UPPSALA
UNIVERSITET

40 Years ICML
International Conference
On Machine Learning

Image Restoration with Mean-Reverting Stochastic Differential Equations

Ziwei Luo, Fredrik K. Gustafsson, Zheng Zhao,
Jens Sjölund, Thomas B. Schön

Uppsala University

Project Page

Our Contributions

UPPSALA UNIVERSITET

Mean-Reverting SDE for image restoration

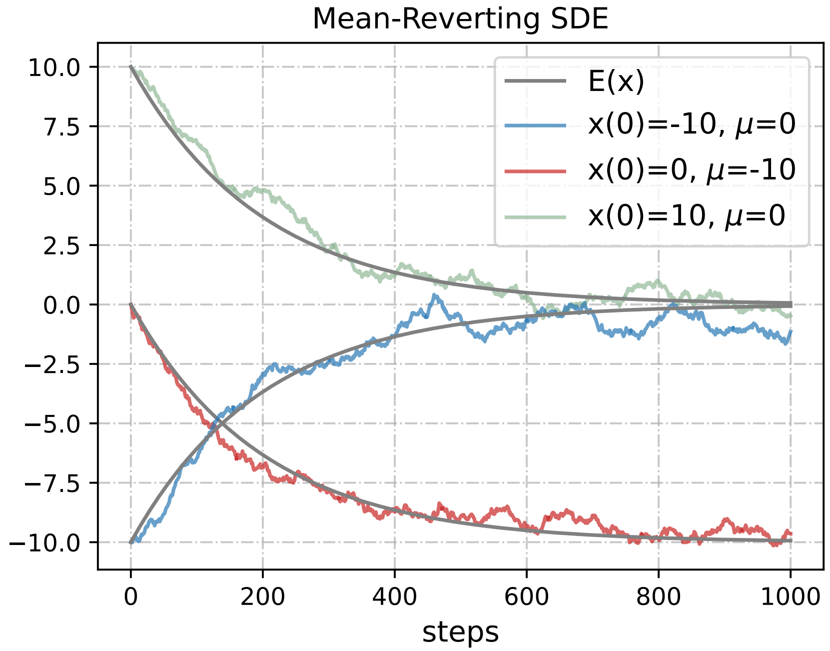

Our Contributions

UPPSALA UNIVERSITET

Mean-Reverting SDE for image restoration

- Tractable *mean* and *variance* → **score**

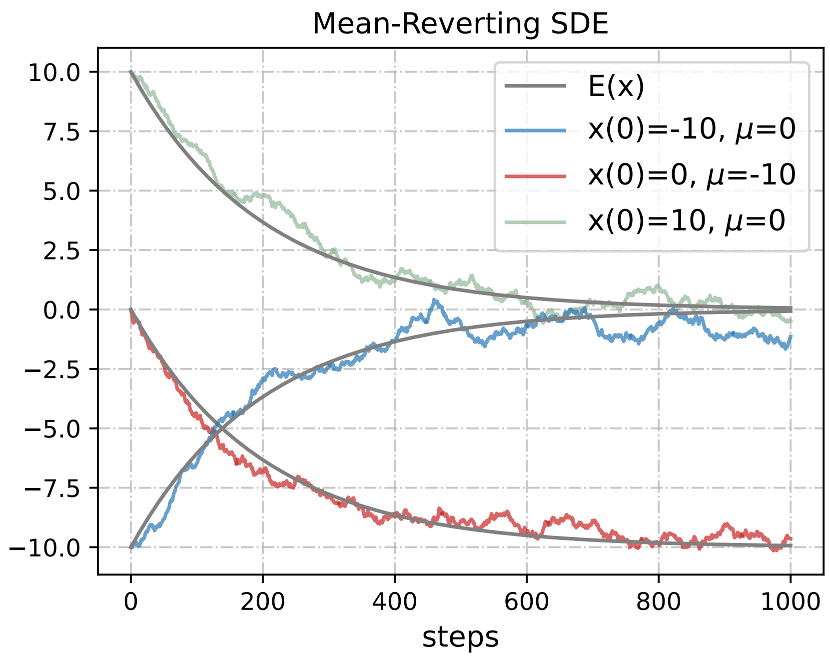

Our Contributions

UPPSALA UNIVERSITET

Mean-Reverting SDE for image restoration

- Tractable *mean* and *variance* → **score**
- Naturally simulate degradation → **no-prior**

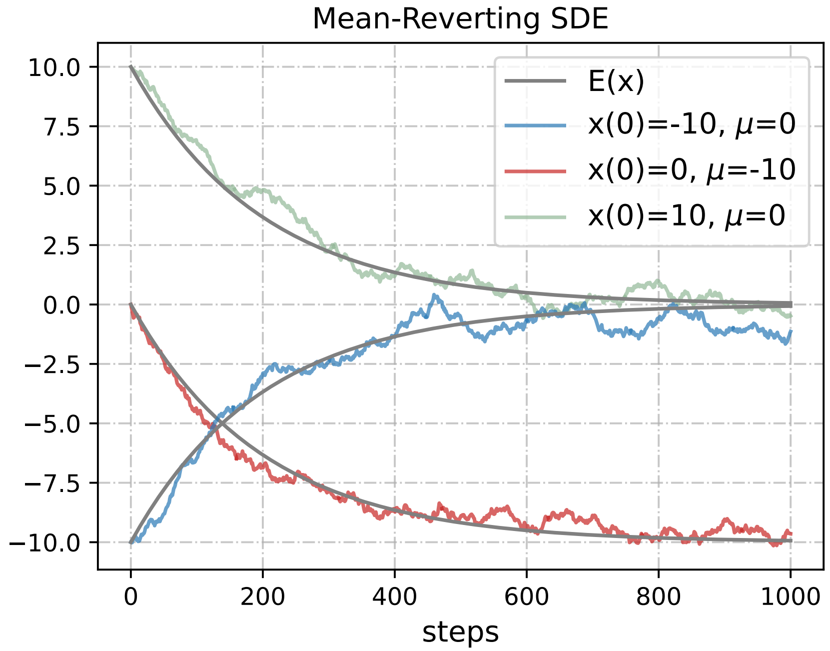

Our Contributions

UPPSALA UNIVERSITET

Mean-Reverting SDE for image restoration

- Tractable *mean* and *variance* → **score**
- Naturally simulate degradation → **no-prior**
- Maximum likelihood learning → **stable training**

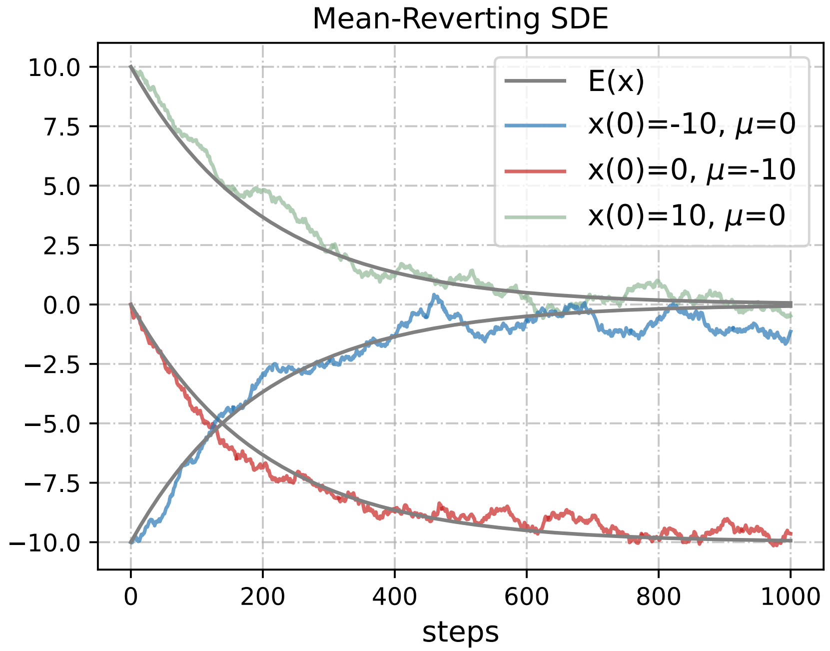

Our Contributions

UPPSALA UNIVERSITET

Mean-Reverting SDE for image restoration

- Tractable *mean* and *variance* → **score**
- Naturally simulate degradation → **no-prior**
- Maximum likelihood learning → **stable training**

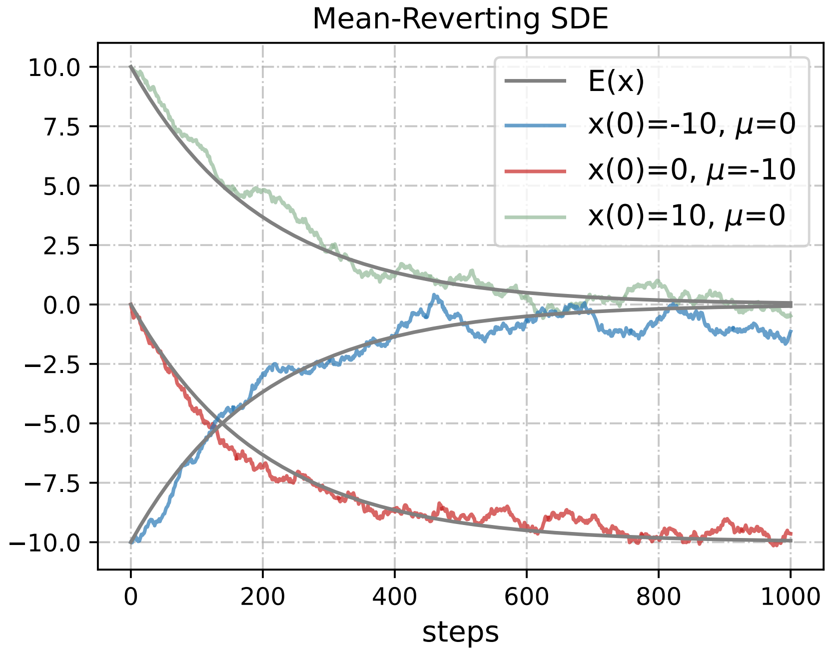

Our approach (**IR-SDE**) achieves highly competitive performance on various tasks.

Forward SDE for Image Degradation

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t \left(\mu - x\right) \mathrm{d}t + \sigma_t \, \mathrm{d}w$

Forward SDE for Image Degradation

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

- We can mathematically prove that this SDE has a closed-form solution when θ_t and σ_t satisfy $\frac{{\sigma_t}^2}{2\theta_t} = \lambda^2$ (stationary variance):

$$p(x(t) \mid x(s)) = \mathcal{N}\big(x(t) \mid m_{s:t}(x(s)), v_{s:t}\big)$$

$$m_{s:t}(x_s) := \mu + (x(s) - \mu)\,\mathrm{e}^{-\bar{\theta}_{s:t}},$$

$$\begin{aligned} v_{s:t} &:= \int_s^t \sigma_z^2\,\mathrm{e}^{-2\bar{\theta}_{z:t}}\,\mathrm{d}z \\ &= \lambda^2\left(1 - \mathrm{e}^{-2\,\bar{\theta}_{s:t}}\right), \quad \text{where } \bar{\theta}_{s:t} := \textstyle\int_s^t \theta_z\,\mathrm{d}z \end{aligned}$$

Proof:

- Itô's formula

Reverse-Time SDE for Image Restoration

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

Reverse-time SDE: $\mathrm{d}x = \big[\theta_t\,(\mu - x) - \sigma_t^2\,\nabla_x \log p_t(x)\big]\,\mathrm{d}t + \sigma_t\,\mathrm{d}\hat{w}$

Reverse-Time SDE for Image Restoration

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

Reverse-time SDE: $\mathrm{d}x = \left[\theta_t\,(\mu - x) - \sigma_t^2 \boxed{\nabla_x \log p_t(x)}\right] \mathrm{d}t + \sigma_t\,\mathrm{d}\hat{w}$

Score

In training: $\nabla_x \log p_t(x \mid x(0)) = -\dfrac{x(t) - m_t(x)}{v_t}$

Reverse-Time SDE for Image Restoration

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

Reverse-time SDE: $\mathrm{d}x = \left[\theta_t\,(\mu - x) - \sigma_t^2\,\boxed{\nabla_x \log p_t(x)}\right]\mathrm{d}t + \sigma_t\,\mathrm{d}\hat{w}$

Score

In training: $\nabla_x \log p_t(x \mid x(0)) = -\dfrac{x(t) - m_t(x)}{v_t}$

Noise Matching Loss: $L_\gamma(\phi) := \sum_{i=1}^{T} \gamma_i\,\mathbb{E}\Big[\big\|\tilde{\epsilon}_\phi(x_i, \mu, i) - \epsilon_i\big\|\Big]$

Framework Overview

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t \left(\mu - x\right) \mathrm{d}t + \sigma_t \, \mathrm{d}w$

Framework Overview

Forward SDE: $\mathrm{d}x = \theta_t \left(\mu - x\right) \mathrm{d}t + \sigma_t \, \mathrm{d}w$

$x(0)$

Framework Overview

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

$x(0)$ —— Image Degradation SDE ——→ $x(T)$

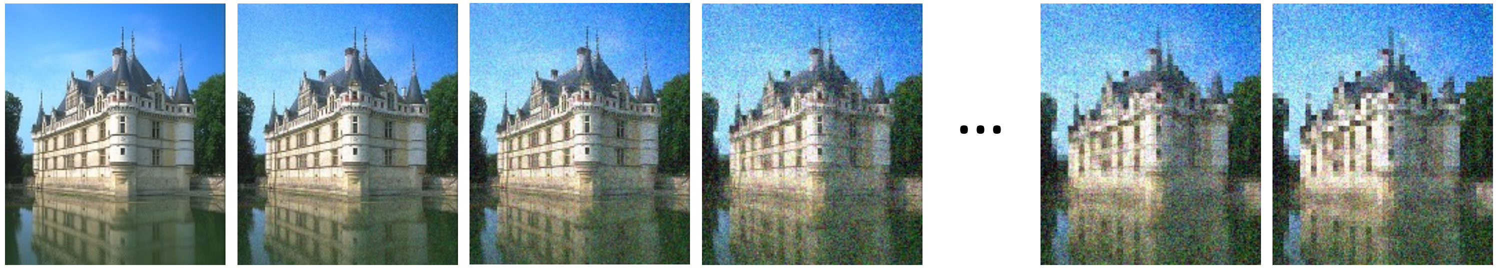

Framework Overview

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t \left(\mu - x\right) \mathrm{d}t + \sigma_t \, \mathrm{d}w$

$x(0)$ —— Image Degradation SDE ——→ $x(T)$ $\quad$ $x(T) \approx \mu + \epsilon$

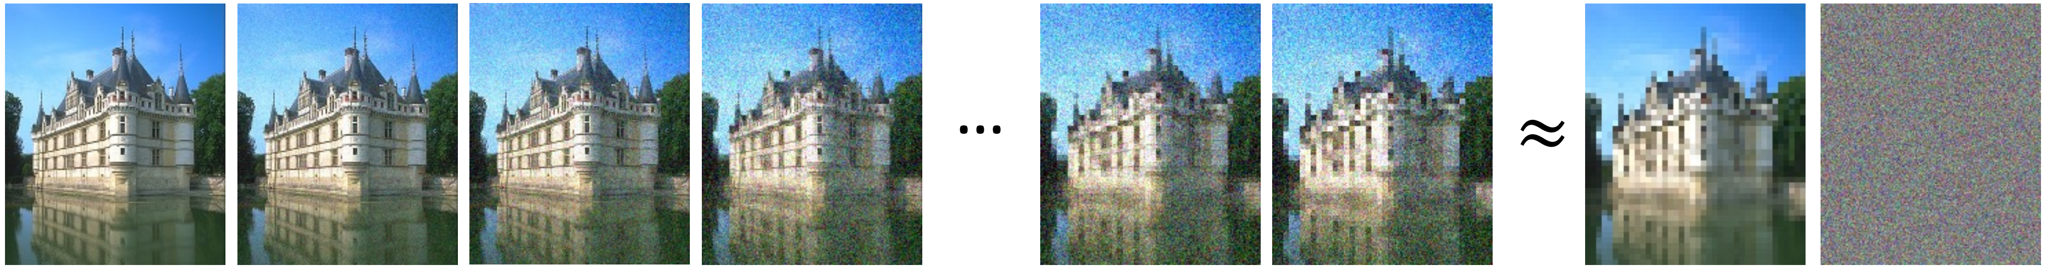

Framework Overview

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

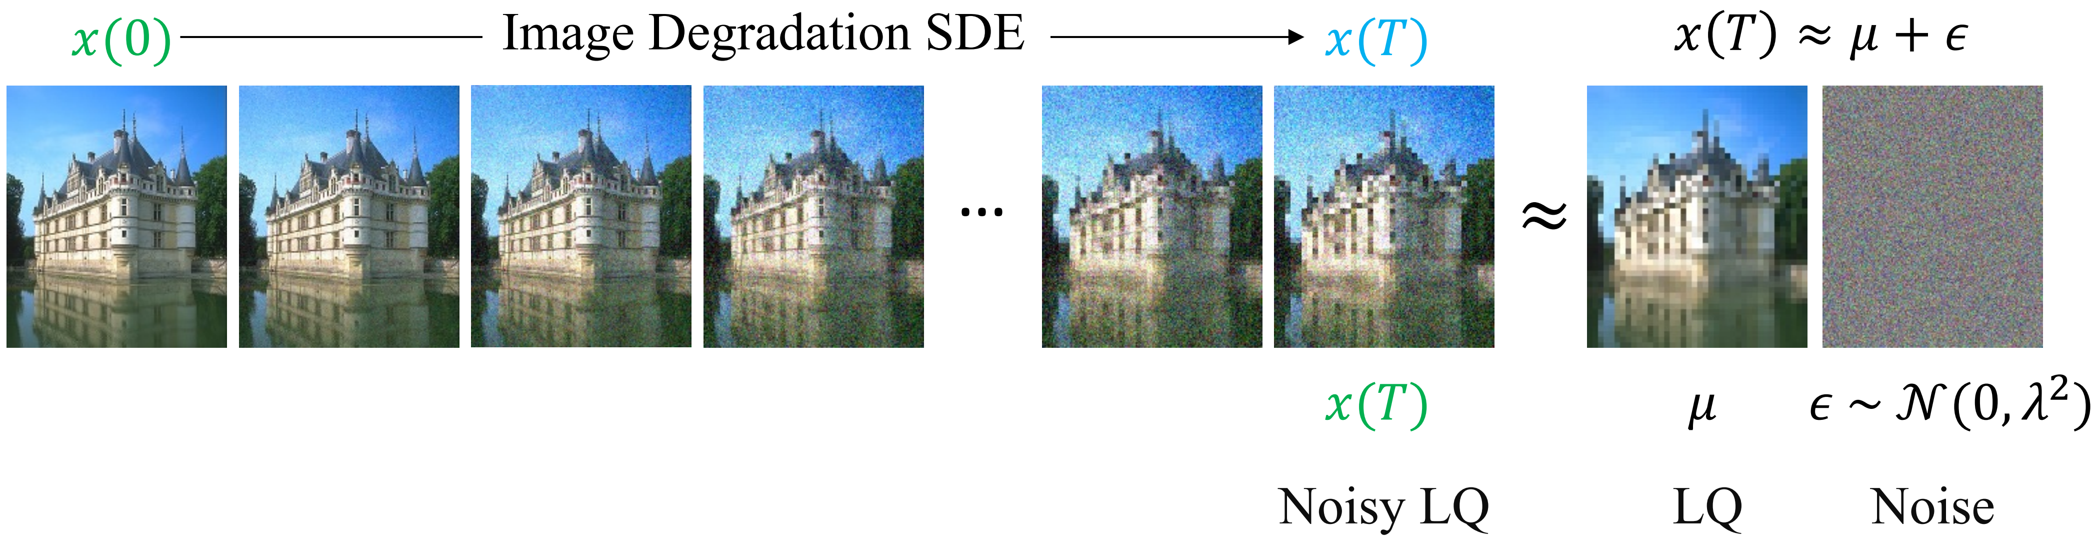

Framework Overview

UPPSALA UNIVERSITET

Forward SDE: $\mathrm{d}x = \theta_t\,(\mu - x)\,\mathrm{d}t + \sigma_t\,\mathrm{d}w$

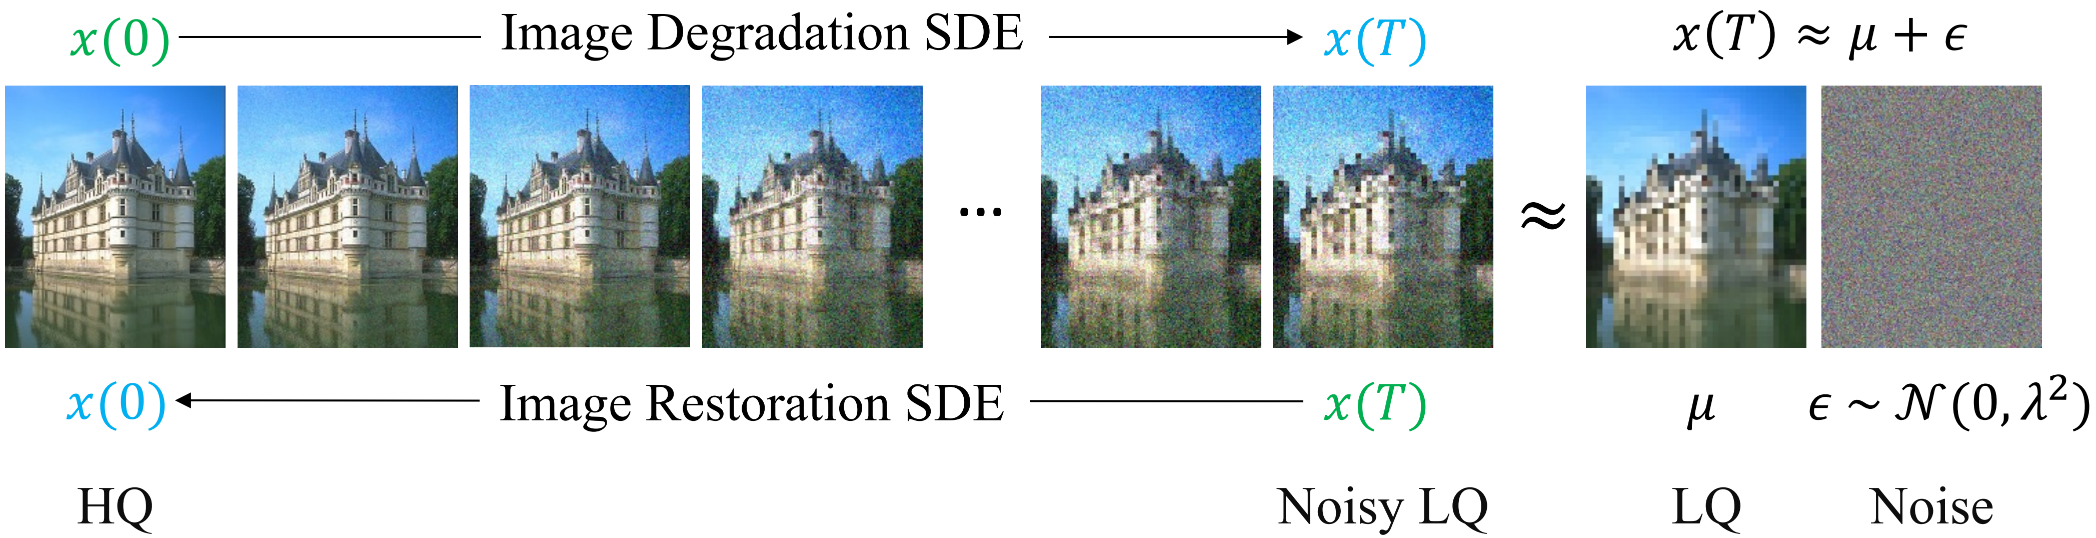

Reverse-time SDE: $\mathrm{d}x = [\theta_t\,(\mu - x) - \sigma_t^2\,\nabla_x \log p_t(x)\,]\mathrm{d}t + \sigma_t\,\mathrm{d}\hat{w}$

Reverse-Time Restoration Process

Maximum Likelihood Loss

UPPSALA
UNIVERSITET

Consider the following reverse process

$$p(x_{1:T} \mid x_0) = p(x_T \mid x_0) \prod_{i=2}^{T} p(x_{i-1} \mid x_i, x_0)$$

Maximum Likelihood Loss

UPPSALA
UNIVERSITET

Consider the following reverse process

$$p(x_{1:T} \mid x_0) = p(x_T \mid x_0) \prod_{i=2}^{T} p(x_{i-1} \mid x_i, x_0)$$

By minimizing the negative log-likelihood, we can get the optimal reverse state

$$x_{i-1}^* = \arg\min_{x_{i-1}} \Big[-\log p\big(x_{i-1} \mid x_i, x_0\big) \Big]$$

$$\boxed{\begin{aligned} &= \frac{1 - \mathrm{e}^{-2\,\bar{\theta}_{i-1}}}{1 - \mathrm{e}^{-2\,\bar{\theta}_i}} \mathrm{e}^{-\theta_i^{\prime}} (x_i - \mu) \\ &\quad + \frac{1 - \mathrm{e}^{-2\,\theta_i^{\prime}}}{1 - \mathrm{e}^{-2\,\bar{\theta}_i}} \mathrm{e}^{-\bar{\theta}_{i-1}} (x_0 - \mu) + \mu. \end{aligned}}$$

Maximum Likelihood Loss

Maximum likelihood loss function:

$$J_\gamma(\phi) := \sum_{i=1}^{T} \gamma_i \, \mathbb{E}\left[\left\| \underbrace{x_i - (\mathrm{d}x_i)_{\tilde{\epsilon}_\phi}}_{\text{reversed } x_{i-1}} - x_{i-1}^* \right\| \right]$$

Maximum Likelihood Loss

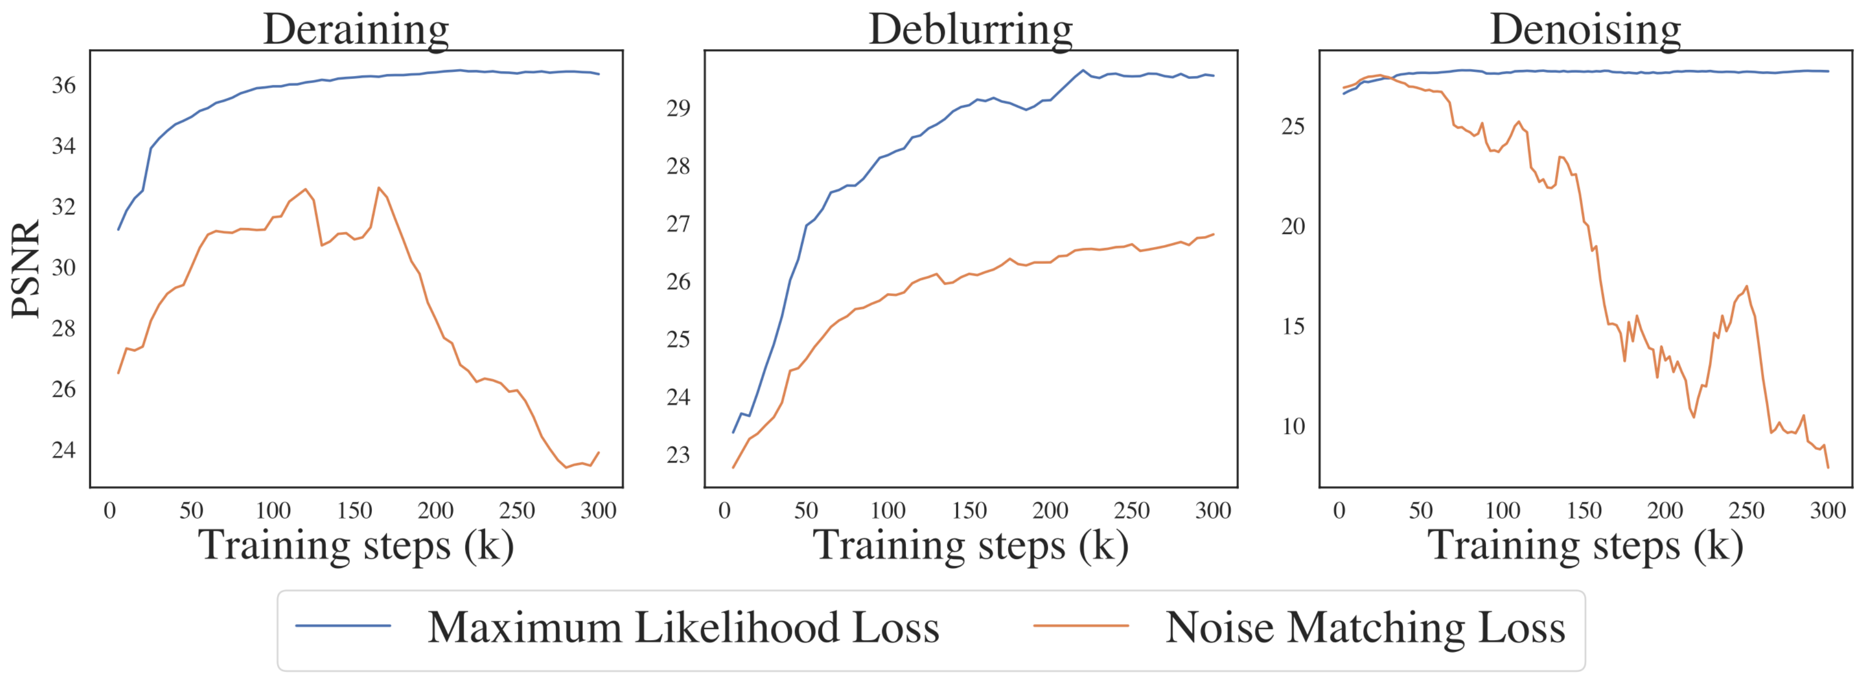

Time-Varying θ Schedules

UPPSALA
UNIVERSITET

$$\mathrm{d}x = \boxed{\theta_t}(\mu - x)\,\mathrm{d}t + \boxed{\sigma_t}\mathrm{d}w$$

Important!

Time-Varying θ Schedules

UPPSALA UNIVERSITET

$$\mathrm{d}x = \boxed{\theta_t}(\mu - x)\,\mathrm{d}t + \boxed{\sigma_t}\,\mathrm{d}w$$

Important!

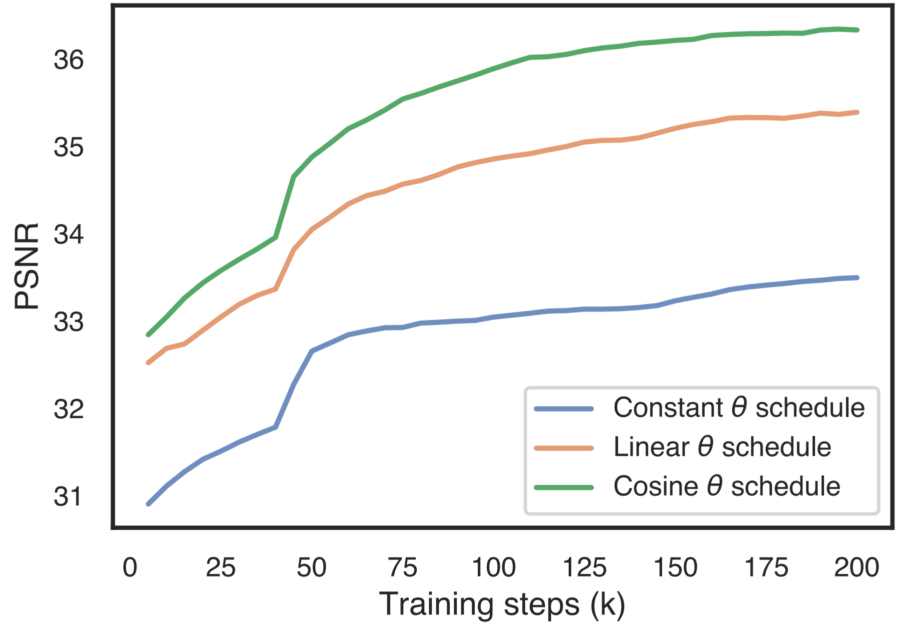

Training curves of three different time-varying θ schedules.

Experiments

- Deraining
- Deblurring
- Denoising
- Super-Resolution
- Inpainting
- Dehazing

Deraining

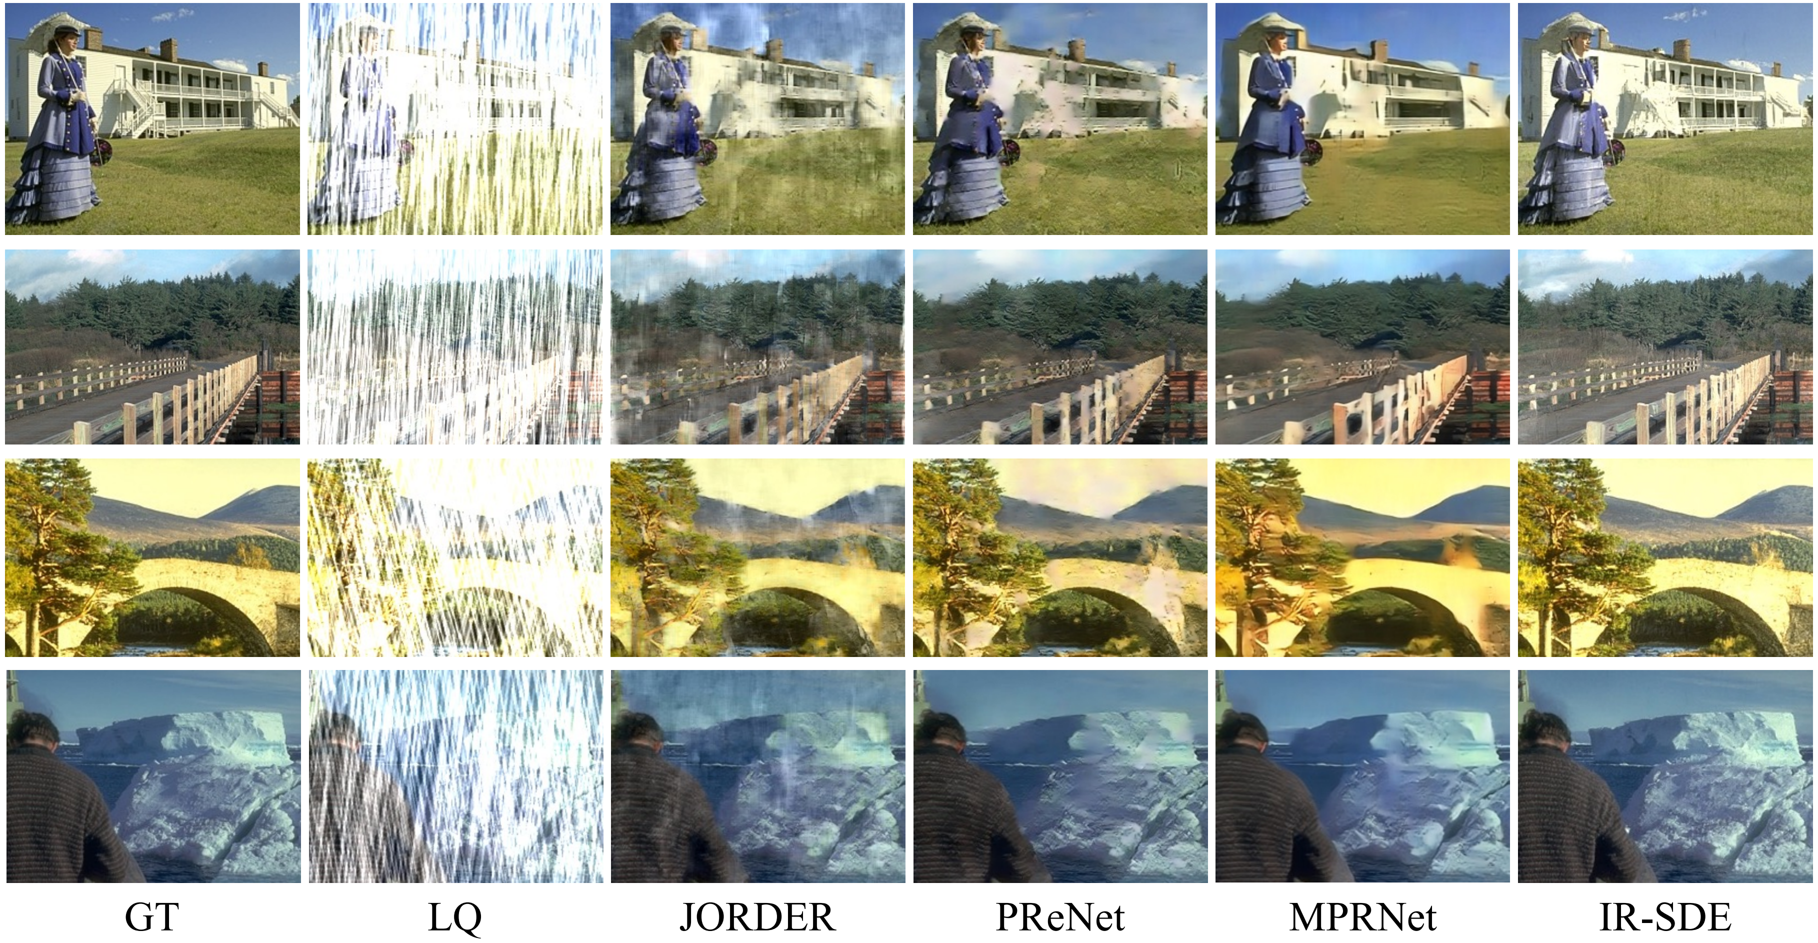

Deraining

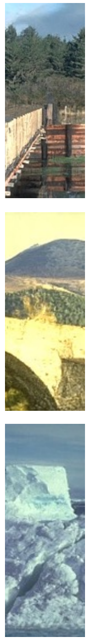

METHOD	DISTORTION		PERCEPTUAL	
	PSNR↑	SSIM↑	LPIPS↓	FID↓
JORDER	26.25	0.8349	0.197	94.58
PRENET	29.46	0.8990	0.128	52.67
MPRNET	30.41	0.8906	0.158	61.59
MAXIM	30.81	0.9027	0.133	58.72
CNN-BASELINE	29.12	0.8824	0.153	57.55
IR-SDE	**31.65**	**0.9041**	**0.047**	**18.64**

Quantitative comparison between the proposed IR-SDE with other image deraining approaches on the Rain100H test set.

Deblurring

GT LQ DeepDeblur DeblurGAN-v2 IR-SDE

Denoising

GT LQ DnCNN FFDNet IR-SDE Denoising-ODE

Super-Resolution

GT Bicubic EDSR IR-SDE

Inpainting

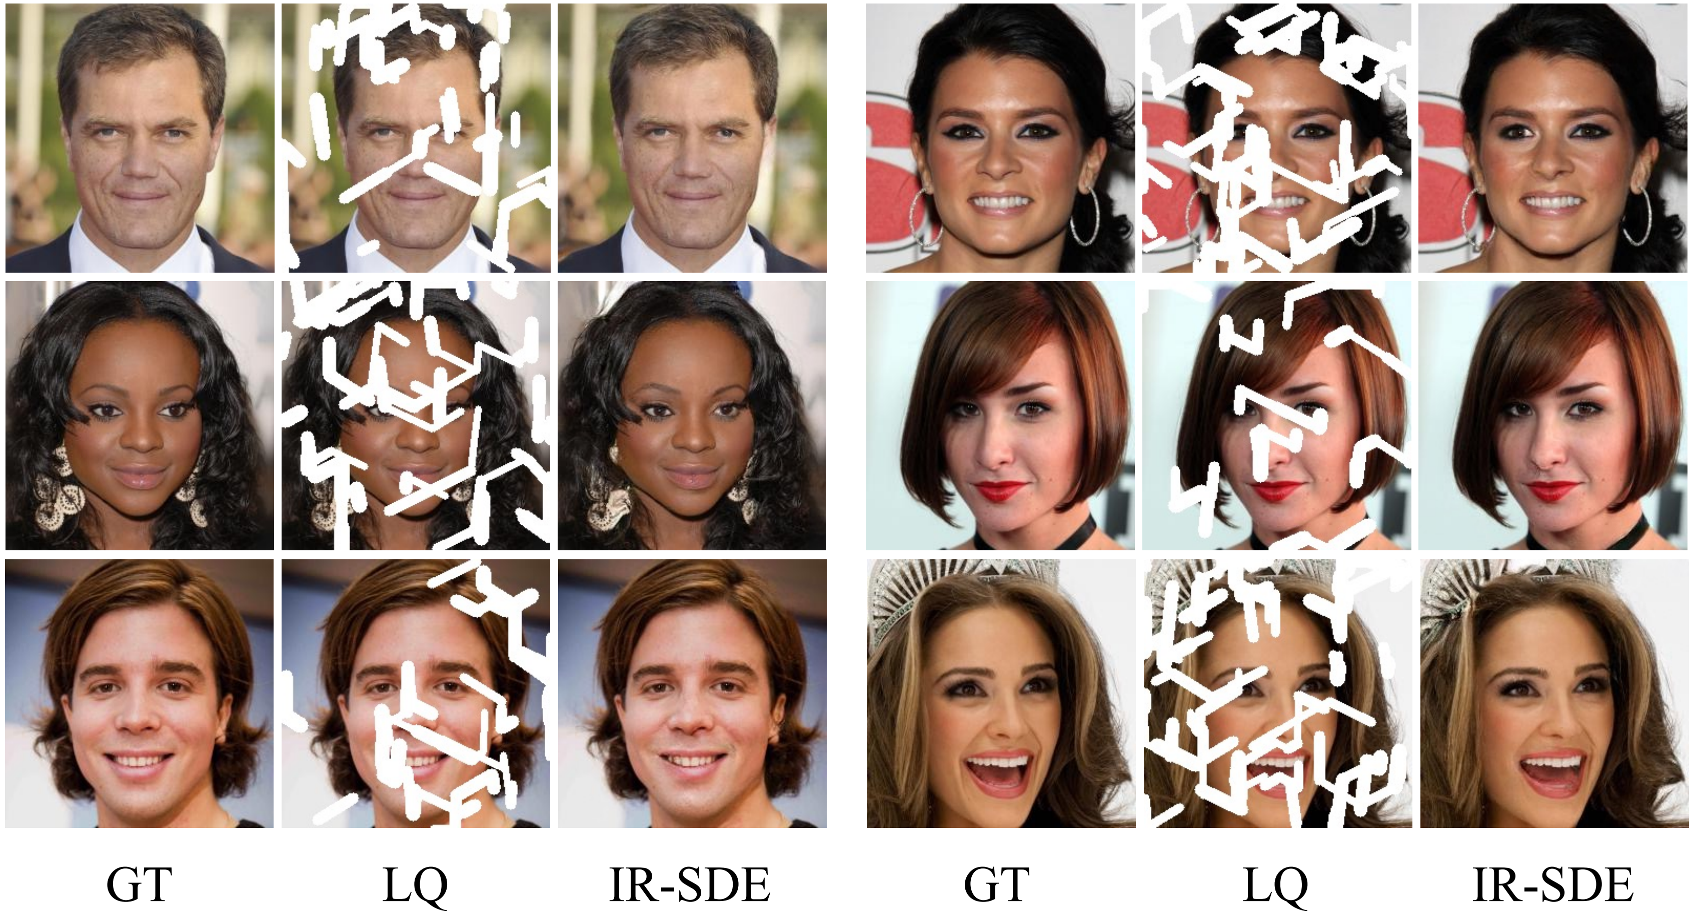

Dehazing

UPPSALA UNIVERSITET

Summary

UPPSALA
UNIVERSITET

Summary

UPPSALA UNIVERSITET

- We proposed a mean-reverting SDE for general-purpose image restoration.

Summary

UPPSALA UNIVERSITET

- We proposed a mean-reverting SDE for general-purpose image restoration.
- We designed a maximum likelihood loss to stabilize training and improve results.

Summary

UPPSALA
UNIVERSITET

- We proposed a mean-reverting SDE for general-purpose image restoration.
- We designed a maximum likelihood loss to stabilize training and improve results.
- Our approach achieves highly competitive performance on diverse image restoration tasks.

Thanks

Project Page