

Reconciling Stochastic and Deterministic Strategies for Zero-shot Image Restoration using Diffusion Model in Dual

CVPR 2025

Chong Wang¹, Lanqing Guo², Zixuan Fu¹, Siyuan Yang¹, Hao Cheng³, Alex C. Kot¹, Bihan Wen^{1*}

¹Nanyang Technological University, Singapore

²The University of Texas at Austin, USA ³Hebei University of Technology, China

Presented by Zejia Fan
2026.6.7

Super resolution

- Downsample formulation
- Consider probability
- Bayesian

$$y = (x * k) \downarrow_s + n$$

y : LR
 x : HR

$$p(x|y)$$

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)}$$




$$p(x|y) \propto p(y|x)p(x)$$

Likelihood

Prior

Super resolution

- Generate HR from LR



Super resolution

- Downsample formulation
- Consider probability
- Bayesian

$$y = (x * k) \downarrow_s + n$$

y : LR
 x : HR

$$p(x|y)$$

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)}$$



$$p(x|y) \propto p(y|x)p(x)$$

Likelihood

Prior

Super resolution

- Maximum A Posteriori (MAP)

$$x^* = \arg \max_x p(x|y)$$


$$x^* = \arg \min_x [-\log p(y|x) - \log p(x)]$$

Diffusion



- **Forward / noising process**

- Sample data $p_0(\mathbf{x}_0)$ → turn to noise




- **Reverse / denoising process**

- Sample noise $\mathbf{x}_T$ → turn into data

Diffusion for Image Restoration

■ Image Super-Resolution via Repeated Refinement (SR3)



- Conditional denoising diffusion model
- Forward diffusion process q , reverse inference process p
- Source image x , target image y_0

Diffusion for Image Restoration

■ Image Super-Resolution via Repeated Refinement (SR3)

**Conditioned reverse
inference process**


$$\begin{aligned} p_{\theta}(\mathbf{y}_{0:T} | \mathbf{x}) &= p(\mathbf{y}_T) \prod_{t=1}^T p_{\theta}(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathbf{x}) \\ p(\mathbf{y}_T) &= \mathcal{N}(\mathbf{y}_T | \mathbf{0}, \mathbf{I}) \\ p_{\theta}(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathbf{x}) &= \mathcal{N}(\mathbf{y}_{t-1} | \mu_{\theta}(\mathbf{x}, \mathbf{y}_t, \gamma_t), \sigma_t^2 \mathbf{I}) \end{aligned}$$

- The reverse inference process conditioned on source image $\mathbf{x}$
- $p_{\theta}(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathbf{x})$ learned by the neural network

Diffusion for Image Restoration

■ Image Super-Resolution via Repeated Refinement (SR3)


- U-Net architecture
- x interpolated to target high resolution
- x concatenate with noisy image y_t



Diffusion for Image Restoration

- Image Super-Resolution via Repeated Refinement (SR3)

$$x^* = \arg \min_x [-\log p(y|x) - \log p(x)]$$




Sampling-based Image Restoration

- Sampling restored images directly from pre-trained models
 - Take the degraded image as guidance

■ Example:

Denoising Diffusion Restoration Models
(DDRM)



Denoising Diffusion Restoration Models

Sampling-based Image Restoration

- Denoising Diffusion Restoration Models

- A simplified degradation model in general:

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(0, \sigma_y^2 \mathbf{I})$$

- The target is to find:

$$p_{\theta}(\mathbf{x}_{0:T}|\mathbf{y}) = p_{\theta}^{(T)}(\mathbf{x}_T|\mathbf{y}) \prod_{t=0}^{T-1} p_{\theta}^{(t)}(\mathbf{x}_t|\mathbf{x}_{t+1}, \mathbf{y})$$

- where θ is the parameter set of a pre-trained diffusion model.

Sampling-based Image Restoration

■ Denoising Diffusion Restoration Models

- The conditional distributions are demonstrated as:

Perform SVD $H = U\Sigma V^\top$

Distribution of the initial noise

$$p_{\theta}^{(T)}(\bar{\mathbf{x}}_T^{(i)} | \mathbf{y}) = \begin{cases} \mathcal{N}(\bar{\mathbf{y}}^{(i)}, \sigma_T^2 - \frac{\sigma_{\mathbf{y}}^2}{s_i^2}) & \text{if } s_i > 0 \\ \mathcal{N}(0, \sigma_T^2) & \text{if } s_i = 0 \end{cases}$$

Correct Sampling Distribution

$$p_{\theta}^{(t)}(\bar{\mathbf{x}}_t^{(i)} | \mathbf{x}_{t+1}, \mathbf{y}) = \begin{cases} \mathcal{N}(\bar{\mathbf{x}}_{\theta,t}^{(i)} + \sqrt{1 - \eta^2} \sigma_t \frac{\bar{\mathbf{x}}_{t+1}^{(i)} - \bar{\mathbf{x}}_{\theta,t}^{(i)}}{\sigma_{t+1}}, \eta^2 \sigma_t^2) & \text{if } s_i = 0 \\ \mathcal{N}(\bar{\mathbf{x}}_{\theta,t}^{(i)} + \sqrt{1 - \eta^2} \sigma_t \frac{\bar{\mathbf{y}}^{(i)} - \bar{\mathbf{x}}_{\theta,t}^{(i)}}{\sigma_{\mathbf{y}}/s_i}, \eta^2 \sigma_t^2) & \text{if } \sigma_t < \frac{\sigma_{\mathbf{y}}}{s_i} \\ \mathcal{N}((1 - \eta_b) \bar{\mathbf{x}}_{\theta,t}^{(i)} + \eta_b \bar{\mathbf{y}}^{(i)}, \sigma_t^2 - \frac{\sigma_{\mathbf{y}}^2}{s_i^2} \eta_b^2) & \text{if } \sigma_t \geq \frac{\sigma_{\mathbf{y}}}{s_i} \end{cases}$$

where $\bar{\mathbf{X}}_t$ is the set of vectors in spectral space: $\bar{\mathbf{x}}_t = V^\top \mathbf{x}_t$. Requires known degradation.

Sampling-based Image Restoration

■ Denoising Diffusion Restoration Models

- The conditional distributions are demonstrated as:

Perform SVD $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$

Correct Sampling Distribution

$$p_{\theta}^{(T)}(\bar{\mathbf{x}}_T^{(i)} | \mathbf{y}) = \begin{cases} \mathcal{N}(\bar{\mathbf{y}}^{(i)}, \sigma_T^2 - \frac{\sigma_{\mathbf{y}}^2}{s_i^2}) & \text{if } s_i > 0 \\ \mathcal{N}(0, \sigma_T^2) & \text{if } s_i = 0 \end{cases}$$

Distribution of one single denoise step

$$p_{\theta}^{(t)}(\bar{\mathbf{x}}_t^{(i)} | \mathbf{x}_{t+1}, \mathbf{y}) = \begin{cases} \mathcal{N}(\bar{\mathbf{x}}_{\theta,t}^{(i)} + \sqrt{1 - \eta^2} \sigma_t \frac{\bar{\mathbf{x}}_{t+1}^{(i)} - \bar{\mathbf{x}}_{\theta,t}^{(i)}}{\sigma_{t+1}}, \eta^2 \sigma_t^2) & \text{if } s_i = 0 \\ \mathcal{N}(\bar{\mathbf{x}}_{\theta,t}^{(i)} + \sqrt{1 - \eta^2} \sigma_t \frac{\bar{\mathbf{y}}^{(i)} - \bar{\mathbf{x}}_{\theta,t}^{(i)}}{\sigma_{\mathbf{y}}/s_i}, \eta^2 \sigma_t^2) & \text{if } \sigma_t < \frac{\sigma_{\mathbf{y}}}{s_i} \\ \mathcal{N}((1 - \eta_b) \bar{\mathbf{x}}_{\theta,t}^{(i)} + \eta_b \bar{\mathbf{y}}^{(i)}, \sigma_t^2 - \frac{\sigma_{\mathbf{y}}^2}{s_i^2} \eta_b^2) & \text{if } \sigma_t \geq \frac{\sigma_{\mathbf{y}}}{s_i} \end{cases}$$


where $\bar{\mathbf{X}}_t$ is the set of vectors in spectral space: $\bar{\mathbf{x}}_t = \mathbf{V}^\top \mathbf{x}_t$. Requires known degradation.

Sampling-based Image Restoration

■ Denoising Diffusion Restoration Models




(a) Inpainting results on cat images.



(b) Deblurring results ($\sigma_y = 0.05$) on bedroom images.

Sampling-based Image Restoration

■ Denoising Diffusion Models for Plug-and-Play Image Restoration (DiffPIR)



Employ a refinement during sampling process:

$$\hat{\mathbf{x}}_0^{(t)} = \arg \min_{\mathbf{x}} \boxed{\|\mathbf{y} - \mathcal{H}(\mathbf{x})\|^2} + \rho_t \|\mathbf{x} - \mathbf{x}_0^{(t)}\|^2$$

Fidelity to degraded image

where $\mathbf{x}_0^{(t)}$ is the predicted noise-free image at timestep t . Requires known degradation.

Sampling-based Image Restoration

■ Denoising Diffusion Models for Plug-and-Play Image Restoration



Employ a refinement during sampling process:



$$\hat{\mathbf{x}}_0^{(t)} = \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathcal{H}(\mathbf{x})\|^2 + \boxed{\rho_t \|\mathbf{x} - \mathbf{x}_0^{(t)}\|^2}$$

Regularization of generated image

where $\mathbf{x}_0^{(t)}$ is the predicted noise-free image at timestep t . Requires known degradation.

Sampling-based Image Restoration

Generative Diffusion Prior (GDP)



Learn the degradation model during the sampling process.

Sampling-based Image Restoration

■ Generative Diffusion Prior (GDP)


At the t -th step, the degradation model will be updated via:

$$\begin{aligned}\mathcal{L}_{\phi, \tilde{\mathbf{x}}_0}^{total} &= \mathcal{L}(\mathbf{y}, \mathcal{D}_{\phi}(\tilde{\mathbf{x}}_0)) + \mathcal{Q}(\tilde{\mathbf{x}}_0) \\ \phi &\leftarrow \phi - l \nabla_{\phi} \mathcal{L}_{\phi, \tilde{\mathbf{x}}_0}^{total}\end{aligned}$$


The noisy image at $(t - 1)$ -th step will be sampled with guidance provided by the loss:

$$\text{Sample } \mathbf{x}_{t-1} \text{ by } \mathcal{N}\left(\mu + s \nabla_{\tilde{\mathbf{x}}_0} \mathcal{L}_{\phi, \tilde{\mathbf{x}}_0}^{total}, \Sigma\right)$$

Sampling-based Image Restoration



Sampling-based Image Restoration




Sampling-based Image Restoration

■ Diffusion Posterior Sampling (DPS)

refine the noisy state at sampling process via:

$$\mathbf{x}_{i-1} \leftarrow \mathbf{x}'_{i-1} - \zeta_i \nabla_{\mathbf{x}_i} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$$




Sampling-based Image Restoration

■ Diffusion Posterior Sampling (DPS)


refine the noisy state at sampling process via:

$$\mathbf{x}_{i-1} \leftarrow \mathbf{x}'_{i-1} - \zeta_i \nabla_{\mathbf{x}_i} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$$


Correction Item



Sampling-based Image Restoration




Sampling-based Image Restoration




Reconciling Diffusion Model in Dual

- Based on pre-trained diffusion model
- Half quadratic splitting: two complementary regularizers, stochastic sampling and deterministic regression




Reconciling Diffusion Model in Dual

Deterministic Regression




$$z_T = y$$

Stochastic Sampling



$$x_T \sim N(0, I)$$

Reconciling Diffusion Model in Dual



Reconciling Diffusion Model in Dual

- HQS algorithm

$$\mathbf{x}_t = \operatorname{argmin}_{\mathbf{x}} \frac{1}{2(\sqrt{\tau\lambda}/\mu)^2} \|\mathbf{x} - \mathbf{z}_t\|_2^2 + \mathcal{R}_\theta^S(\mathbf{x}),$$

$$\mathbf{z}_{t-1} = \operatorname{argmin}_{\mathbf{z}} \frac{1}{2} \|\mathbf{y} - A\mathbf{z}\|_2^2 + \frac{\mu\sigma_n^2}{2} \|\mathbf{z} - \mathbf{x}_t\|_2^2 \\ + (1 - \tau)\lambda\sigma_n^2 \mathcal{R}_\theta^D(\mathbf{z}).$$

- Overall

$$\min_{\mathbf{z}, \mathbf{x}} \frac{1}{2\sigma_n^2} \|\mathbf{y} - A\mathbf{z}\|_2^2 + \frac{\mu}{2} \|\mathbf{z} - \mathbf{x}\|_2^2 + \boxed{\tau\lambda\mathcal{R}_\theta(\mathbf{x})} + \boxed{(1 - \tau)\lambda\mathcal{R}_\theta(\mathbf{z})}$$

Fidelity & perceptual

Reconciling Diffusion Model in Dual

Algorithm 1 RDMD

Require: $\mathbf{y}, \mathcal{N}_\theta, \lambda, \eta, \tau, \zeta, \{\sigma'_t\}_{t=1}^T$ defined in Section 3.1.

Initialization: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \mathbf{z}_T = A^\top \mathbf{y}$

for $t = T, \dots, 1$ **do**

pre-calculate t' for deterministic resgression

$$t' = \operatorname{argmin}_{i \in [1, T]} \left| \sigma'_t - \frac{\sqrt{1 - \bar{\alpha}_i}}{\sqrt{\bar{\alpha}_i}} \right|$$

stochastic estimation

$$\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t; t))$$

deterministic estimation

$$f_\theta(\mathbf{z}_t; t') = \frac{1}{\sqrt{\bar{\alpha}_{t'}}} (\mathbf{z}_t - \sqrt{1 - \bar{\alpha}_{t'}} \boldsymbol{\epsilon}_\theta(\mathbf{z}_t; t'))$$

main update process

$$\mathbf{z}_{t-1} = \mathbf{z}_t - \eta \left(\nabla_{\mathbf{z}_t} \|\mathbf{y} - A\mathbf{z}_t\|_2^2 + \tau \lambda \frac{\sigma_n^2}{\sigma_t^2} (\mathbf{z}_t - \mathbf{x}_{0|t}) + (1 - \tau) \lambda \sigma_n^2 (\mathbf{z}_t - f_\theta(\mathbf{z}_t; t')) \right)$$

renoising

$$\hat{\boldsymbol{\epsilon}}_t = \frac{1}{\sqrt{1 - \bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \mathbf{z}_{t-1})$$

$$\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$


$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \mathbf{z}_{t-1} + \sqrt{1 - \bar{\alpha}_{t-1}} (\sqrt{1 - \zeta} \hat{\boldsymbol{\epsilon}}_t + \sqrt{\zeta} \boldsymbol{\epsilon}_t)$$

end for

return $\mathbf{x}_0$

Experiment

- Fidelity & perceptual



Experiment

- Fidelity & perceptual

		$\tau = 0.0$ PSNR: 26.44 dB ✓ LPIPS: 0.3363 ✗
		$\tau = 0.1$ PSNR: 26.56 dB ✓ LPIPS: 0.3073 ✓
		$\tau = 1.0$ PSNR: 24.30dB ✗ LPIPS: 0.3158 ✓
Original	Restorations	

Experiment

- Quantitative Comparison - FFHQ


FFHQ	Deblur (<i>Gaussian</i>)			Deblur (<i>motion</i>)			Super-resolution ($2\times$)			Super-resolution ($4\times$)			Super-resolution ($8\times$)		
Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
MCG[8]	11.19	0.0968	0.8042	11.33	0.0831	0.8088	19.57	0.3482	0.5648	17.34	0.2153	0.7190	16.81	0.1829	0.7676
DPS[9]	23.78	0.6776	0.2829	20.64	0.5732	0.3461	25.06	0.7418	0.2738	22.85	0.6503	0.3215	20.54	0.5597	0.3726
Resample[29]	23.50	0.4962	0.4612	18.60	0.2566	0.6455	23.07	0.4754	0.4828	18.28	0.2563	0.6691	18.74	0.2712	0.6807
BIRD[7]	26.14	0.7009	0.2928	24.21	0.6419	0.3412	27.04	0.6150	0.3310	22.51	0.3870	0.5190	19.67	0.3070	0.6340
RED[25](Ours - deterministic)	<u>27.84</u>	0.8096	0.3170	26.09	0.7372	0.3613	27.87	0.7919	0.2562	<u>27.05</u>	<u>0.7960</u>	0.3277	<u>22.22</u>	<u>0.6431</u>	0.4883
DiffPIR[43](Ours - stochastic)	27.30	0.7672	0.2372	<u>26.63</u>	<u>0.7399</u>	0.2529	<u>29.39</u>	<u>0.8060</u>	<u>0.2172</u>	26.27	0.7332	0.2674	21.08	0.5634	<u>0.3666</u>
Ours	28.29	<u>0.8057</u>	<u>0.2770</u>	27.05	0.7776	<u>0.2767</u>	31.03	0.8715	0.1807	27.84	0.8012	<u>0.2768</u>	23.31	0.6524	0.3614

Experiment


- Quantitative Comparison - ImageNet

ImageNet	Deblur (<i>Gaussian</i>)			Deblur (<i>motion</i>)			Super-resolution ($2\times$)			Super-resolution ($4\times$)			Super-resolution ($8\times$)		
Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
MCG[8]	11.11	0.0982	0.7338	11.05	0.0362	0.7535	15.12	0.2217	0.6259	16.31	0.2150	0.6727	15.87	0.1507	0.7271
DPS[9]	20.09	0.4774	0.4272	20.14	0.4871	0.4336	21.31	0.5514	0.4392	19.47	0.4510	0.4961	17.65	0.3522	<u>0.5505</u>
Resample[29]	18.88	0.2957	0.5579	15.49	0.1857	0.6571	21.01	0.4674	0.4334	16.46	0.2235	0.6163	17.55	0.2207	0.6968
BIRD[7]	22.04	0.5360	0.4107	20.79	0.4976	0.4499	24.94	0.6841	0.3047	21.37	0.4521	0.5001	19.67	0.3073	0.6334
RED[25](Ours - deterministic)	<u>23.24</u>	0.6090	0.4282	22.52	0.5674	0.4963	<u>25.87</u>	<u>0.7199</u>	0.3023	<u>23.16</u>	<u>0.6027</u>	0.4324	<u>20.42</u>	<u>0.4634</u>	0.6024
DiffPIR[43](Ours - stochastic)	21.97	0.5363	0.3827	23.34	0.6276	0.3736	25.33	0.7092	<u>0.2900</u>	22.80	0.5967	0.3871	19.44	0.4292	0.5411
Ours	23.30	<u>0.6004</u>	<u>0.4079</u>	<u>22.77</u>	<u>0.6037</u>	<u>0.4247</u>	25.94	0.7563	0.2648	23.28	0.6203	<u>0.4000</u>	20.83	0.4736	0.5690

Experiment



Experiment



Sampling-based Restoration


DiTVR: Zero-Shot Diffusion Transformer for Video Restoration

• CVPR 26

Sicheng Gao Nancy Mehta Zongwei Wu Radu Timofte

Computer Vision Lab, CAIDAS & IFI, University of Würzburg

sicheng.gao@uni-wuerzburg.de



Conclusion

- Restoration problem can be formulated as 2 main parts
 - Likelihood - fidelity
 - Prior – perceptual quality
- Diffusion prior can act as prior regularization
- Balance between fidelity and perceptual quality is still tricky

Thanks for your listening!