai, F. Yu, W. Ouyang, Y. Qiao, and C. Dong, "Diffbir: Towards blind image restoration with generative diffusion prior," *arXiv preprint arXiv:2308.15070*, 2023.
- [132] Y. Jiang, Z. Zhang, T. Xue, and J. Gu, "Autodir: Automatic all-in-one image restoration with latent diffusion," *arXiv preprint arXiv:2310.10123*, 2023.
- [133] C. Guo, R. Wu, X. Jin, L. Han, W. Zhang, Z. Chai, and C. Li, "Underwater ranker: Learn which is better and how to be better," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 1, 2023, pp. 702–709.
- [134] B. Øksendal and B. Øksendal, *Stochastic differential equations*. Springer, 2003.
- [135] D. Kingma, T. Salimans, B. Poole, and J. Ho, "Variational diffusion models," *Advances in Neural Information Processing Systems*, vol. 34, pp. 21 696–21 707, 2021.

Zhiyu Zhu received the B.E. and M.E. degrees in Mechatronic Engineering, both from Harbin Institute of Technology, in 2017 and 2019, respectively. He has also received a Ph.D. degree in Computer Science from the City University of Hong Kong in 2023, where he currently holds a postdoctoral position. His research interests include generative models and computer vision.

Jinhui Hou the B.E. and M.E. degrees in Communication Engineering from Huaqiao University, Xiamen, China, in 2017 and 2020, respectively. He is pursuing a Ph.D. in Computer Science at the City University of Hong Kong. His research interests include hyperspectral image processing and computer vision.

Hui Liu is currently an Assistant Professor in the School of Computing and Information Sciences at Saint Francis University, Hong Kong. She received the B.Sc. degree in Communication Engineering from Central South University, Changsha, China, the M.Eng. degree in Computer Science from Nanyang Technological University, Singapore, and the Ph.D. degree from the Department of Computer Science, City University of Hong Kong, Hong Kong. Her research interests include image processing and machine learning.

Huanqiang Zeng received the B.S. and M.S. degrees in electrical engineering from Huaqiao University, China, and the Ph.D. degree in electrical engineering from Nanyang Technological University, Singapore.

He is currently a Full Professor at the School of Engineering and the School of Information Science and Engineering, Huaqiao University. Before that, he was a Postdoctoral Fellow at The Chinese University of Hong Kong, Hong Kong. He has published more than 100 papers in well-known journals and conferences, including three best poster/paper awards (in the International Forum of Digital TV and Multimedia Communication 2018 and the Chinese Conference on Signal Processing 2017/2019). His research interests include image processing, video coding, machine learning, and computer vision. He has also been actively serving as the General Co-Chair for IEEE International Symposium on Intelligent Signal Processing and Communication Systems 2017 (ISPACS2017), the Co-Organizer for ICME2020 Workshop on 3D Point Cloud Processing, Analysis, Compression, and Communication, the Technical Program Co-Chair for Asia-Pacific Signal and Information Processing Association Annual Summit and Conference 2017 (APSIPA-ASC2017), the Area Chair for IEEE International Conference on Visual Communications and Image Processing (VCIP2015 and VCIP2020), and a technical program committee member for multiple flagship international conferences. He has been actively serving as an Associate Editor for IEEE Transactions on Image Processing, IEEE Transactions on Circuits and Systems for Video Technology, and Electronics Letters (IET). He has been actively serving as a Guest Editor for Journal of Visual Communication and Image Representation, Multimedia Tools and Applications, and Journal of Ambient Intelligence and Humanized Computing.

Junhui Hou (Senior Member, IEEE) is an Associate Professor with the Department of Computer Science, City University of Hong Kong. He holds a B.Eng. degree in information engineering (Talented Students Program) from the South China University of Technology, Guangzhou, China (2009), an M.Eng. degree in signal and information processing from Northwestern Polytechnical University, Xi'an, China (2012), and a Ph.D. degree from the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore (2016). His research interests are multi-dimensional visual computing.

Dr. Hou received the Early Career Award (3/381) from the Hong Kong Research Grants Council in 2018 and the NSFC Excellent Young Scientists Fund in 2024. He has served or is serving as an Associate Editor for *IEEE Transactions on Visualization and Computer Graphics*, *IEEE Transactions on Image Processing*, *IEEE Transactions on Multimedia*, and *IEEE Transactions on Circuits and Systems for Video Technology*.

APPENDIX A
PROBABILISTIC FLOW OF RECTIFIED FLOW

In this section, we analyze the probabilistic flow of FLUX [103] to prove: ① the given integral formulation of Eq. (13) has the stochastic formulation of Eq. (15); and ② given ODE Eq. (6) and SDE formulation (15) in our reinforcement alignment, belonging to the same probabilistic flow. Specifically, rectified flow is designed to directly learn the following velocity as

$$dx_t = (x_0 - x_T) dt. \quad (22)$$

Deterministic formulation of rectified flow is shown as:

$$x_{t-\Delta_t} = x_t - \Delta_t \frac{dx_t}{dt}. \quad (23)$$

Moreover, the stochastic formulation of rectified flow:

$$x_{t-\Delta_t} = \frac{[x_t - \alpha_t \Delta_t \frac{dx_t}{dt} - \beta_k \epsilon]}{(1 + \alpha_t \Delta_t - t) + \sqrt{(t - \alpha_t \Delta_t)^2 + \beta_k^2}}, \quad (24)$$

where α_t is a scalar ($\alpha_t > 1$), and β_k is formulated as

$$\beta_k = \sqrt{\frac{(t - \Delta_t)^2 [1 - (t - \alpha_t \Delta_t)]^2}{[1 - (t - \Delta_t)]^2} - (t - \alpha_t \Delta_t)^2}. \quad (25)$$

Then, we will illustrate that the aforementioned deterministic and stochastic formulations, i.e., Eq. (23) and Eq. (24), respectively, correspond to the same probabilistic flow. For the probabilistic flow of Eq. (23), we can formulate it by substituting Eq. (24) into the Kolmogorov's forward equation as

$$\begin{aligned} \frac{dp(x)}{dt} &= -\frac{d(f(x, t)p(x))}{dx} + \frac{1}{2} \frac{d^2(g^2(x, t)p(x))}{dx^2}, \\ &= -\frac{d((x_0 - x_T)p(x))}{dx}. \end{aligned} \quad (26)$$

To derive the probabilistic flow of stochastic equation, we first substitute Eq. (25) into the denominator expression of Eq. (24), obtaining

$$\begin{aligned} &(1 + \alpha_t \Delta_t - t) + \sqrt{(t - \alpha_t \Delta_t)^2 + \beta_k^2} \\ &= (1 - (t - \alpha_t \Delta_t)) + \frac{(t - \Delta_t) [1 - (t - \alpha_t \Delta_t)]}{[1 - (t - \Delta_t)]} \\ &= \frac{1 - (t - \alpha_t \Delta_t)}{1 - (t - \Delta_t)}. \end{aligned}$$

Moreover, we also have

$$\beta_k \stackrel{\Delta_t \rightarrow dt}{\approx} \sqrt{2(\alpha_t - 1)} \sqrt{dt}.$$

Then, we can substitute the aforementioned two results accompanied with $x_t = (1 - t)x_0 + tx_T$ together into Eq. (24). We have

$$\begin{aligned} x_{t-\Delta_t} &= \frac{[1 - (t - \Delta_t)]}{[1 - (t - \alpha_t \Delta_t)]} [x_t - \alpha_t \Delta_t \frac{dx_t}{dt} - \sqrt{2t\alpha_t} \epsilon], \\ &\frac{[1 - (t - \alpha_t \Delta_t)]}{[1 - (t - \Delta_t)]} (x_{t-\Delta_t} - x_t) \\ &= \frac{(1 - \alpha_t) dt}{1 - (t - \Delta_t)} x_t + \alpha_t (x_0 - x_T) dt + \sqrt{2t\alpha_t} d\omega, \\ -dx &= \frac{(1 - \alpha_t) dt}{1 - t} ((1 - t)x_0 + tx_T) \\ &\quad + \alpha_t (x_0 - x_T) dt + \sqrt{2t\alpha_t} d\omega, \end{aligned}$$

$$\textcircled{1} \quad dx = x_0 dt + \frac{t - \alpha_t}{1 - t} x_T dt + \sqrt{2(\alpha_t - 1)} d\omega. \quad (27)$$

The corresponding Kolmogorov's forward equation can be written as

$$\begin{aligned} \frac{dp(x)}{dt} &= -\frac{d(f(x, t)p(x))}{dx} + \frac{1}{2} \frac{d^2(g^2(x, t)p(x))}{dx^2}, \\ &= -\frac{d\left(\left(x_0 dt + \frac{t - \alpha_t}{1 - t} x_T dt\right) p(x)\right)}{dx} + (\alpha_t - 1) \frac{d \log p(x)}{dx} p(x). \end{aligned}$$

Since we have $\frac{d \log p(x)}{dx} = \frac{x_T}{1 - t}$, we can conclude the probabilistic flow of Eq. (24) is

$$\textcircled{2} \quad \frac{dp(x)}{dt} = -\frac{d((x_0 - x_T)p(x))}{dx},$$

which is same with Eq. (23). It indicates that the stochastic process of Eq. (27) and deterministic process Eq.(22) represent the identical probabilistic flow. Thus, the corresponding integral formulations of Eqs. (23) and (24) also represent the same probabilistic transition process.

APPENDIX B
MODULATED-SDE AND ITS INTEGRAL SOLVER

A. Modulated-SDE as a general expression of reverse diffusion derivative equation

The proof generally follows the diffusion ODE from [24]. Considering a general formulation of forward diffusion SDE as

$$d\mathbf{X} = f(\mathbf{X}, t)dt + g(t)d\omega,$$

the marginal probability $\mathbf{P}(\mathbf{X}_t)$ evolves with the following Kolmogorov's forward equation [134],

$$\begin{aligned} \frac{dp(x)}{dt} &= -\frac{d(f(x, t)p(x))}{dx} + \frac{1}{2} \frac{d^2(g^2(x, t)p(x))}{dx^2}, \\ &= -\frac{d(f(x, t)p(x))}{dx} + \left[\frac{1 + \gamma^2(t)}{2} - \frac{\gamma(t)^2}{2} \right] \frac{d^2(g^2(x, t)p(x))}{dx^2} \\ &= -\frac{d}{dx} \left[f(x, t)p(x) - \frac{1 + \gamma^2(t)}{2} \left(\frac{p(x)dg^2(x, t)}{dx} \right. \right. \\ &\quad \left. \left. + \frac{g^2(x, t)dp(x)}{dx} \right) \right] - \frac{1}{2} \frac{d^2 \left[[\gamma(t)g(x, t)]^2 p(x) \right]}{dx^2}. \end{aligned}$$

Since for the diffusion process, we $g(t)$ is independent with \mathbf{X} . Thus, we have

$$\frac{dp(x)}{dt} = -\frac{d}{dx} \left[\left(f(x, t) - \frac{1 + \gamma^2(t)}{2} \frac{g^2(x, t) d \log p(x)}{dx} \right) p(x) \right] - \frac{1}{2} \frac{d^2}{dx^2} \left[[\gamma(t)g(x, t)]^2 p(x) \right].$$

Considering for the reverse-time SDE with timestamp of \hat{t} , $d\hat{t} = -dt$. Then

$$\frac{dp(x)}{d\hat{t}} = -\frac{d}{dx} - \left[\left(f(x, t) - \frac{1 + \gamma^2(t)}{2} \frac{g^2(x, t) d \log p(x)}{dx} \right) p(x) \right] + \frac{1}{2} \frac{d^2}{dx^2} \left[[\gamma(t)g(x, t)]^2 p(x) \right].$$

It actually corresponds to the Kolmogorov's forward equation with the following differential equations.

$$\begin{aligned} d\mathbf{X} &= \hat{f}(\mathbf{X}, t)d\hat{t} + g(\hat{t})d\hat{\omega}, \\ \hat{f}(\mathbf{X}, t) &= - \left[f(x, t) - \frac{1 + \gamma^2(t)}{2} \frac{g^2(x, t) d \log p(x)}{dx} \right], \\ g(\hat{t}) &= \gamma(t)g(x, t). \end{aligned}$$

By substituting $d\hat{t} = -dt$ into above equation, we have

$$d\mathbf{X} = \left[f(x, t) - \frac{1 + \gamma^2(t)}{2} g^2(x, t) \nabla_x \log p(x) \right] dt + \gamma(t)g(x, t)d\hat{\omega}; \mathbf{X}_t^* = \mathbf{X}_t + \gamma\epsilon, (\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})).$$

which is exactly the Modulated-SDE as we mentioned. We want to note that the same SDE formulation has been introduced in [52], [77]. However, our proof follows [24] is more straightforward and complete.

B. Integral Solver of Modulated-SDE

Here we give the calculation of the integral solver for Modulated-SDE. To make such of semi-linear property as [38], we introduce the surrogate function $\mathcal{F}(\mathbf{X}_t, \alpha_t) = \frac{\mathbf{X}_t}{\alpha_t}$. Furthermore, by substituting $f(t) = \frac{d \log \alpha_t}{dt}$ and $g^2(t) = \frac{d\sigma_t^2}{dt} - 2\frac{d \log \alpha_t}{dt} \sigma_t^2$ from [135], we have

$$\begin{aligned} d\mathcal{F} &= \frac{d^{(\gamma)} \mathbf{X}_t}{\alpha_t} - \frac{\mathbf{X}_t d\alpha_t}{\alpha_t^2} \\ &= \frac{1 + \gamma^2(t)}{2\alpha_t \sigma_t} g^2(t) \epsilon_\theta dt + \gamma(t)g(t)d\hat{\omega} \end{aligned}$$

We take the first-order integral solver for an example. Through making integral from both sides

$$\begin{aligned} \mathbf{X}_{t-\Delta t} &= \frac{\alpha_{t-\Delta t}}{\alpha_t} \mathbf{X}_t - [1 + \gamma^2(t)] \epsilon_\theta \left(\frac{\alpha_{t-\Delta t}}{\alpha_t} \sigma_t - \sigma_{t-\Delta t} \right) \\ &\quad - \gamma(t) \epsilon \sqrt{\int_{t-\Delta t}^t (\sigma_t d\sigma_t - \frac{\sigma_t^2}{\alpha_t} d\alpha_t)}. \end{aligned}$$

To derive the close-form solution, we consider the specific VP and VE diffusion models. For the VP diffusion model ($\alpha_t^2 + \sigma_t^2 = 1$), we have

$$\begin{aligned} \mathbf{X}_{t-\Delta t} &= \frac{\alpha_{t-\Delta t}}{\alpha_t} \mathbf{X}_t - [1 + \gamma^2(t)] \epsilon_\theta \left(\frac{\alpha_{t-\Delta t}}{\alpha_t} \sigma_t - \sigma_{t-\Delta t} \right) \\ &\quad - \gamma(t) \epsilon \alpha_{t-\Delta t} \sqrt{\log \frac{\alpha_{t-\Delta t}}{\alpha_t}}. \end{aligned}$$

Moreover, for the VE diffusion model, we have

$$\begin{aligned} \mathbf{X}_{t-\Delta t} &= \frac{\alpha_{t-\Delta t}}{\alpha_t} \mathbf{X}_t - [1 + \gamma^2(t)] \epsilon_\theta \left(\frac{\alpha_{t-\Delta t}}{\alpha_t} \sigma_t - \sigma_{t-\Delta t} \right) \\ &\quad - \gamma(t) \epsilon \alpha_{t-\Delta t} \sqrt{\alpha_{t-\Delta t}^2 - \alpha_t^2}. \end{aligned}$$

APPENDIX C

ADAPTIVELY ADJUSTING NOISE INTENSITY IS NECESSARY FOR ALIGNMENT OF DIFFUSION TRAJECTORIES

In this section, we analyze that for a reinforcement alignment process of diffusion models, we need to adaptively adjusting the intensity to compensate the score estimation error. Considering the time-stamp of the to be aligned feature map as t , according to the DPM-Solver, we can calculate \mathbf{X}_0 by \mathbf{X}_t as

$$\mathbf{X}_0 = \frac{\alpha_0}{\alpha_t} \mathbf{X}_t + \left(\frac{\sigma_\tau}{\alpha_\tau} \Big|_{\tau=t}^{\tau=0} \right) \alpha_0 \epsilon_\theta(\mathbf{X}_t, t).$$

Then, considering for the ground-truth \mathbf{X}_0^* , we have

$$\mathbf{X}_0^* = \frac{\alpha_0}{\alpha_t} \mathbf{X}_t^* + \left(\frac{\sigma_\tau}{\alpha_\tau} \Big|_{\tau=t}^{\tau=0} \right) \alpha_0 \epsilon_\theta(\mathbf{X}_t^*, t),$$

where \mathbf{X}_t^* is obtained via adding random perturbation to \mathbf{X}_t , i.e.; $\mathbf{X}_t^* = \mathbf{X}_t + \gamma\epsilon$, ($\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$). Thus, we have

$$\begin{aligned} &\|\mathbf{X}_0 - \mathbf{X}_0^*\|_2 \\ &= \left\| \frac{\alpha_0}{\alpha_t} (\mathbf{X}_t - \mathbf{X}_t^*) + \left(\frac{\sigma_\tau}{\alpha_\tau} \Big|_{\tau=t}^{\tau=0} \right) \alpha_0 (\epsilon_\theta(\mathbf{X}_t, t) - \epsilon_\theta(\mathbf{X}_t^*, t)) \right\|_2 \\ &\stackrel{\textcircled{1}}{\approx} \left\| \frac{\alpha_0}{\alpha_t} (\mathbf{X}_t - \mathbf{X}_t^*) + \left(\frac{\sigma_\tau}{\alpha_\tau} \Big|_{\tau=t}^{\tau=0} \right) \alpha_0 k (\mathbf{X}_t - \mathbf{X}_t^*) \right\|_2 \\ &= \left| \frac{\alpha_0}{\alpha_t} + \alpha_0 k \left(\frac{\sigma_\tau}{\alpha_\tau} \Big|_{\tau=t}^{\tau=0} \right) \right| \|\mathbf{X}_t - \mathbf{X}_t^*\|_2, \end{aligned}$$

where $\textcircled{1}$ by assuming that the noise prediction network $\epsilon_\theta(\cdot)$ can accurately predict all the noise. Thus, we have $\epsilon_\theta(\mathbf{X}_t, t) - \epsilon_\theta(\mathbf{X}_t^*, t) \approx k(\mathbf{X}_t - \mathbf{X}_t^*)$. Then, we have

$$\|\mathbf{X}_t - \mathbf{X}_t^*\|_2 \approx \frac{\|\mathbf{X}_0 - \mathbf{X}_0^*\|_2}{\left| \frac{\alpha_0}{\alpha_t} - k\alpha_0 \left(\frac{\sigma_\tau}{\alpha_\tau} \Big|_{\tau=0}^{\tau=t} \right) \right|}.$$

Considering the parameterization $\mathbf{X}_t^* = \mathbf{X}_t + \gamma\epsilon$, ($\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$), then we have $\|\mathbf{X}_t - \mathbf{X}_t^*\|_2 = \gamma$. We can find that γ is correlated with reconstruction error $\|\mathbf{X}_0 - \mathbf{X}_0^*\|_2$ and the timestamp t to be adjusted. Thus, we parameterize a learnable function $\gamma_\phi(\|\mathbf{X}_0 - \mathbf{X}_0^*\|_2, t)$.

APPENDIX D

TRAJECTORY DISTILLATION COST

A. Calculation of Distillation Cost

In this section, we illustrate the detailed steps for the calculation of distillation cost. As we define the distillation cost as

$$\mathcal{C} = \sum_{i=k}^0 \left\| \check{\epsilon} \left(\frac{\mathbf{X}_i^{i-1}}{t_i^{i-1}} \Big| \frac{d\mathbf{X}_\epsilon}{dt} \right) - \epsilon_\theta(\mathbf{X}_{t_i}, t_i) \right\|_2. \quad (28)$$

We then start by calculating $\frac{d\mathbf{X}}{d\alpha_t}$ via the DPM-Solver as,

$$\frac{d\mathbf{X}}{d\alpha_t} = \frac{\mathbf{X}_t}{\alpha_t} - \frac{(\sigma_{t-\Delta t} - \frac{\alpha_{t-\Delta t}}{\alpha_t}\sigma_t)\epsilon_\theta}{\alpha_t - \alpha_{t-\Delta t}},$$

Subsequently, we can obtain

$$\check{\epsilon} = \left[\frac{\mathbf{X}_t}{\alpha_t} - \frac{\mathbf{X}_i^{i-1}}{\alpha_i^{i-1}} \right] \frac{\alpha_t - \alpha_{t-\Delta t}}{\sigma_{t-\Delta t} - \frac{\alpha_{t-\Delta t}}{\alpha_t}\sigma_t}.$$

While, for the VE diffusion model, the gradient is given by

$$\frac{d\mathbf{X}}{d\sigma_t} = -\epsilon_\theta.$$

This leads to the inverted noise

$$\check{\epsilon} = -\frac{\mathbf{X}_i^{i-1}}{\alpha_i^{i-1}}.$$

Finally, the \mathcal{C} can be derived by calculating the L2 Norm of $\check{\epsilon} - \epsilon_\theta$.

B. Proof of Existence

In this section, we first illustrate the existence of a low-cost distillation strategy. For an arbitrary continuous trajectory and $k \geq 2$, there will always be a distillation with a lower cost than straight flow-based distillation, e.g., a rectified flow and consistency model. While for $k = 1$, it results in the same cost as straight flow-based distillation. Denote by the \mathbf{X}_T and \mathbf{X}_0 as the clean image and initialized noise, with T and 0 as corresponding timestamps. Moreover, $T_1 \in (0, T)$ is a mid timestamp. Then, the distillation cost of the straight flow can be formulated as

$$\begin{aligned} \left\| \frac{d\mathbf{X}}{dt_0} - \frac{\mathbf{X}_T - \mathbf{X}_0}{T} \right\|_2^2 &= \left\| \frac{d\mathbf{X}}{dt_0} - \frac{\mathbf{X}_T - \mathbf{X}_0}{T} \right\|_2^2 \\ &= \|\mathbf{A} + \mathbf{B}\|_2^2, \\ \mathbf{A} &= \frac{T_1 \frac{d\mathbf{X}}{dt_0} - \int_0^{T_1} \frac{d\mathbf{X}}{dt} dt}{T} + \frac{(T - T_1) \frac{d\mathbf{X}}{dt_1} - \int_{T_1}^T \frac{d\mathbf{X}}{dt} dt}{T}, \\ \mathbf{B} &= \frac{(T - T_1) \left(\frac{d\mathbf{X}}{dt_0} - \frac{d\mathbf{X}}{dt_1} \right)}{T}. \end{aligned}$$

Here, we utilize T_1 to divide $[0, T]$ into 2 segments. Considering a single variable x , we have

$$\begin{aligned} \mathbf{a} &= \frac{T_1 \frac{dx}{dt_0} - \int_0^{T_1} \frac{dx}{dt} dt}{T} + \frac{(T - T_1) \frac{dx}{dt_1} - \int_{T_1}^T \frac{dx}{dt} dt}{T}, \\ &\stackrel{\textcircled{1}}{=} \frac{T_1}{T} \frac{dx}{dt_0} + \frac{T - T_1}{T} \frac{dx}{dt_1} - \frac{dx}{dt_i} \\ \mathbf{b} &= \frac{(T - T_1) \left(\frac{dx}{dt_0} - \frac{dx}{dt_1} \right)}{T}. \end{aligned}$$

① for the mean value theorem of integral. If we let $t_1 = t_i$, we then have

$$\mathbf{a} = \frac{T_1}{T} \left(\frac{dx}{dt_0} - \frac{dx}{dt_i} \right).$$

Thus, we then have $|\mathbf{a} + \mathbf{b}| = \left| \frac{dx}{dt_0} - \frac{dx}{dt_i} \right| \geq \frac{T_i}{T} \left| \frac{dx}{dt_0} - \frac{dx}{dt_i} \right| = |\mathbf{a}|$. Thus, there will always be a low-cost distillation point for $k = 2$ than $k = 1$. We can then easily derive similar results for higher k with iterative applying aforementioned process in sub-intervals.

APPENDIX E

SYNTHESIZING NOISE LATENT HELPS IR DIFFUSION

Assuming the noise-prediction neural $\epsilon_\theta(\cdot)$ has Lipschitz continuity, i.e., $\|\epsilon_\theta(\mathbf{X}_0) - \epsilon_\theta(\mathbf{X}_1)\|_2 \leq k \|\mathbf{X}_0 - \mathbf{X}_1\|_2$. Then, we take a one-step distillation model for example

$$\mathbf{X}_{T \rightarrow 0} = \frac{\alpha_0}{\alpha_T} \mathbf{X}_T + \alpha_0 e^{-\lambda} \Big|_{\lambda_T}^{\lambda_0} \hat{\epsilon}_\theta(\hat{\mathbf{X}}_T, T),$$

$$\mathbf{X}_{T-\delta \rightarrow 0} = \frac{\alpha_0}{\alpha_{T-\delta}} \mathbf{X}_{T-\delta} + \alpha_0 e^{-\lambda} \Big|_{\lambda_{T-\delta}}^{\lambda_0} \hat{\epsilon}_\theta(\hat{\mathbf{X}}_{T-\delta}, T - \delta),$$

We can then get the optimal noise estimation via making $\mathbf{X}_{T \rightarrow 0} \mathbf{X}_{T-\delta \rightarrow 0}$ to be the ground-truth value \mathbf{X}_0 . Then the optimal noise can be formulated as

$$\begin{aligned} \epsilon_{T-\delta \rightarrow 0} &= \frac{\frac{\mathbf{X}}{\sigma_0} - SNR_0 \mathbf{Y} - \frac{SNR_0}{SNR_{T-\delta}} \epsilon}{1 - \frac{SNR_0}{SNR_{T-\delta}}}, \\ \epsilon_{T \rightarrow 0} &= \frac{\frac{\mathbf{X}}{\sigma_0} - \frac{SNR_0}{SNR_T} \epsilon}{1 - \frac{SNR_0}{SNR_T}}, \end{aligned}$$

where SNR indicates the signal to noise ratio, i.e., $SNR_0 = \frac{\alpha_0}{\sigma_0}$. We suppose that the potential error of a neural network \mathcal{E} is positively correlated with the shift between the target and input, i.e., $\mathcal{E} = k \|\mathbf{X}_{in} - \mathbf{X}_{out}\|_2$, where $k > 0$. Thus, we can easily measure the magnitude of error by calculating the following ratio:

$$\begin{aligned} \frac{\mathcal{E}_{T-\delta}}{\mathcal{E}_T} &= \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \frac{\|\epsilon_{T-\delta \rightarrow 0} - \epsilon\|_2}{\|\epsilon_{T \rightarrow 0} - \epsilon\|_2} \\ &= \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \frac{\left\| \frac{\frac{\mathbf{X}}{\sigma_0} - SNR_0 \mathbf{Y} - \epsilon}{1 - \frac{SNR_0}{SNR_{T-\delta}}} \right\|_2}{\left\| \frac{\frac{\mathbf{X}}{\sigma_0} - \epsilon}{1 - \frac{SNR_0}{SNR_T}} \right\|_2} \\ &\stackrel{\textcircled{1}}{\approx} \frac{SNR_{T-\delta}}{SNR_T} \frac{\mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left\| \frac{\mathbf{X}}{\sigma_0} - SNR_0 \mathbf{Y} - \epsilon \right\|_2}{\mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left\| \frac{\mathbf{X}}{\sigma_0} - \epsilon \right\|_2} \\ &= \frac{SNR_{T-\delta}}{SNR_T} \frac{\sqrt{\left\| \frac{\mathbf{X}}{\sigma_0} - SNR_0 \mathbf{Y} \right\|_2^2 + \dim(\epsilon)}}{\sqrt{\left\| \frac{\mathbf{X}}{\sigma_0} \right\|_2^2 + \dim(\epsilon)}}, \\ &\stackrel{\textcircled{2}}{\approx} \frac{SNR_{T-\delta}}{SNR_T} \frac{\left\| \frac{\mathbf{X}}{\sigma_0} - SNR_0 \mathbf{Y} \right\|_2}{\left\| \frac{\mathbf{X}}{\sigma_0} \right\|_2} \stackrel{\textcircled{3}}{\approx} \frac{SNR_{T-\delta}}{SNR_T} \frac{\|\mathbf{X} - \mathbf{Y}\|_2}{\|\mathbf{X}\|_2} \end{aligned}$$

where $\dim(\cdot)$ indicates the number of elements in the input tensor. ① for $\frac{SNR_0}{SNR_{T-\delta}} \gg 1$, $\frac{SNR_0}{SNR_T} \gg 1$. ② for $\sigma_0 \rightarrow 0$ thus $\left\| \frac{\mathbf{X}}{\sigma_0} \right\|_2^2 \gg \dim(\epsilon)$ and ③ for $\alpha_0 \rightarrow 1$. Meanwhile, for the image restoration tasks, we usually have $\|\mathbf{X} - \mathbf{Y}\|_2 \leq k_1 \|\mathbf{X}\|_2$, where $0 \leq k_1 \leq 1$. Thus, we have $\frac{\mathcal{E}_{T-\delta}}{\mathcal{E}_T} \leq k_1$. If we utilize a small network to adaptively learn an initialization value to replace \mathbf{Y} , we can further reduce k_1 .