Provable Diffusion Posterior Sampling for Bayesian Inversion

Chenguang Duan
RWTH Aachen

Joint work with **Jinyuan Chang** (SWUFE & CAS), **Yuling Jiao** (WHU), **Ruoxuan Li** (WHU), **Jerry Zhijian Yang** (WHU), and **Cheng Yuan** (WHU)

December 10, 2025

Outline

Applications of Inverse Problems


(a) Medical Imaging


(b) Weather Prediction


(c) Image Processing

Mathematical Formulation of Bayesian Inverse Problems

Consider the measurement model:

$$\mathbf{Y} = \mathcal{F}(\mathbf{X}) + \mathbf{n}.$$

- $\mathcal{F} : \mathbb{R}^d \to \mathbb{R}^n$ is a known forward map.
- $\mathbf{X} \in \mathbb{R}^d$ is an unknown signal with prior density π.
- π is typically intractable but one has $\mathbf{X}_1, \ldots, \mathbf{X}_N \sim^{\text{i.i.d.}} \pi$.
- $\mathbf{n} \in \mathbb{R}^n$ is a random noise independent of $\mathbf{X}$ and with known density.

Mathematical Formulation of Bayesian Inverse Problems

Consider the measurement model:

$$\mathbf{Y} = \mathcal{F}(\mathbf{X}) + \mathbf{n}.$$

- $\mathcal{F} : \mathbb{R}^d \to \mathbb{R}^n$ is a known forward map.
- $\mathbf{X} \in \mathbb{R}^d$ is an unknown signal with prior density π.
- π is typically intractable but one has $\mathbf{X}_1, \ldots, \mathbf{X}_N \sim^{\text{i.i.d.}} \pi$.
- $\mathbf{n} \in \mathbb{R}^n$ is a random noise independent of $\mathbf{X}$ and with known density.

Given noisy measurement $\mathbf{Y} = \mathbf{y}$, the goal of Bayesian inverse problem is to estimate the posterior density

$$q(\mathbf{x}|\mathbf{y}) \propto \exp(-\ell_{\mathbf{y}}(\mathbf{x}))\pi(\mathbf{x}),$$

where $\ell_{\mathbf{y}}(\cdot)$ is the negative log-likelihood.

Mainstream Methods for Bayesian Inverse Problems

Optimization-based approach

Maximum-a-posteriori (MAP): Find a signal that minimizes the posterior density:

$$\arg\max_{\mathbf{x}} \log q(\mathbf{x}|\mathbf{y}) = \arg\max_{\mathbf{x}} -\ell_{\mathbf{y}}(\mathbf{x}) + \log \pi(\mathbf{x})$$

Mainstream Methods for Bayesian Inverse Problems

Optimization-based approach

Maximum-a-posteriori (MAP): Find a signal that minimizes the posterior density:

$$\arg\max_{\mathbf{x}} \log q(\mathbf{x}|\mathbf{y}) = \arg\max_{\mathbf{x}} -\ell_{\mathbf{y}}(\mathbf{x}) + \log \pi(\mathbf{x})$$

- Connection with Tikhonov regularization
- Point estimation, cannot quantify the uncertainty

Mainstream Methods for Bayesian Inverse Problems

Optimization-based approach

Maximum-a-posteriori (MAP): Find a signal that minimizes the posterior density:

$$\arg\max_{\mathbf{x}} \log q(\mathbf{x}|\mathbf{y}) = \arg\max_{\mathbf{x}} -\ell_{\mathbf{y}}(\mathbf{x}) + \log \pi(\mathbf{x})$$

- Connection with Tikhonov regularization
- Point estimation, cannot quantify the uncertainty

Sampling-based approach

Sampling from the posterior density $q(\cdot|\mathbf{y})$:

$$\mathbf{X}_1^{\mathbf{y}}, \ldots, \mathbf{X}_n^{\mathbf{y}} \sim^{\text{i.i.d.}} q(\cdot|\mathbf{y}).$$

Mainstream Methods for Bayesian Inverse Problems

Optimization-based approach

Maximum-a-posteriori (MAP): Find a signal that minimizes the posterior density:

$$\arg\max_{\mathbf{x}} \log q(\mathbf{x}|\mathbf{y}) = \arg\max_{\mathbf{x}} -\ell_{\mathbf{y}}(\mathbf{x}) + \log \pi(\mathbf{x})$$

- Connection with Tikhonov regularization
- Point estimation, cannot quantify the uncertainty

Sampling-based approach

Sampling from the posterior density $q(\cdot|\mathbf{y})$:

$$\mathbf{X}_1^{\mathbf{y}}, \ldots, \mathbf{X}_n^{\mathbf{y}} \sim^{\text{i.i.d.}} q(\cdot|\mathbf{y}).$$

- Uncertainty quantification

Key Questions in Posterior Sampling

Q1. *How to estimate the prior distribution using samples drawn from it?*

Key Questions in Posterior Sampling

Q1. *How to estimate the prior distribution using samples drawn from it?*

- Manual selection of the prior: cannot capture the complex structure of the prior distribution.
- Generative prior (Vishal Purohit et al. 2025)
- Prior score matching (Hyungjin Chung et al. 2023, Jiaming Song et al. 2023)
- ...

Key Questions in Posterior Sampling

Q1. *How to estimate the prior distribution using samples drawn from it?*

- Manual selection of the prior: cannot capture the complex structure of the prior distribution.
- Generative prior (Vishal Purohit et al. 2025)
- Prior score matching (Hyungjin Chung et al. 2023, Jiaming Song et al. 2023)
- ...

Q2. *How to sample from the multi-modal posterior distribution?*

Key Questions in Posterior Sampling

Q1. *How to estimate the prior distribution using samples drawn from it?*

- Manual selection of the prior: cannot capture the complex structure of the prior distribution.
- Generative prior (Vishal Purohit et al. 2025)
- Prior score matching (Hyungjin Chung et al. 2023, Jiaming Song et al. 2023)
- ...

Q2. *How to sample from the multi-modal posterior distribution?*

- Langevin dynamics (Vishal Purohit et al. 2025, Zhao Ding et al. 2025)
- Diffusion models (Hyungjin Chung et al. 2023, Jiaming Song et al. 2023)
- ...

Key Questions in Posterior Sampling

Q1. *How to estimate the prior distribution using samples drawn from it?*

- Manual selection of the prior: cannot capture the complex structure of the prior distribution.
- Generative prior (Vishal Purohit et al. 2025)
- Prior score matching (Hyungjin Chung et al. 2023, Jiaming Song et al. 2023)
- ...

Q2. *How to sample from the multi-modal posterior distribution?*

- Langevin dynamics (Vishal Purohit et al. 2025, Zhao Ding et al. 2025)
- Diffusion models (Hyungjin Chung et al. 2023, Jiaming Song et al. 2023)
- ...

Prior Score + Diffusion Models represent a highly promising solution!

Diffusion Posterior Sampling

Forward process

$$\mathrm{d}\mathbf{X}_t = -\mathbf{X}_t \mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \mathbf{X}_0 \sim \pi, \quad t \in (0, T).$$

Time-reversal process

$$\mathrm{d}\bar{\mathbf{X}}_t = (\bar{\mathbf{X}}_t + 2\nabla_{\mathbf{x}} \log \pi_{T-t}(\bar{\mathbf{X}}_t))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0 \sim \pi_T, \; t \in (0, T). \qquad \bar{\mathbf{X}}_t \sim \pi_{T-t}$$

Diffusion Posterior Sampling

Forward process

$$\mathrm{d}\mathbf{X}_t = -\mathbf{X}_t \mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \mathbf{X}_0 \sim \pi, \quad t \in (0, T).$$

Time-reversal process

$$\mathrm{d}\bar{\mathbf{X}}_t = (\bar{\mathbf{X}}_t + 2\nabla_{\mathbf{x}} \log \pi_{T-t}(\bar{\mathbf{X}}_t))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0 \sim \pi_T, \ t \in (0, T). \qquad \bar{\mathbf{X}}_t \sim \pi_{T-t}$$

Posterior time-reversal process

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y}))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y}) \qquad \bar{\mathbf{X}}_t \sim q_{T-t}(\cdot|\mathbf{y})$$

Diffusion Posterior Sampling

Forward process

$$\mathrm{d}\mathbf{X}_t = -\mathbf{X}_t \mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \mathbf{X}_0 \sim \pi, \quad t \in (0, T).$$

Time-reversal process

$$\mathrm{d}\bar{\mathbf{X}}_t = (\bar{\mathbf{X}}_t + 2\nabla_{\mathbf{x}} \log \pi_{T-t}(\bar{\mathbf{X}}_t))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0 \sim \pi_T, \ t \in (0, T). \qquad \bar{\mathbf{X}}_t \sim \pi_{T-t}$$

Posterior time-reversal process

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y}))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y}) \qquad \bar{\mathbf{X}}_t \sim q_{T-t}(\cdot|\mathbf{y})$$

How to estimate the posterior score $\nabla \log q_t(\cdot|\mathbf{y})$ for $t \in (0, T)$?

Existing Approximation Methods for Posterior Score Estimation

$$\nabla \log q_t(\mathbf{x}_t|\mathbf{y}) = \nabla \log p(\mathbf{y}|\mathbf{x}_t) + \nabla \log \pi_t(\mathbf{x}_t)$$

Existing Approximation Methods for Posterior Score Estimation

$$\nabla \log q_t(\mathbf{x}_t|\mathbf{y}) = \nabla \log p(\mathbf{y}|\mathbf{x}_t) + \nabla \log \pi_t(\mathbf{x}_t)$$

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}_t) &= \int p(\mathbf{y}|\mathbf{x}_0, \mathbf{x}_t) p(\mathbf{x}_0|\mathbf{x}_t) \mathrm{d}\mathbf{x}_0 \\ &= \int \exp(-\ell_{\mathbf{y}}(\mathbf{x}_0)) p(\mathbf{x}_0|\mathbf{x}_t) \mathrm{d}\mathbf{x}_0 \end{aligned}$$

Existing Approximation Methods for Posterior Score Estimation

$$\nabla \log q_t(\mathbf{x}_t|\mathbf{y}) = \nabla \log p(\mathbf{y}|\mathbf{x}_t) + \nabla \log \pi_t(\mathbf{x}_t)$$

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}_t) &= \int p(\mathbf{y}|\mathbf{x}_0,\mathbf{x}_t)p(\mathbf{x}_0|\mathbf{x}_t)\mathrm{d}\mathbf{x}_0 \\ &= \int \exp(-\ell_{\mathbf{y}}(\mathbf{x}_0))p(\mathbf{x}_0|\mathbf{x}_t)\mathrm{d}\mathbf{x}_0 \end{aligned}$$

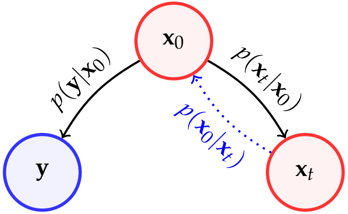

The time-dependent likelihood $p(\mathbf{y}|\mathbf{x}_t)$ is intractable.

Existing Approximation Methods for Posterior Score Estimation

$$\nabla \log q_t(\mathbf{x}_t|\mathbf{y}) = \nabla \log p(\mathbf{y}|\mathbf{x}_t) + \nabla \log \pi_t(\mathbf{x}_t)$$

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}_t) &= \int p(\mathbf{y}|\mathbf{x}_0, \mathbf{x}_t) p(\mathbf{x}_0|\mathbf{x}_t) \mathrm{d}\mathbf{x}_0 \\ &= \int \exp(-\ell_{\mathbf{y}}(\mathbf{x}_0)) p(\mathbf{x}_0|\mathbf{x}_t) \mathrm{d}\mathbf{x}_0 \end{aligned}$$

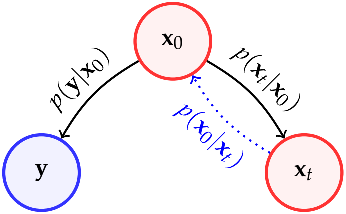

The time-dependent likelihood $p(\mathbf{y}|\mathbf{x}_t)$ is intractable.

Approximation Methods

DPS (Hyungjin Chung et al. 2023)	Dirac approximation	$p(\mathbf{x}_0|\mathbf{x}_t) \approx \delta_{\mathbb{E}[\mathbf{X}_0|\mathbf{x}_t]}(\mathbf{x}_0)$
LGD (Jiaming Song et al. 2023)	Gaussian approximation	$p(\mathbf{x}_0|\mathbf{x}_t) \approx N(\mathbb{E}[\mathbf{X}_0|\mathbf{x}_t], r_t^2 \mathbf{I}_d)$
...	...	...

Outline

Posterior Score Estimation via Monte Carlo

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y})}_{\text{posterior score}})\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y})$$

Posterior Score Estimation via Monte Carlo

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y})}_{\text{posterior score}})\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y})$$

Conditional Tweedie's formula

For each $t \in (0, T)$, the score function of the posterior density satisfies

$$\nabla_{\mathbf{x}} \log q_t(\mathbf{x}|\mathbf{y}) = -\frac{1}{\sigma_t^2}\mathbf{x} + \frac{\mu_t}{\sigma_t^2}\mathbf{D}(t, \mathbf{x}, \mathbf{y}), \quad (\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^n,$$

where the posterior denoisor $\mathbf{D}(t, \mathbf{x}, \mathbf{y})$ is defined as

$$\mathbf{D}(t, \mathbf{x}, \mathbf{y}) := \mathbb{E}[\mathbf{X}_0|\mathbf{X}_t = \mathbf{x}, \mathbf{Y} = \mathbf{y}].$$

Posterior Score Estimation via Monte Carlo

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y})}_{\text{posterior score}})\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y})$$

Conditional Tweedie's formula

For each $t \in (0, T)$, the score function of the posterior density satisfies

$$\nabla_{\mathbf{x}} \log q_t(\mathbf{x}|\mathbf{y}) = -\frac{1}{\sigma_t^2}\mathbf{x} + \frac{\mu_t}{\sigma_t^2}\mathbf{D}(t, \mathbf{x}, \mathbf{y}), \quad (\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^n,$$

where the posterior denoisor $\mathbf{D}(t, \mathbf{x}, \mathbf{y})$ is defined as

$$\mathbf{D}(t, \mathbf{x}, \mathbf{y}) := \mathbb{E}[\mathbf{X}_0|\mathbf{X}_t = \mathbf{x}, \mathbf{Y} = \mathbf{y}].$$

Expectation with respect to $\quad p_t(\mathbf{x}_0|\mathbf{x}, \mathbf{y}) \propto \pi(\mathbf{x}_0) \exp\left(-\frac{\|\mathbf{x} - \mu_t \mathbf{x}_0\|_2^2}{2\sigma_t^2} - \ell_{\mathbf{y}}(\mathbf{x}_0)\right)$

Sampling from the posterior denoising density

- Langevin dynamics

$$\mathrm{d}\mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t} = \Big(\nabla \log \pi(\mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t}) + \frac{\mu_t}{\sigma_t^2}(\mathbf{x} - \mu_t \mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t}) - \nabla \ell_{\mathbf{y}}(\mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t})\Big)\mathrm{d}s + \sqrt{2}\mathrm{d}\mathbf{B}_s, \quad s \in (0, S)$$

- Langevin dynamics with estimated prior score

$$\mathrm{d}\widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t} = \Big(\widehat{\mathbf{s}}_\pi(\widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t}) + \frac{\mu_t}{\sigma_t^2}(\mathbf{x} - \mu_t \widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t}) - \nabla \ell_{\mathbf{y}}(\widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t})\Big)\mathrm{d}s + \sqrt{2}\mathrm{d}\mathbf{B}_s, \quad s \in (0, S).$$

Sampling from the posterior denoising density

- Langevin dynamics

$$\mathrm{d}\mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t} = \Big(\nabla \log \pi(\mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t}) + \frac{\mu_t}{\sigma_t^2}(\mathbf{x} - \mu_t \mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t}) - \nabla \ell_{\mathbf{y}}(\mathbf{X}_{0,s}^{\mathbf{x},\mathbf{y},t})\Big)\mathrm{d}s + \sqrt{2}\mathrm{d}\mathbf{B}_s, \quad s \in (0, S)$$

- Langevin dynamics with estimated prior score

$$\mathrm{d}\widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t} = \Big(\widehat{\mathbf{s}}_{\pi}(\widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t}) + \frac{\mu_t}{\sigma_t^2}(\mathbf{x} - \mu_t \widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t}) - \nabla \ell_{\mathbf{y}}(\widehat{\mathbf{X}}_{0,s}^{\mathbf{x},\mathbf{y},t})\Big)\mathrm{d}s + \sqrt{2}\mathrm{d}\mathbf{B}_s, \quad s \in (0, S).$$

Monte Carlo approximation

- Estimation of the posterior score

$$\widehat{\mathbf{s}}_m^S(t, \mathbf{x}, \mathbf{y}) \triangleq -\frac{\mathbf{x}}{\sigma_t^2} + \frac{\mu_t}{\sigma_t^2}\frac{1}{m}\sum_{i=1}^m \widehat{\mathbf{X}}_{0,S,i}^{\mathbf{x},\mathbf{y},t} \approx \nabla_{\mathbf{x}} \log q_t(\mathbf{x}|\mathbf{y}), \quad \widehat{\mathbf{X}}_{0,S,1}^{\mathbf{x},\mathbf{y},t}, \ldots, \widehat{\mathbf{X}}_{0,S,m}^{\mathbf{x},\mathbf{y},t} \sim^{\text{i.i.d.}} \widehat{p}_t^S(\cdot|\mathbf{x}, \mathbf{y}).$$

Algorithm 1: Posterior score estimation via Monte Carlo

Input: Measurement $\mathbf{y}$, prior score estimator $\widehat{\mathbf{s}}_\pi$, noise level t, evaluation point $\mathbf{x}$, simulation horizon S, initial density $\widehat{p}_t^0(\cdot|\mathbf{y})$ for Langevin dynamics.

Output: Posterior score estimator $\widehat{\mathbf{s}}_m^S(t, \mathbf{x}, \mathbf{y})$.

for $i = 1, \ldots, m$ **do**

  Initialize particle by $\widehat{\mathbf{X}}_{0,0,i}^{\mathbf{x},\mathbf{y},t} \sim \widehat{p}_t^0(\cdot|\mathbf{x}, \mathbf{y})$.

  Evolve particle to time S by simulating Langevin dynamics to obtain $\widehat{\mathbf{X}}_{0,S,i}^{\mathbf{x},\mathbf{y},t}$.

end

Compute the posterior denoiser estimate $\widehat{\mathbf{D}}_m^S(t, \mathbf{x}, \mathbf{y})$ using $\{\widehat{\mathbf{X}}_{0,S,i}^{\mathbf{x},\mathbf{y},t}\}_{i=1}^m$ via Monte Carlo method.

Compute the posterior score estimate $\widehat{\mathbf{s}}_m^S(t, \mathbf{x}, \mathbf{y})$ using $\widehat{\mathbf{D}}_m^S(t, \mathbf{x}, \mathbf{y})$.

return $\widehat{\mathbf{s}}_m^S(t, \mathbf{x}, \mathbf{y})$

- **A1.** Semi-log-concavity of the posterior distribution.
  For a fixed $\mathbf{y} \in \mathbb{R}^n$, there exist a constant $\alpha > 0$, such that $q_0(\cdot|\mathbf{y}) \in C^2(\mathbb{R}^d)$ satisfies $-\nabla^2 \log q_0(\mathbf{x}_0|\mathbf{y}) \succeq -\alpha \mathbf{I}_d$ for each $\mathbf{x}_0 \in \mathbb{R}^d$.

Log-concavity of posterior denoising density

Suppose **A1** holds, then for each $t > 0$,

$$\left(\frac{\mu_t^2}{\sigma_t^2} - \alpha\right)\mathbf{I}_d \preceq -\nabla_{\mathbf{x}_0}^2 \log p_t(\mathbf{x}_0|\mathbf{x}, \mathbf{y}), \quad (\mathbf{x}_0, \mathbf{x}) \in \mathbb{R}^d \times \mathbb{R}^d.$$

Further, the posterior denoising density is log-concave provided that $t < \bar{t} := \frac{1}{2}\log(1 + \alpha^{-1})$.

- **A1.** Semi-log-concavity of the posterior distribution.
  For a fixed $\mathbf{y} \in \mathbb{R}^n$, there exist a constant $\alpha > 0$, such that $q_0(\cdot|\mathbf{y}) \in C^2(\mathbb{R}^d)$ satisfies $-\nabla^2 \log q_0(\mathbf{x}_0|\mathbf{y}) \succeq -\alpha \mathbf{I}_d$ for each $\mathbf{x}_0 \in \mathbb{R}^d$.

Log-concavity of posterior denoising density

Suppose **A1** holds, then for each $t > 0$,

$$\Big(\frac{\mu_t^2}{\sigma_t^2} - \alpha\Big)\mathbf{I}_d \preceq -\nabla_{\mathbf{x}_0}^2 \log p_t(\mathbf{x}_0|\mathbf{x}, \mathbf{y}), \quad (\mathbf{x}_0, \mathbf{x}) \in \mathbb{R}^d \times \mathbb{R}^d.$$

Further, the posterior denoising density is log-concave provided that $t < \bar{t} := \frac{1}{2}\log(1 + \alpha^{-1})$.

Monte Carlo-based approach for posterior score estimation works for small $t > 0$.

Warm-Start Strategy

Goal of warm-start: Sample from the terminal posterior density $q_T(\cdot|\mathbf{y})$.

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y}))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y})}_{\text{warm-start}}$$

- For sufficiently large T, we can use Gaussian approximation $q_T(\cdot|\mathbf{y}) \approx N(\mathbf{0}, \mathbf{I}_d)$.
- However, Monte Carlo-based approach for posterior score estimation requires small T.

Warm-Start Strategy

Goal of warm-start: Sample from the terminal posterior density $q_T(\cdot|\mathbf{y})$.

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y}))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y})}_{\text{warm-start}}$$

- For sufficiently large T, we can use Gaussian approximation $q_T(\cdot|\mathbf{y}) \approx N(\mathbf{0}, \mathbf{I}_d)$.
- However, Monte Carlo-based approach for posterior score estimation requires small T.

Warm-start strategy

Langevin dynamics with estimated posterior score:

$$\mathrm{d}\widehat{\mathbf{X}}_{T,u}^{\mathbf{y}} = \widehat{\mathbf{s}}_m^S(T, \widehat{\mathbf{X}}_{T,u}^{\mathbf{y}}, \mathbf{y})\mathrm{d}u + \sqrt{2}\mathrm{d}\mathbf{B}_u, \quad u \in (0, U).$$

Here $\widehat{\mathbf{s}}_m^S(T, \cdot, \mathbf{y})$ is obtained by Monte Carlo-based posterior score estimation.

Algorithm 2: Warm-start for diffusion posterior sampling

Input: Measurement $\mathbf{y}$, simulation horizon U, terminal density $\widehat{q}_T^0(\cdot|\mathbf{y})$ for Langevin dynamics.

Output: Sample $\widehat{\mathbf{X}}_{T,U}^{\mathbf{y}}$ approximately distributed according to $q_T(\cdot|\mathbf{y})$.

1 Initialize particle by $\widehat{\mathbf{X}}_{T,0}^{\mathbf{y}} \sim \widehat{q}_T^0(\cdot|\mathbf{y})$.

2 Simulate the approximate Langevin dynamics from 0 to U using the estimated posterior score function $\widehat{\mathbf{s}}_m^S(t, \mathbf{x}, \mathbf{y})$ to obtain $\widehat{\mathbf{X}}_{T,U}^{\mathbf{y}}$.

3 **return** $\widehat{\mathbf{X}}_{T,U}^{\mathbf{y}}$

- **A2.** Sub-Gaussian tails.
  The log-prior density $\log \pi_0 \in C^2(\mathbb{R}^d)$, and there exist $V_{\mathrm{SG}} > 0$ and $C_{\mathrm{SG}} > 0$, such that

$$\int \exp\left(\frac{\|\mathbf{x}_0\|_2^2}{V_{\mathrm{SG}}^2}\right)\pi_0(\mathbf{x}_0)\mathrm{d}\mathbf{x}_0 \leq C_{\mathrm{SG}}.$$

Log-Sobolev inequality of posterior

Suppose **A2** holds, then for each $T > \underline{t} := \frac{1}{2}\log(1+2V_{\mathrm{SG}}^2)$,

$$C_{\mathrm{LSI}}(q_T(\cdot|\mathbf{y})) \leq 12\sigma_T^2 \exp\left(2\frac{\sigma_T^2 + 2\mu_T^2 V_{\mathrm{SG}}^2}{\sigma_T^2 - 2\mu_T^2 V_{\mathrm{SG}}^2}\log(\kappa_{\mathbf{y}}^2 C_{\mathrm{SG}}^2)\right).$$

- **A2.** Sub-Gaussian tails.
  The log-prior density $\log \pi_0 \in C^2(\mathbb{R}^d)$, and there exist $V_{\mathrm{SG}} > 0$ and $C_{\mathrm{SG}} > 0$, such that

$$\int \exp\left(\frac{\|\mathbf{x}_0\|_2^2}{V_{\mathrm{SG}}^2}\right)\pi_0(\mathbf{x}_0)\mathrm{d}\mathbf{x}_0 \leq C_{\mathrm{SG}}.$$

Log-Sobolev inequality of posterior

Suppose **A2** holds, then for each $T > \underline{t} := \frac{1}{2}\log(1 + 2V_{\mathrm{SG}}^2)$,

$$C_{\mathrm{LSI}}(q_T(\cdot|\mathbf{y})) \leq 12\sigma_T^2 \exp\left(2\frac{\sigma_T^2 + 2\mu_T^2 V_{\mathrm{SG}}^2}{\sigma_T^2 - 2\mu_T^2 V_{\mathrm{SG}}^2}\log(\kappa_{\mathbf{y}}^2 C_{\mathrm{SG}}^2)\right).$$

The terminal posterior density $q_T(\cdot|\mathbf{y})$ satisfies LSI for large T.

Diffusion Posterior Sampling

$$\mathrm{d}\widehat{\mathbf{X}}_t^{\mathbf{y}} = \big(\widehat{\mathbf{X}}_t^{\mathbf{y}} + 2 \underbrace{\widehat{\mathbf{s}}_m^S(T-t, \widehat{\mathbf{X}}_t^{\mathbf{y}}, \mathbf{y})}_{\text{posterior score estimation}}\big)\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\widehat{\mathbf{X}}_0^{\mathbf{y}} \sim \widehat{q}_T^U(\cdot|\mathbf{y})}_{\text{warm-start}}, \ t \in (0, T - T_0),$$

Diffusion Posterior Sampling

$$\mathrm{d}\widehat{\mathbf{X}}_t^{\mathbf{y}} = (\widehat{\mathbf{X}}_t^{\mathbf{y}} + 2 \underbrace{\widehat{\mathbf{s}}_m^S(T-t, \widehat{\mathbf{X}}_t^{\mathbf{y}}, \mathbf{y})}_{\text{posterior score estimation}})\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\widehat{\mathbf{X}}_0^{\mathbf{y}} \sim \widehat{q}_T^U(\cdot|\mathbf{y})}_{\text{warm-start}}, \ t \in (0, T-T_0),$$

Algorithm 4: Diffusion-based posterior sampling

Input: Measurement $\mathbf{y}$, diffusion terminal time T, early-stopping time T_0.

Output: Posterior sample $\widehat{\mathbf{X}}_{T-T_0}^{\mathbf{y}}$ approximately drawn from $q_0(\cdot|\mathbf{y})$.

1 **Warm-start:** Set the initial particle $\widehat{\mathbf{X}}_0^{\mathbf{y}} = \widehat{\mathbf{X}}_{T,U}^{\mathbf{y}}$, where $\widehat{\mathbf{X}}_{T,U}^{\mathbf{y}}$ is a warm-start sample.

2 **Reverse diffusion:** Simulate the approximate time-reversal stochastic process over the time interval $[0, T-T_0]$, employing the estimated score function $\widehat{\mathbf{s}}_m^S(t, \mathbf{x}, \mathbf{y})$, to generate $\widehat{\mathbf{X}}_{T-T_0}^{\mathbf{y}}$.

3 **return** $\widehat{\mathbf{X}}_{T-T_0}^{\mathbf{y}}$

Diffusion Posterior Sampling

$$\mathrm{d}\widehat{\mathbf{X}}_t^{\mathbf{y}} = \Big(\widehat{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\widehat{\mathbf{s}}_m^S(T-t, \widehat{\mathbf{X}}_t^{\mathbf{y}}, \mathbf{y})}_{\text{posterior score estimation}}\Big)\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\widehat{\mathbf{X}}_0^{\mathbf{y}} \sim \widehat{q}_T^U(\cdot|\mathbf{y})}_{\text{warm-start}}, \ t \in (0, T-T_0),$$

- Posterior score estimation: $T \in (0, \bar{t})$
- Warm-start: $T \in (\underline{t}, +\infty)$

Diffusion Posterior Sampling

$$\mathrm{d}\widehat{\mathbf{X}}_t^{\mathbf{y}} = (\widehat{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\widehat{\mathbf{s}}_m^S(T-t, \widehat{\mathbf{X}}_t^{\mathbf{y}}, \mathbf{y})}_{\text{posterior score estimation}})\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\widehat{\mathbf{X}}_0^{\mathbf{y}} \sim \widehat{q}_T^U(\cdot|\mathbf{y})}_{\text{warm-start}}, \ t \in (0, T-T_0),$$

- Posterior score estimation: $T \in (0, \bar{t})$
- Warm-start: $T \in (\underline{t}, +\infty)$

Does there exist a terminal time T such that both the posterior score estimation and the terminal posterior sampling procedures are theoretically guaranteed to converge?

Diffusion Posterior Sampling

$$\mathrm{d}\widehat{\mathbf{X}}_t^{\mathbf{y}} = \Big(\widehat{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\widehat{\mathbf{s}}_m^S(T-t, \widehat{\mathbf{X}}_t^{\mathbf{y}}, \mathbf{y})}_{\text{posterior score estimation}}\Big)\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\widehat{\mathbf{X}}_0^{\mathbf{y}} \sim \widehat{q}_T^U(\cdot|\mathbf{y})}_{\text{warm-start}}, \ t \in (0, T-T_0),$$

- Posterior score estimation: $T \in (0, \bar{t})$
- Warm-start: $T \in (\underline{t}, +\infty)$

Does there exist a terminal time T such that both the posterior score estimation and the terminal posterior sampling procedures are theoretically guaranteed to converge?

Suppose **A1** and **A2** hold. If $2\alpha V_{\mathrm{SG}}^2 \leq 1$, then there exists a terminal time T such that:

- the posterior denoising density $p_t(\cdot|\mathbf{x}, \mathbf{y})$ is log-concave for all $t \in (0, T)$; and
- the posterior density $q_T(\cdot|\mathbf{y})$ satisfies LSI.

Outline

Error Decomposition

The posterior time-reversal process

$$\mathrm{d}\bar{\mathbf{X}}_t^{\mathbf{y}} = (\bar{\mathbf{X}}_t^{\mathbf{y}} + 2\nabla_{\mathbf{x}} \log q_{T-t}(\bar{\mathbf{X}}_t^{\mathbf{y}}|\mathbf{y}))\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{X}}_0^{\mathbf{y}} \sim q_T(\cdot|\mathbf{y}),\ t \in (0, T).$$

The approximate time-reversal process

$$\mathrm{d}\widehat{\mathbf{X}}_t^{\mathbf{y}} = (\widehat{\mathbf{X}}_t^{\mathbf{y}} + 2\underbrace{\widehat{\mathbf{s}}_m^S(T-t, \widehat{\mathbf{X}}_t^{\mathbf{y}}, \mathbf{y})}_{\text{posterior score estimation}})\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \underbrace{\widehat{\mathbf{X}}_0^{\mathbf{y}} \sim \widehat{q}_T^U(\cdot|\mathbf{y})}_{\text{warm-start}},\ \underbrace{t \in (0, T-T_0)}_{\text{early-stopping}}.$$

Sources of errors

- posterior score estimation error
- warm-start error
- early-stopping error

Additional Assumptions

- **A3.** For a fixed measurement $\mathbf{y} \in \mathbb{R}^n$, there exist constants $r \geq 1$, $B > 1$ and $L > 1$ such that for each $\mathbf{x}_0 \in \mathbb{R}^d$,
$$\|\nabla \log \pi_0(\mathbf{x}_0)\|_2 \leq B(1 + \|\mathbf{x}_0\|_2^r).$$
- **A4.** For a fixed measurement $\mathbf{y} \in \mathbb{R}^n$, there exists a constant $H > 1$, such that $\|\nabla \log q_0(\mathbf{0}|\mathbf{y})\|_2 \leq H$.
- **A5.** Let $\widehat{\mathbf{s}}_{\mathrm{prior}}$ be an estimator of the prior score function $\nabla \log \pi_0$. There exists $\varepsilon_{\mathrm{prior}} \in (0, 1)$, such that
$$\int \|\nabla \log \pi_0(\mathbf{x}_0) - \widehat{\mathbf{s}}_{\mathrm{prior}}(\mathbf{x}_0)\|_2^2 \pi_0(\mathbf{x}_0) \mathrm{d}\mathbf{x}_0 \leq \varepsilon_{\mathrm{prior}}^2.$$
- **A6.** For a fixed measurement $\mathbf{y} \in \mathbb{R}^n$, there exist constants $G > 1$, such that $\widehat{\mathbf{s}}_{\mathrm{prior}} - \nabla \ell_{\mathbf{y}}$ is G-Lipschitz.

Error Decomposition

Suppose **A2** hold.

- Let $T_0 \in (0, \frac{1}{2})$ and $R^2 = (4\mu_{T_0}^2 V_{\mathrm{SG}}^2 + 16\sigma_{T_0}^2)\log(\kappa_{\mathbf{y}}\varepsilon_{T_0}^{-1})$. For $\varepsilon_{T_0} \in (0,1)$, we have

$$\mathbb{W}_2^2(q_0(\cdot|\mathbf{y}), \mathcal{M}(\mu_{T_0}^{-1})\sharp\widehat{q}_{T-T_0}^{R}(\cdot|\mathbf{y})) \leq C\Big\{\frac{\sigma_{T_0}^2}{\mu_{T_0}^2} + \varepsilon_{T_0}\log\Big(\frac{\kappa_{\mathbf{y}}}{\varepsilon_{T_0}}\Big)\Big\},$$

  provided that $\|q_{T_0}(\cdot|\mathbf{y}) - \widehat{q}_{T-T_0}(\cdot|\mathbf{y})\|_{\mathrm{TV}}^2 = \varepsilon_{T_0}^2$. Here C is a constant only depending on d, V_{SG}, and C_{SG}.
- For $0 < T_0 < T$, it holds that

$$\begin{aligned} &\|q_{T_0}(\cdot|\mathbf{y}) - \widehat{q}_{T-T_0}(\cdot|\mathbf{y})\|_{\mathrm{TV}}^2 \\ &\leq \underbrace{\int_{T_0}^{T} \mathbb{E}\big[\|\nabla \log q_t(\mathbf{X}_t^{\mathbf{y}}|\mathbf{y}) - \widehat{\mathbf{s}}_m^{S}(t, \mathbf{X}_t^{\mathbf{y}}, \mathbf{y})\|_2^2\big]\,\mathrm{d}t}_{\text{posterior score estimation}} + \underbrace{2\|q_T(\cdot|\mathbf{y}) - \widehat{q}_0^{U}(\cdot|\mathbf{y})\|_{\mathrm{TV}}^2}_{\text{warm-start}}. \end{aligned}$$

Error of Posterior Score Estimation

Suppose **A1**, **A2**, **A3**, **A5**, and **A6** hold. Let $0 < T_0 < T < \frac{1}{2}\log(1+\alpha^{-1})$. For each time $t \in (T_0, T)$,

$$
\begin{aligned}
&\mathbb{E}\left[\|\nabla \log q_t(\mathbf{X}_t^{\mathbf{y}}|\mathbf{y}) - \widehat{\mathbf{s}}_m^S(t, \mathbf{X}_t^{\mathbf{y}}, \mathbf{y})\|_2^2\right] \\
&\leq \underbrace{C\frac{\mu_{T_0}^2}{\sigma_{T_0}^4}\frac{\kappa_{\mathbf{y}}}{m}}_{\text{Monte Carlo}} + \underbrace{C\frac{\mu_{T_0}^2}{\sigma_{T_0}^4}\exp\left(-\frac{2(\mu_T^2-\alpha\sigma_T^2)}{\sigma_T^2}S\right)\eta_{\mathbf{y}}^2}_{\text{convergence of Langevin dynamics}} \\
&\quad + \underbrace{C\frac{\mu_{T_0}^2}{\sigma_{T_0}^4}\left(\frac{\sigma_T^2\eta_{\mathbf{y}}^2}{\mu_T^2-\alpha\sigma_T^2}+S\right)^{\frac{1}{2}}\exp\left(2\left(G+\frac{\mu_{T_0}^2}{\sigma_{T_0}^2}\right)S\right)\kappa_{\mathbf{y}}^{\frac{1}{2}}\varepsilon_{\mathrm{prior}}^{\frac{1}{2}}}_{\text{prior score estimation error}},
\end{aligned}
$$

where C is a constant only depending on d, B, V_{SG}, and C_{SG}. Here $\kappa_{\mathbf{y}}$ is the condition number, and $\eta_{\mathbf{y}}$ is the initial discrepancy of Langevin dynamics.

Error of Warm-Start

Suppose **A1**, **A2**, **A3**, **A4**, **A5**, and **A6** hold. Assume $2\alpha V_{\mathrm{SG}}^2 \leq 1$. Let the terminal time satisfy $T \in (\frac{1}{2}\log(1+2V_{\mathrm{SG}}^2), \frac{1}{2}\log(1+\alpha^{-1}))$. Then for $\varepsilon \in (0,1)$, it holds that

$$\|q_T(\cdot|\mathbf{y}) - \widehat{q}_T^U(\cdot|\mathbf{y})\|_{\mathrm{TV}}^2 \leq C\varepsilon,$$

provided that

$$U = \Theta\Big(\exp\Big(2\frac{\sigma_T^2 + 2\mu_T^2 V_{\mathrm{SG}}^2}{\sigma_T^2 - 2\mu_T^2 V_{\mathrm{SG}}^2}\log(\kappa_{\mathbf{y}}^2 C_{\mathrm{SG}}^2)\Big)\log\Big(\frac{\zeta_{\mathbf{y}}^2}{\varepsilon}\Big)\Big),$$

$$\varepsilon_{\mathrm{post}} = \Theta\Big(\exp\Big(-2\frac{\sigma_T^2 + 2\mu_T^2 V_{\mathrm{SG}}^2}{\sigma_T^2 - 2\mu_T^2 V_{\mathrm{SG}}^2}\log(\kappa_{\mathbf{y}}^2 C_{\mathrm{SG}}^2)\Big)\Big(\zeta_{\mathbf{y}}^2 + \log\Big(\frac{\zeta_{\mathbf{y}}^2}{\varepsilon}\Big)\Big)^{-1}(\mu_T^2 - \alpha\sigma_T^2)^3\frac{\varepsilon^2}{\kappa_{\mathbf{y}}^{1/2}}\Big),$$

where C is a constant only depending on d, B, H, V_{SG}, and C_{SG}. The initial divergence $\zeta_{\mathbf{y}}^2$ and the posterior score matching error at the terminal time T are defined, respectively, as:

$$\zeta_{\mathbf{y}}^2 := \chi^2(\widehat{q}_T^0(\cdot|\mathbf{y})\|q_T(\cdot|\mathbf{y})), \quad \varepsilon_{\mathrm{post}}^2 := \mathbb{E}\big[\|\nabla\log q_T(\mathbf{X}_T^{\mathbf{y}}|\mathbf{y}) - \widehat{\mathbf{s}}_m^S(T,\mathbf{X}_T^{\mathbf{y}},\mathbf{y})\|_2^2\big].$$

Convergence of Posterior Sampling

Suppose **A1**, **A2**, **A3**, **A4**, **A5**, and **A6** hold. Assume $2\alpha V_{\mathrm{SG}}^2 \leq 1$. Let the terminal time satisfy $T \in (\frac{1}{2}\log(1+2V_{\mathrm{SG}}^2), \frac{1}{2}\log(1+\alpha^{-1}))$. For each time $t \in (T_0, T)$, Then for $\varepsilon \in (0,1)$, the following inequality holds:

$$\begin{aligned} &\mathbb{E}\left[\mathbb{W}_2^2(q_0(\cdot|\mathbf{y}), \mathcal{M}(\mu_{T_0}^{-1})\sharp\widehat{q}_{T-T_0}^R(\cdot|\mathbf{y}))\right] \\ &\leq C\left\{\frac{\sigma_{T_0}^2}{\mu_{T_0}^2} + \left(\frac{\mu_{T_0}}{\sigma_{T_0}^2}\varepsilon + \varepsilon^{\frac{1}{2}}\right)\log\left(\kappa_{\mathbf{y}}\left(\frac{\mu_{T_0}}{\sigma_{T_0}^2}\varepsilon + \varepsilon^{\frac{1}{2}}\right)^{-1}\right)\right\}, \end{aligned}$$

where C is a constant only depending on d, $\kappa_{\mathbf{y}}$, B, H, V_{SG}, and C_{SG}. Furthermore, by setting $T_0 = \sqrt{\varepsilon_0}$, we have

$$\mathbb{E}\left[\mathbb{W}_2^2(q_0(\cdot|\mathbf{y}), \mathcal{M}(\mu_{T_0}^{-1})\sharp\widehat{q}_{T-T_0}^R(\cdot|\mathbf{y}))\right] \leq C'\varepsilon^{\frac{1}{2}}\log\varepsilon^{-1}.$$

Outline

Numerical Experiemnts

Dataset: `FFHQ64`

Motion deblurring

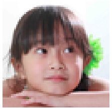

face 1

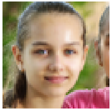

face 2

Gaussian deblurring

face 3

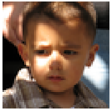

face 4

Nonlinear deblurring

face 5

face 6

	Gaussian Deblur			Motion Deblur			Nonlinear Deblur		
	TV	DPS	**Ours**	TV	DPS	**Ours**	TV	DPS	**Ours**
PSNR	23.95	24.15	**26.42**	24.65	26.66	**28.86**	19.70	20.93	**28.44**
SSIM	0.81	0.81	**0.87**	0.80	0.88	**0.92**	0.53	0.68	**0.91**

Gaussian Deblurring

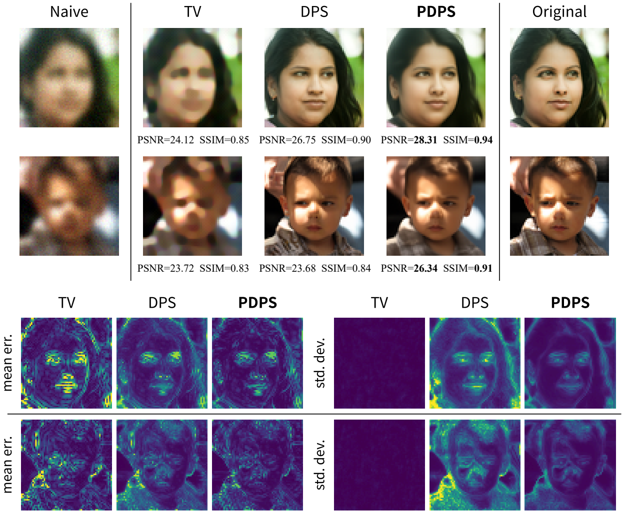

Motion Deblurring

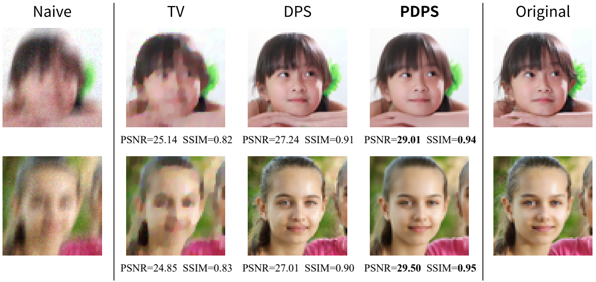

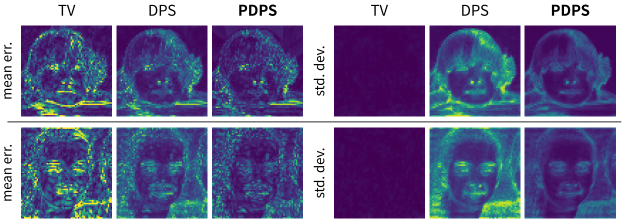

Nonlinear Deblurring

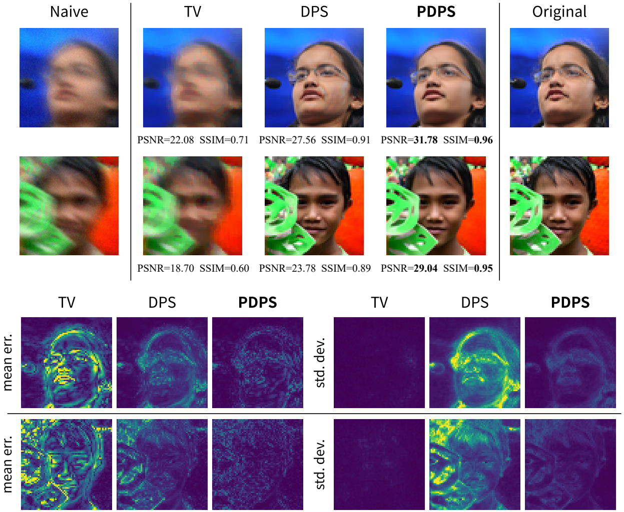

Cross-Dataset Generalization Assessment

- Prior score is learning using FFHQ64.
- Groundtruth image is from AFHQ64.

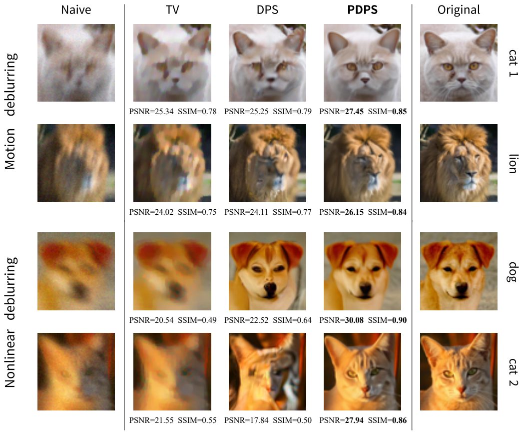

Ablation Study of Terminal Time

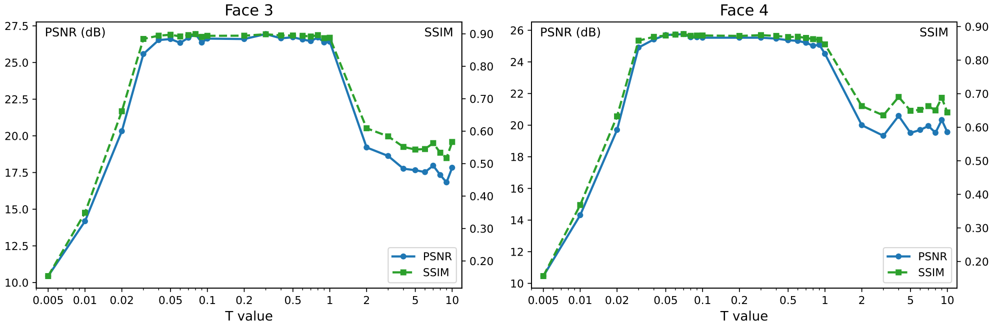

- For too small T, the terminal posterior density $q_T(\cdot|\mathbf{y})$ does not satisfies LSI.
- For too large T, the posterior denoising density $p_t(\cdot|\mathbf{x}, \mathbf{y})$ is non-log-concave for t close to T.

Concluding Remarks

Summary

- A diffusion-based posterior sampling method without heuristic approximations.
- Non-asymptotic error analysis for posterior sampling.
- Experiemntal performance in image restoration.

Concluding Remarks

Summary

- A diffusion-based posterior sampling method without heuristic approximations.
- Non-asymptotic error analysis for posterior sampling.
- Experiemntal performance in image restoration.

Further work

- Derivative-free Bayesian inference.
- Blind inverse problems.
- Inverse problem in inifinite-dimensional problems.

Reference

Jinyuan Chang, Chenguang Duan, Yuling Jiao, Ruoxuan Li, Jerry Zhijian Yang, and Cheng Yuan. Provable Diffusion Posterior Sampling for Bayesian Inversion. arXiv:2512.08022

Thanks for your attention!

Email: duan@igpm.rwth-aachen.de
Homepage:
https://chenguangduan.github.io/
Google Scholar:
https://scholar.google.com/citations?user=RpmGgyMAAAAJ