Solving Inverse Problems in Imaging with Diffusion Models

Matthew M. Sato
Stanford University
EE 367 Final Project, Winter 2025

Motivation

- Captured images are noisy
- Diffusion models are great at *unconditionally* generating images
- For inverse problems, need to **condition diffusion process on measurement**
- Conditioning is computationally intractable, requiring heuristics/approximations

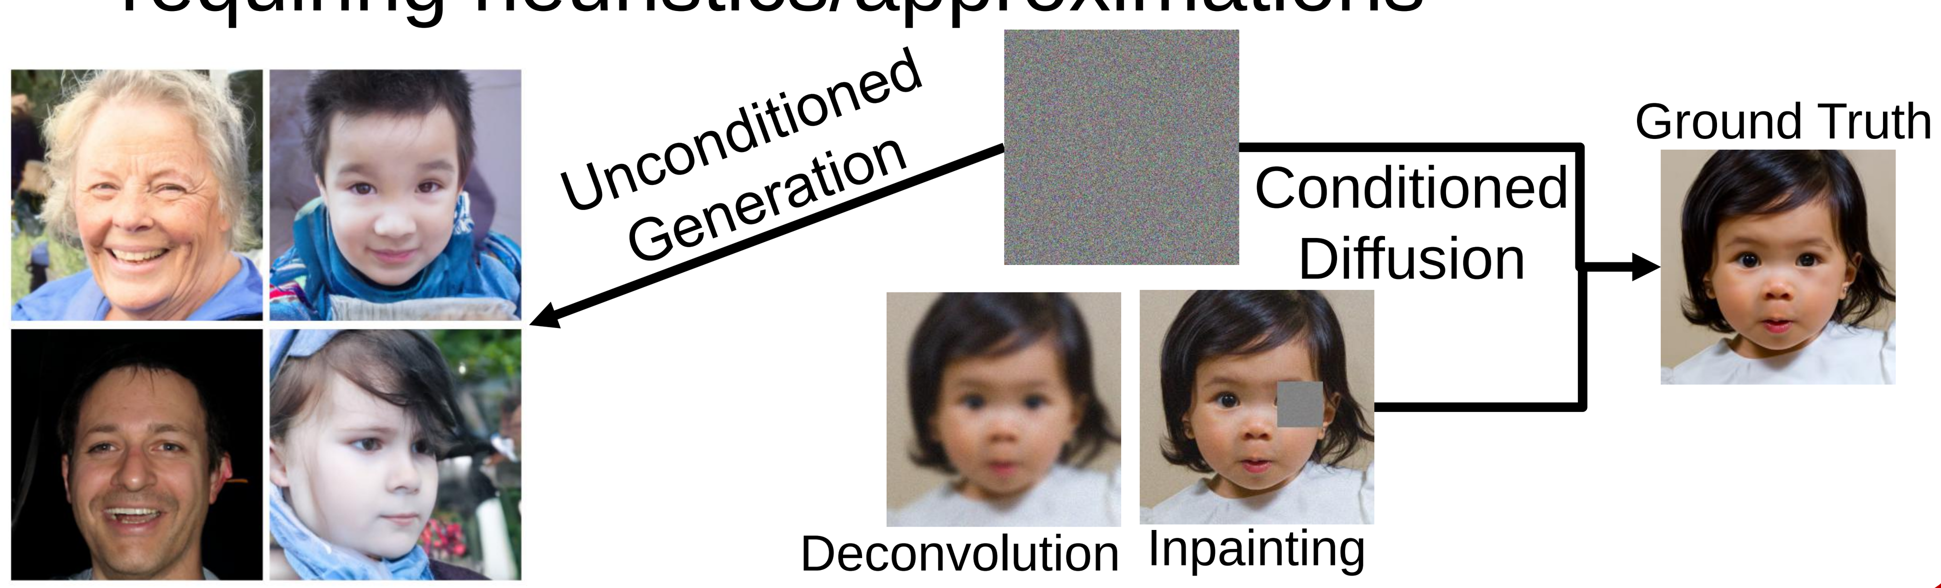

Related Work

- Optimization based methods: ADMM & HQS [1]
  - Convergence can be slow
- Neural network supervised approaches [2]
  - Fails when out of distribution

References

[1] Boyd et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," Foundation and Trends in Machine Learning, 2001
[2] Zhang et al., "Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising," IEEE Trans. Imag. Proc., 2017.
[3] Meng et al., "SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations," ICLR 2022.
[4] Choi et al., "ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models," ICCV 2021.
[5] Jalal et al., "Robust Compressed Sensing MRI with Deep Generative Priors," NeurIps 2021.
[6] Chung et al., "Diffusion Posterior Sampling for General Noisy Inverse Problems," ICLR 2023.

Methods

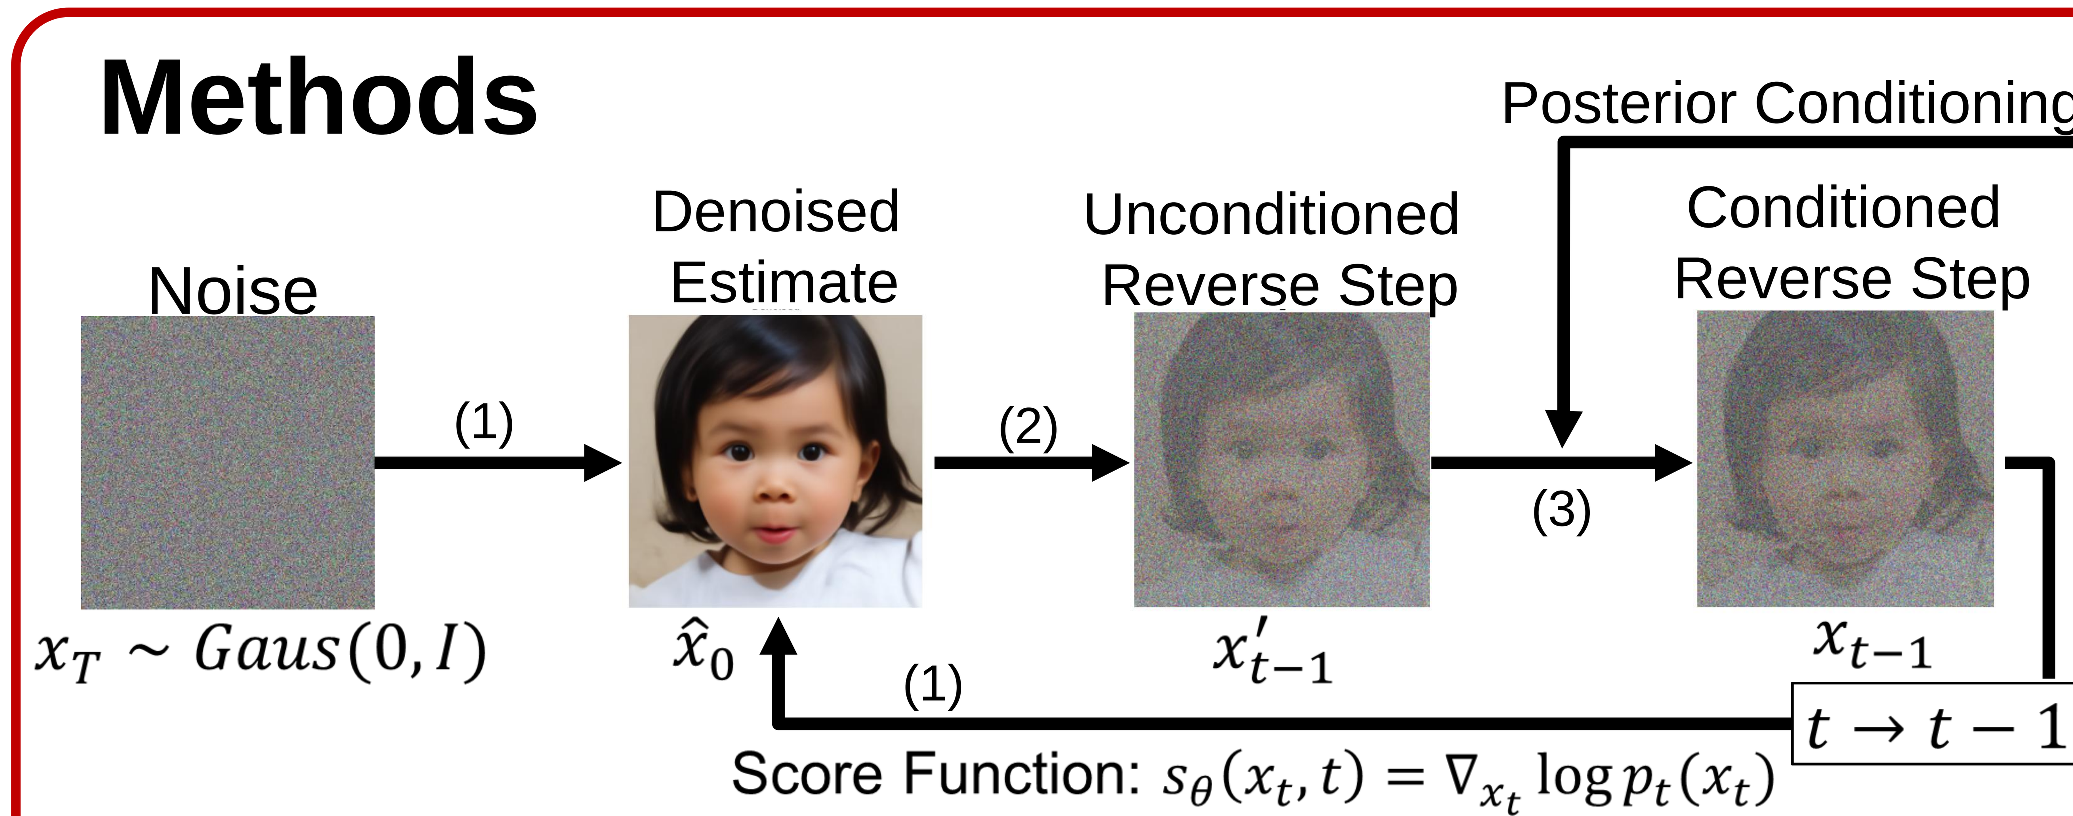

Unconditioned Reverse Diffusion

$$\nabla_{x_t} \log p_t(x_t)$$

Conditioned Reverse Diffusion

$$\nabla_{x_t} \log p_t(x_t|y) = \nabla_{x_t} \log p_t(x_t) + \nabla_{x_t} \log p_t(y|x_t)$$

Posterior Correction Term

$$\zeta_t g(x_t, y) \approx \nabla_{x_t} \log p_t(y|x_t)$$

(1) $\hat{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t + (1 - \bar{\alpha}_t) s_\theta(x_t, t))$

(2) $x'_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t} x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1-\alpha_t)}{1-\bar{\alpha}_t} \hat{x}_0 + \sigma z$

(3) $x_{t-1} = x'_{t-1} + \zeta_t g(x_t, y)$

SDEdit [3]	ILVR [4]	ScoreALD [5]	DPS [6]
$\zeta_t = 0$ $\quad g(x_t, y) = 0$	$\zeta_t = 1$	$\zeta_t = \frac{1}{\sigma^2 + \gamma_t^2}$	$\zeta_t = \frac{\zeta}{\|y - A(\hat{x}_0)\|_2}$
$x_T = \sqrt{\bar{\alpha}_T} y + \sqrt{1 - \bar{\alpha}_T} z$	$g(x_t, y) = \phi_N(y_{t-1}) - \phi_N(x'_{t-1})$	$g(x_t, y) = -\nabla_{x_t} \|y - A(x_t)\|_2^2$	$g(x_t, y) = -\nabla_{x_t} \|y - A(\hat{x}_0)\|_2^2$
$z \sim Gaus(0, I)$	$\phi_N(\cdot)$ = Downsample + Upsample by factor N	γ_t = Annealing Schedule	

Experimental Results

PSNR/LPIPS