Hybrid Regularization Improves Diffusion-based Inverse Problem Solving

Hongkun Dou, Zeyu Li, Jinyang Du, Lijun Yang, Wen Yao, and Yue Deng

https://github.com/deng-ai-lab/HRDIS

April 24, 2025

Background: Inverse problems

Super-resolution

Inpainting

Deblurring

Observations (Inputs)
$\mathbf{y}$

Outputs from our method
$\mathbf{x}_0$

$$\mathbf{y} = \mathcal{A}(\mathbf{x}_0) + \boldsymbol{\eta}, \quad \mathbf{y}, \boldsymbol{\eta} \in \mathbb{R}^n, \mathbf{x}_0 \in \mathbb{R}^m$$

Background: Diffusion models

Denoising Diffusion Models

Learning to generate by denoising

Denoising diffusion models consist of two processes:

- Forward diffusion process that gradually adds noise to input
- Reverse denoising process that learns to generate data by denoising

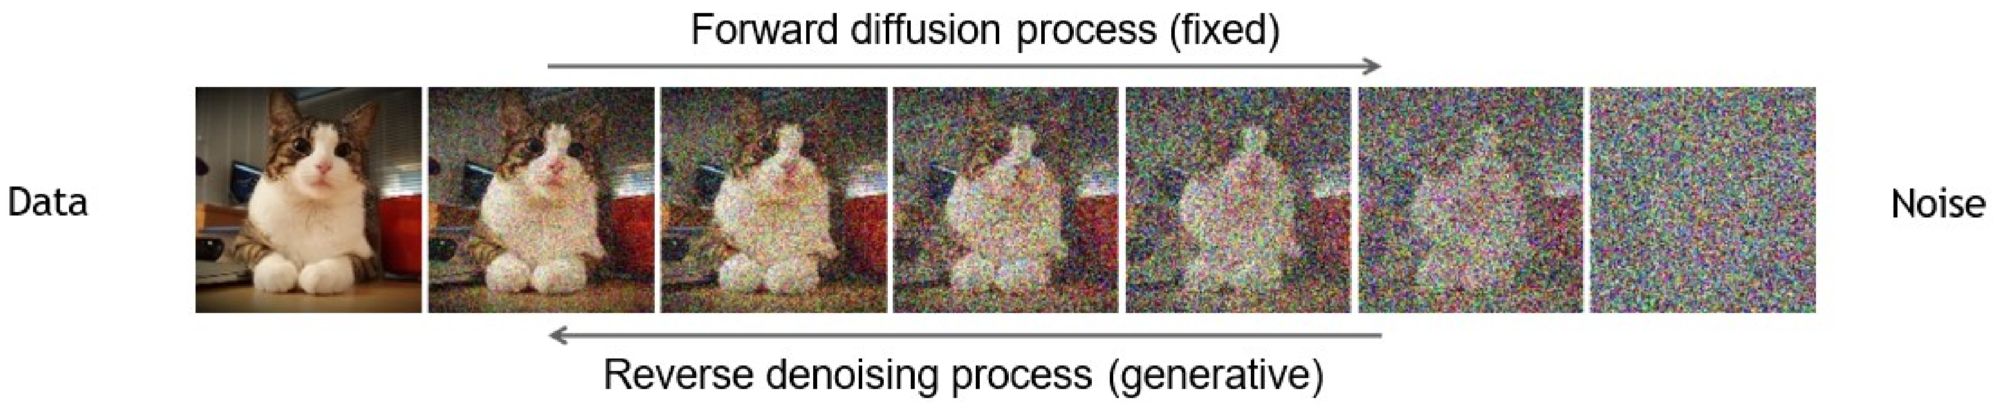

[Sohl-Dickstein et al., Deep Unsupervised Learning using Nonequilibrium Thermodynamics, ICML 2015]
[Ho et al., Denoising Diffusion Probabilistic Models, NeurIPS 2020]
[Song et al., Score-Based Generative Modeling through Stochastic Differential Equations, ICLR 2021]

Background: Diffusion models

$$\mathbf{dx}_t = -\frac{1}{2}\beta(t)\mathbf{x}_t\,\mathbf{d}t + \sqrt{\beta(t)}\,\mathbf{d}\boldsymbol{\omega}_t$$

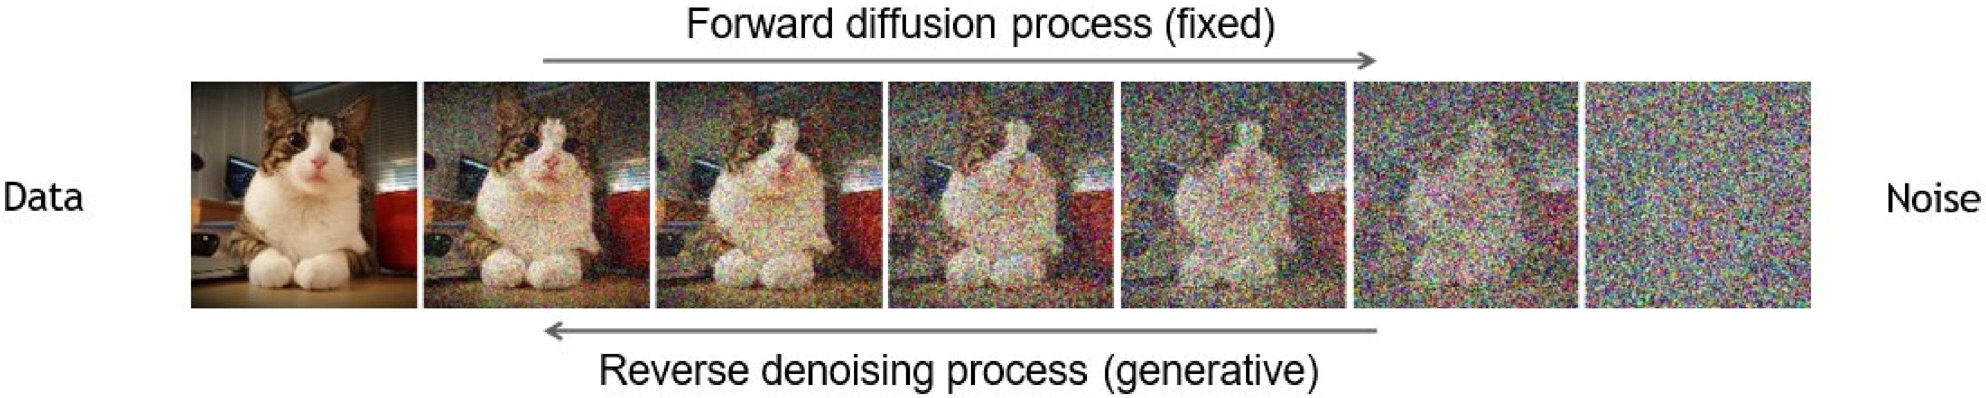

$$\mathrm{d}\mathbf{x}_t = \overbrace{\left[-\frac{1}{2}\beta(t)\mathbf{x}_t - \beta(t)\underbrace{\nabla_{\mathbf{x}_t}\log q_t(\mathbf{x}_t)}_{\text{“Score Function”}}\right]}^{\text{drift term}}\mathrm{d}t + \overbrace{\sqrt{\beta(t)}\,\mathrm{d}\bar{\boldsymbol{\omega}}_t}^{\text{diffusion term}}$$

Sohl-Dickstein et al., Deep Unsupervised Learning using Nonequilibrium Thermodynamics, ICML 2015
Ho et al., Denoising Diffusion Probabilistic Models, NeurIPS 2020
Song et al., Score-Based Generative Modeling through Stochastic Differential Equations, ICLR 2021

Diffusion models for inverse problem

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \boldsymbol{\eta}$$

$$\mathbf{x} \sim p(\mathbf{x}|\mathbf{y})$$

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

$$\log p(\mathbf{x}|\mathbf{y}) = \log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x}) - \log p(\mathbf{y})$$

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y}) = \underbrace{\nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x})}_{\text{Measurement process}} + \underbrace{\nabla_{\mathbf{x}} \log p(\mathbf{x})}_{\simeq s_{\theta^*}(\mathbf{x})}$$

Diffusion Posterior Sampling (DPS)

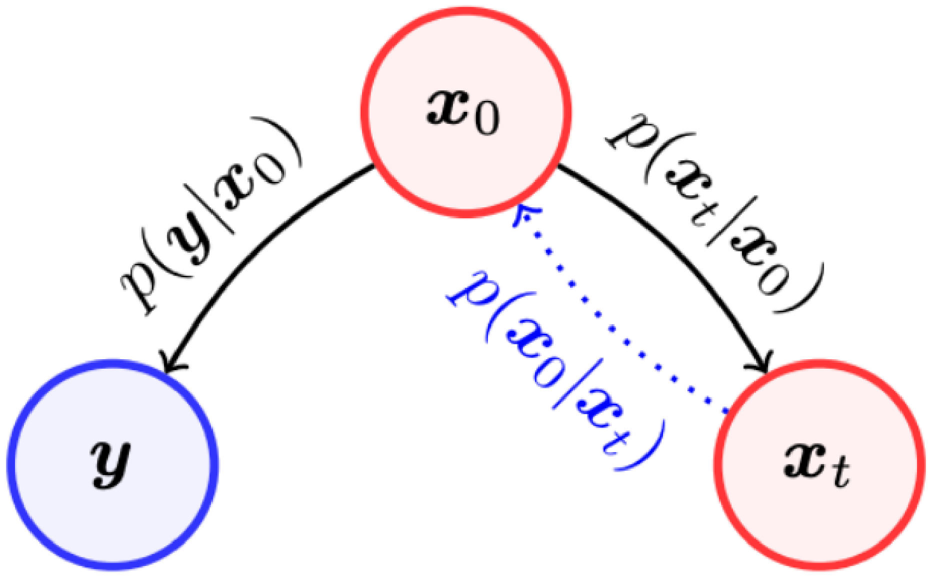

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$$
$$\simeq s_{\theta^*}(\mathbf{x}_t, t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\hat{\mathbf{x}}_0)$$

$$\hat{\mathrm{x}}_0 = \mathbb{E}[\mathrm{x}_0|\mathrm{x}_t]$$

Requires denoising network backpropagation!

Chung, Hyungjin, et al. "Diffusion Posterior Sampling for General Noisy Inverse Problems." ICLR 2023

RED-diff

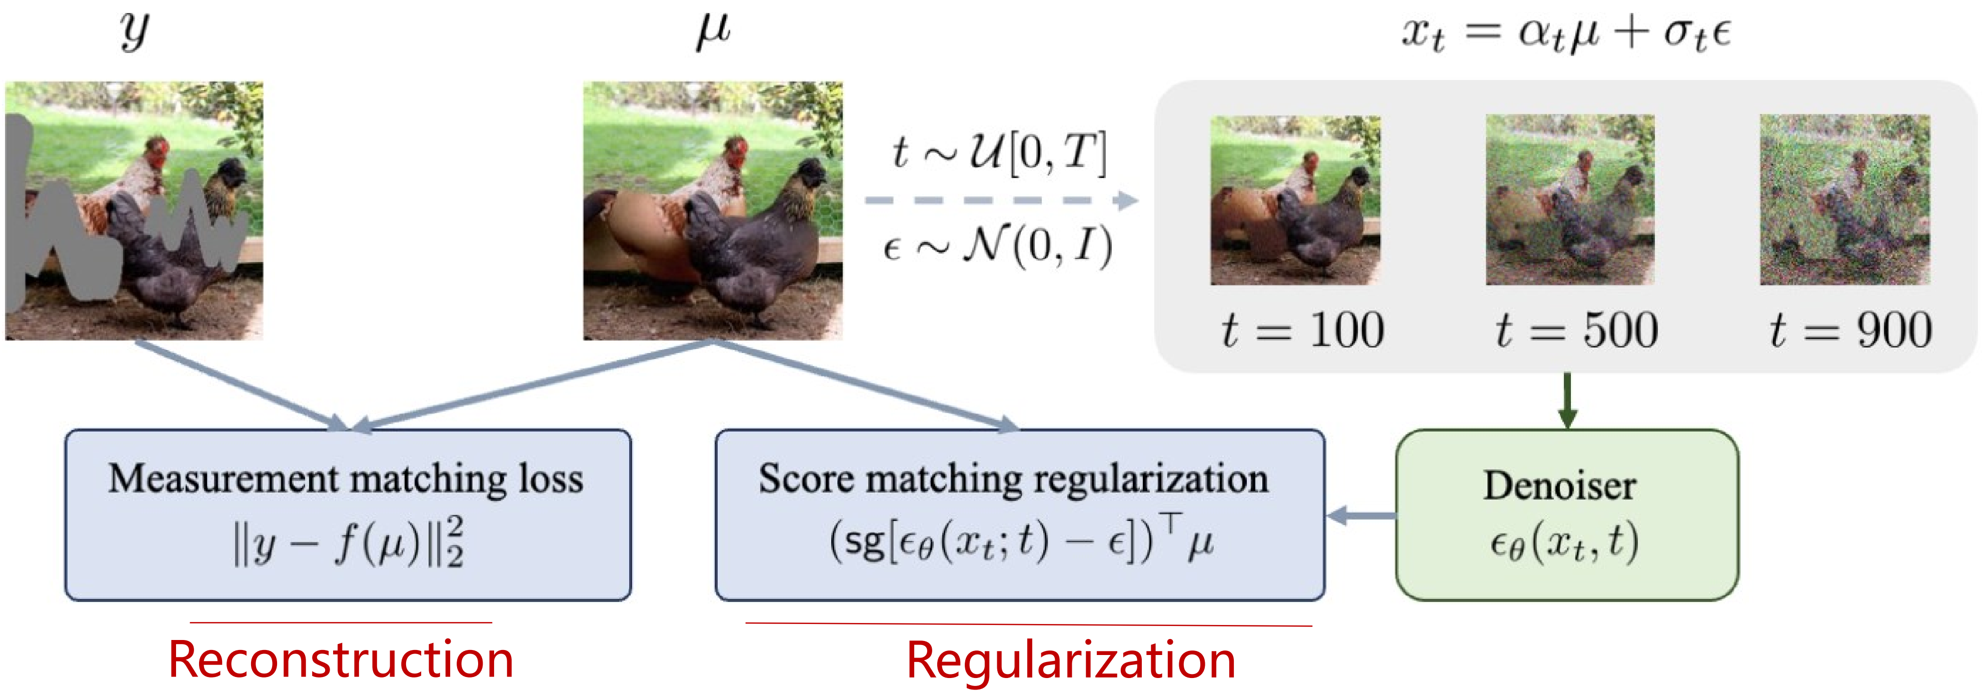

Mardani, Morteza, et al. "A Variational Perspective on Solving Inverse Problems with Diffusion Models." ICLR 2024.
Poole, Ben, et al. "DreamFusion: Text-to-3D using 2D Diffusion." ICLR 2023.

RED-diff

$$\nabla_\mu \mathcal{L}_{\text{DR}} = \nabla_\mu \|\mathbf{y} - \mathcal{A}(\mu)\|_2^2 + \boxed{\begin{aligned}&\mathbb{E}_{t,\boldsymbol{\epsilon}}[\omega_t(\boldsymbol{\epsilon}_\theta(\mu_t, t) - \boldsymbol{\epsilon})] \\ &= (\mu_t - \sigma_t\boldsymbol{\epsilon})/\alpha_t - (\mu_t - \sigma_t\boldsymbol{\epsilon}_\theta(\mu_t, t))/\alpha_t \\ &= \mu - \hat{\mu}_{0|t} = \nabla_\mu \|\mu - \text{sg}(\hat{\mu}_{0|t})\|_2^2,\end{aligned}}$$

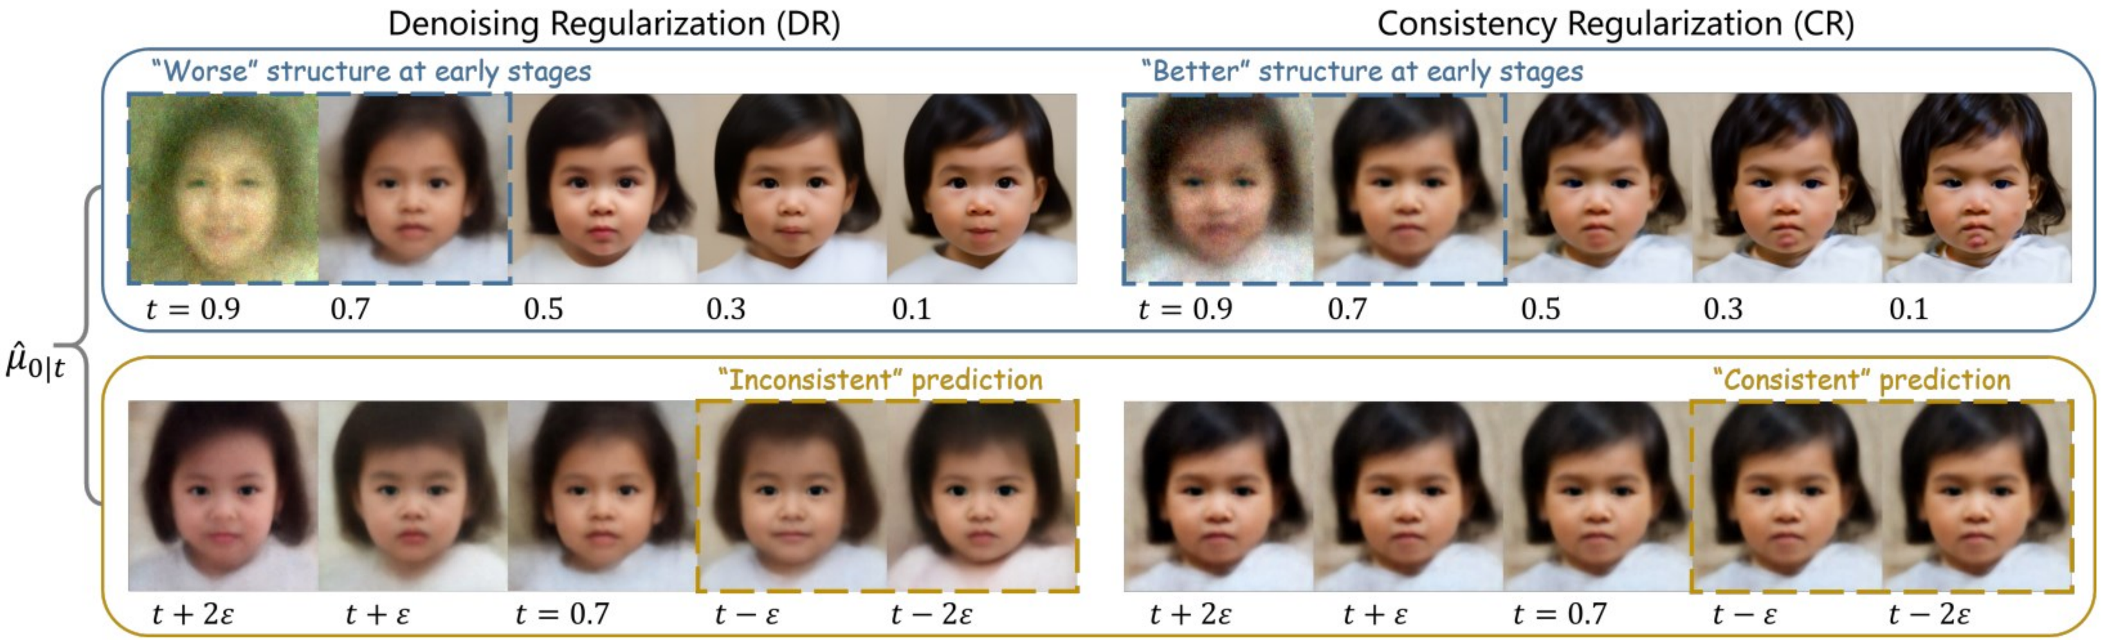

Mardani, Morteza, et al. "A Variational Perspective on Solving Inverse Problems with Diffusion Models." ICLR 2024.
Poole, Ben, et al. "DreamFusion: Text-to-3D using 2D Diffusion." ICLR 2023.

Consistency Regularization (CR)

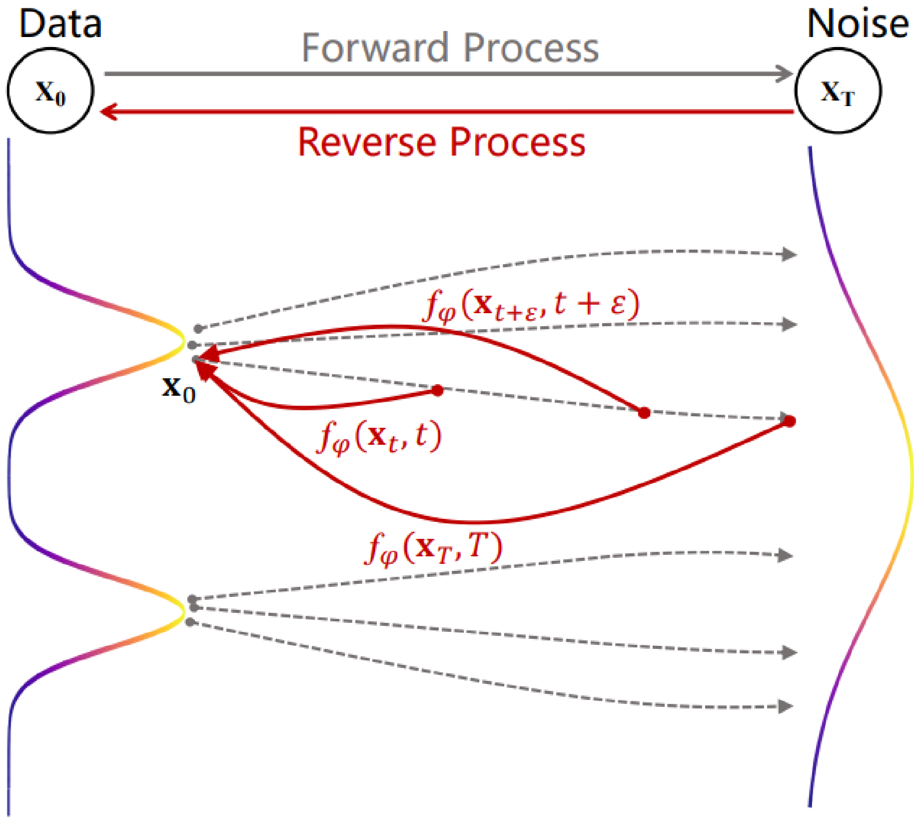

Consistency Distillation (CD)

Original objective

$$\|f_\varphi(\mathbf{x}_t, t) - \mathbf{x}_0\|_2^2$$

Simulation-free objective

$$\|f_\varphi(\mathbf{x}_{t+\varepsilon}, t+\varepsilon) - \mathrm{sg}(f_\varphi(\mathbf{x}_t, t))\|_2^2$$

Optimize φ

Consistency Regularization (CR)

Original objective

$$\|\mu - \mathbf{x}_0\|_2^2$$

Simulation-free objective

$$\mathrm{sg}(\hat{\mu}_{0|t+\varepsilon} - \hat{\mu}_{0|t})^\top \mu$$

Optimize μ

Self-consistency property

$$\nabla_\mu \mathcal{L}_{\mathrm{CR}} = \nabla_\mu \|\mathbf{y} - \mathcal{A}(\mu)\|_2^2 + \mathbb{E}_{t,\epsilon}\left[\boxed{\hat{\mu}_{0|t+\varepsilon} - \hat{\mu}_{0|t}}\right]$$

$$= \hat{\mu}_{0|t+\varepsilon} - (\mu_t - \sigma_t \boldsymbol{\epsilon}_\theta(\mu_t, t))/\alpha_t$$
$$= \omega_t(\boldsymbol{\epsilon}_\theta(\mu_t, t) - \boldsymbol{\epsilon}_\theta(\mu_{t+\varepsilon}, t+\varepsilon))$$

Song, Yang, et al. "Consistency Models." ICML, 2023.

Hybrid Regularization (HR)

DDIM Solver

$$\mathbf{x}_t = \alpha_t \underbrace{\hat{\mathbf{x}}_{0:t+\varepsilon}}_{\text{predicted } \mathbf{x}_0} + \underbrace{\sqrt{\sigma_t^2 - \beta_t^2}\,\boldsymbol{\epsilon}_\theta(\mathbf{x}_{t+\varepsilon}, t+\varepsilon)}_{\text{deterministic noise}} + \underbrace{\beta_t \boldsymbol{\epsilon}}_{\text{random noise}}$$

Total variance σ_t^2

$$\sqrt{\sigma_t^2 - \beta_t^2})^2 + \beta_t^2 = \sigma_t^2$$

Hybrid noise

$$\nabla_\mu \mathcal{L}_{\text{HR}} = \nabla_\mu \|\mathbf{y} - \mathcal{A}(\mu)\|_2^2 + \mathbb{E}_{t,\boldsymbol{\epsilon}}[\omega_t(\boldsymbol{\epsilon}_\theta(\mu_t, t) - \boxed{\boldsymbol{\epsilon}_{\text{hybrid}}})]$$

$$\boldsymbol{\epsilon}_{\text{hybrid}} \triangleq \sqrt{1-\beta}\,\boldsymbol{\epsilon}_\theta(\mu_{t+\varepsilon}, t+\varepsilon) + \sqrt{\beta}\,\boldsymbol{\epsilon}$$

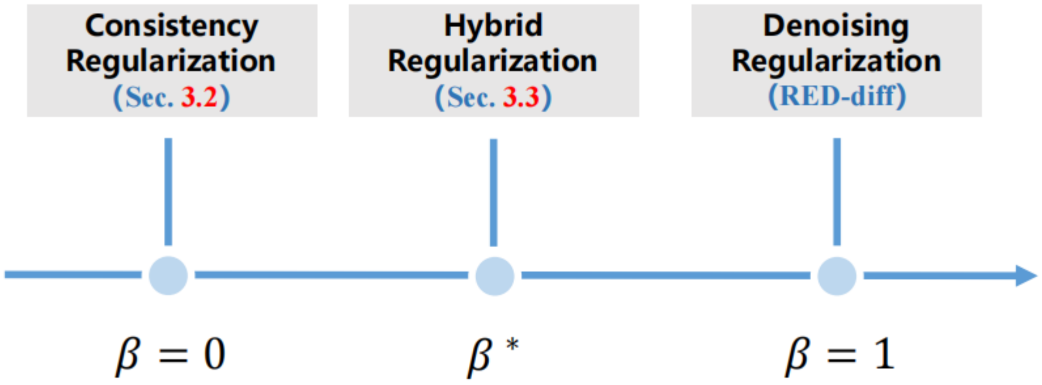

[Song, Jiaming, Chenlin Meng, and Stefano Ermon. "Denoising Diffusion Implicit Models." ICLR, 2021.](#)

HRDIS

Algorithm 1 Sampling procedure for HRDIS.

Input: observation $\mathbf{y}$, measurement operator $\mathcal{A}(\cdot)$, number of iterations N, timesteps sampling strategy $\{s_n\}_{n=1}^{N}$, pretrained model $\boldsymbol{\epsilon}_\theta(\cdot,\cdot)$, β, ω_t

```
Initialize $\mu \leftarrow \mathcal{A}^{-1}(\mathbf{y})$
for $n = 1, \cdots, N$ do
    $t \leftarrow s_n$
    if $n = 1$ then
        Initialize hybrid noise $\boldsymbol{\epsilon}_{\text{hybrid}} \sim \mathcal{N}(0, \boldsymbol{I})$
    end if
    Forward perturb $\mu_t \leftarrow \alpha_t \mu + \sigma_t \boldsymbol{\epsilon}_{\text{hybrid}}$
    Calculate gradient $d_\mu \leftarrow \nabla_\mu \|\mathbf{y} - \mathcal{A}(\mu)\|_2^2 + \omega_t(\boldsymbol{\epsilon}_\theta(\mu_t, t) - \boldsymbol{\epsilon}_{\text{hybrid}})$
    Optimize mean $\mu \leftarrow \text{AdamUpdate}(\mu, d_\mu)$
    Sample fresh noise $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$
    Calculate hybrid noise $\boldsymbol{\epsilon}_{\text{hybrid}} \leftarrow \sqrt{1-\beta}\,\boldsymbol{\epsilon}_\theta(\mu_t, t) + \sqrt{\beta}\,\boldsymbol{\epsilon}$
end for
```

Output: μ

HRDIS

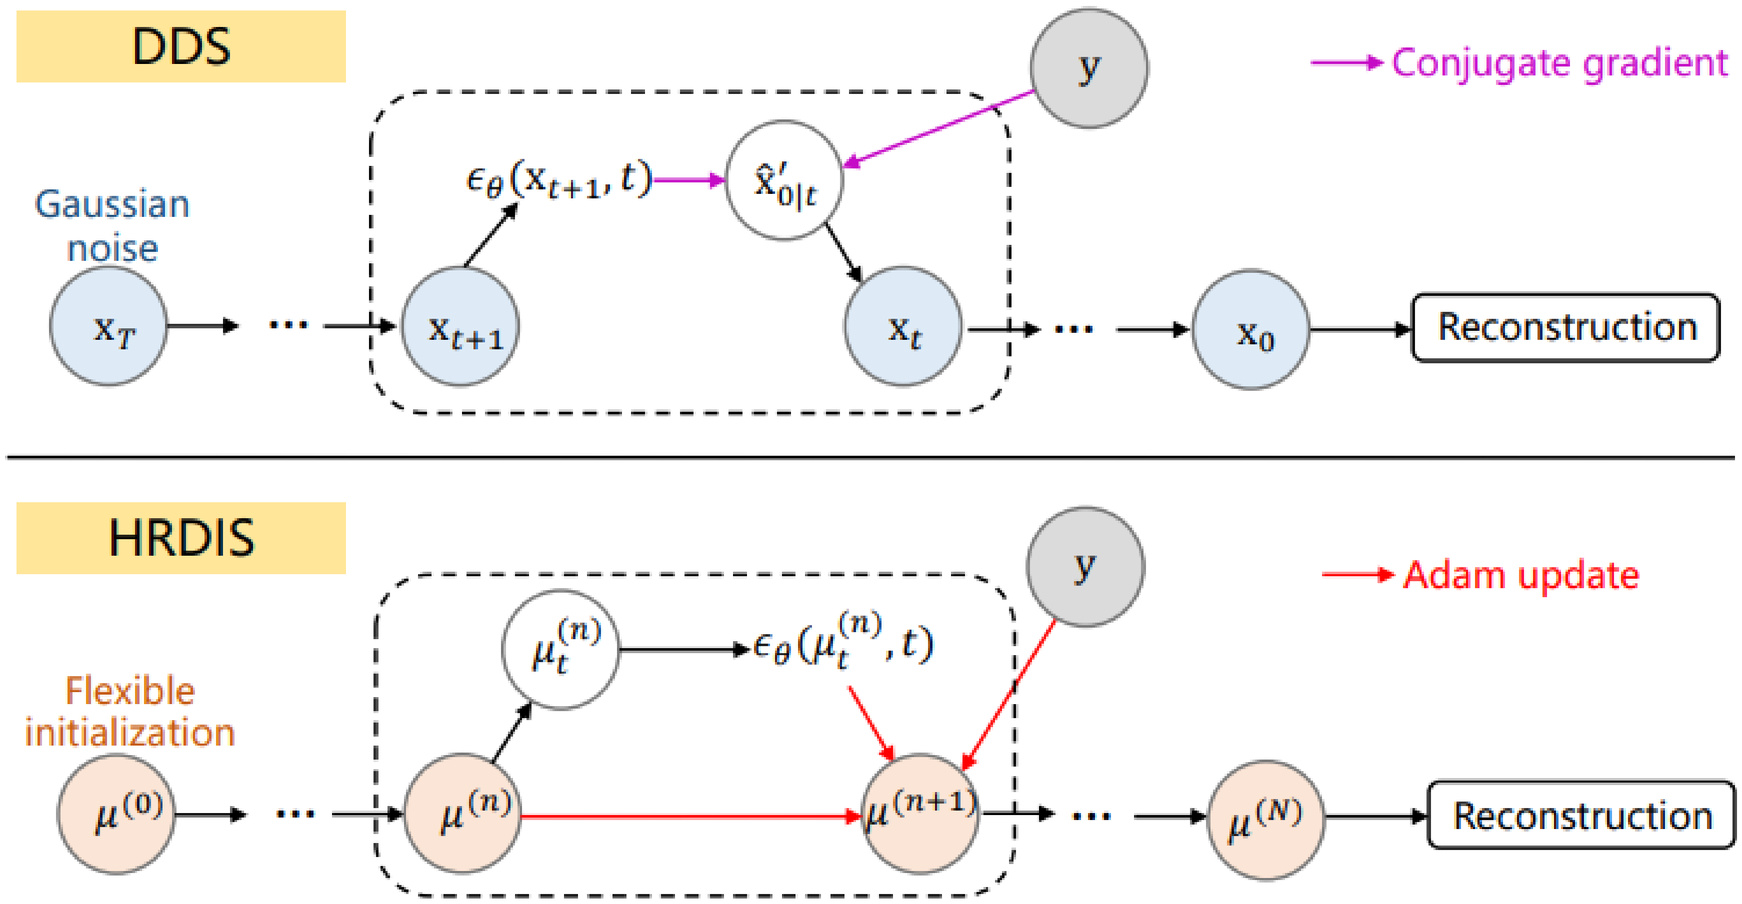

DDS
Conjugate gradient
y
Gaussian noise
$\epsilon_\theta(x_{t+1}, t)$
$\hat{x}'_{0|t}$
x_T
x_{t+1}
x_t
x_0
Reconstruction
HRDIS
Adam update
y
Flexible initialization
$\mu_t^{(n)}$
$\epsilon_\theta(\mu_t^{(n)}, t)$
$\mu^{(0)}$
$\mu^{(n)}$
$\mu^{(n+1)}$
$\mu^{(N)}$
Reconstruction

HRDIS

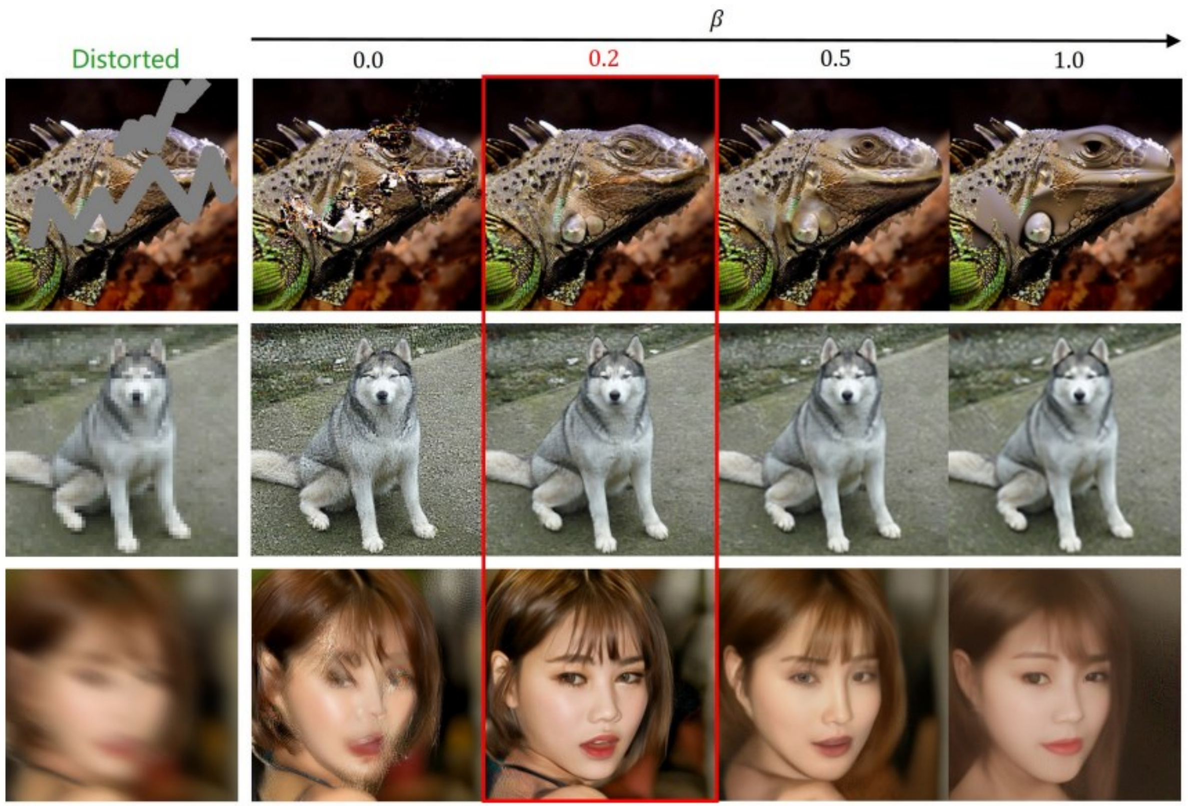

Ablation of β

HRDIS can produce diverse and clear inversion results

Conclusion

We address the challenges associated with solving diffusion-based inverse problems using Denoising Regularization (DR):

- We introduce the Consistency Regularization (CR) method, which effectively mitigates the issue of inaccurate gradient estimations.
- We explore the integration of hybrid noise, resulting in the development of a unified HRDIS framework that fosters synergy between the two regularization techniques
- The proposed framework is versatile, making it applicable to both linear and nonlinear inverse problems. HRDIS not only surpasses current state-of-the-art methods but also maintains high computational efficiency

Future works

- An enticing avenue for future exploration is the extension of our approach to tackling blind inverse problems.

- Additionally, a promising direction for advancement involves extending our methodology to encompass latent diffusion models.

- Since HRDIS involves stochastic optimization for non-convex problems, its convergence behavior remains an open theoretical question and warrants further investigation.

Thank you for listening!

Thank you for listening!