Solving Inverse Problems Using Diffusion Models as a Prior

Daniel Song, Alex Liu
Computer Science, Stanford University

Motivation

- How can we generate a set of "good" images?
- Can we use this technique to recover the original, clean image from a corrupted (blurred, noisy, incomplete, etc.) image?

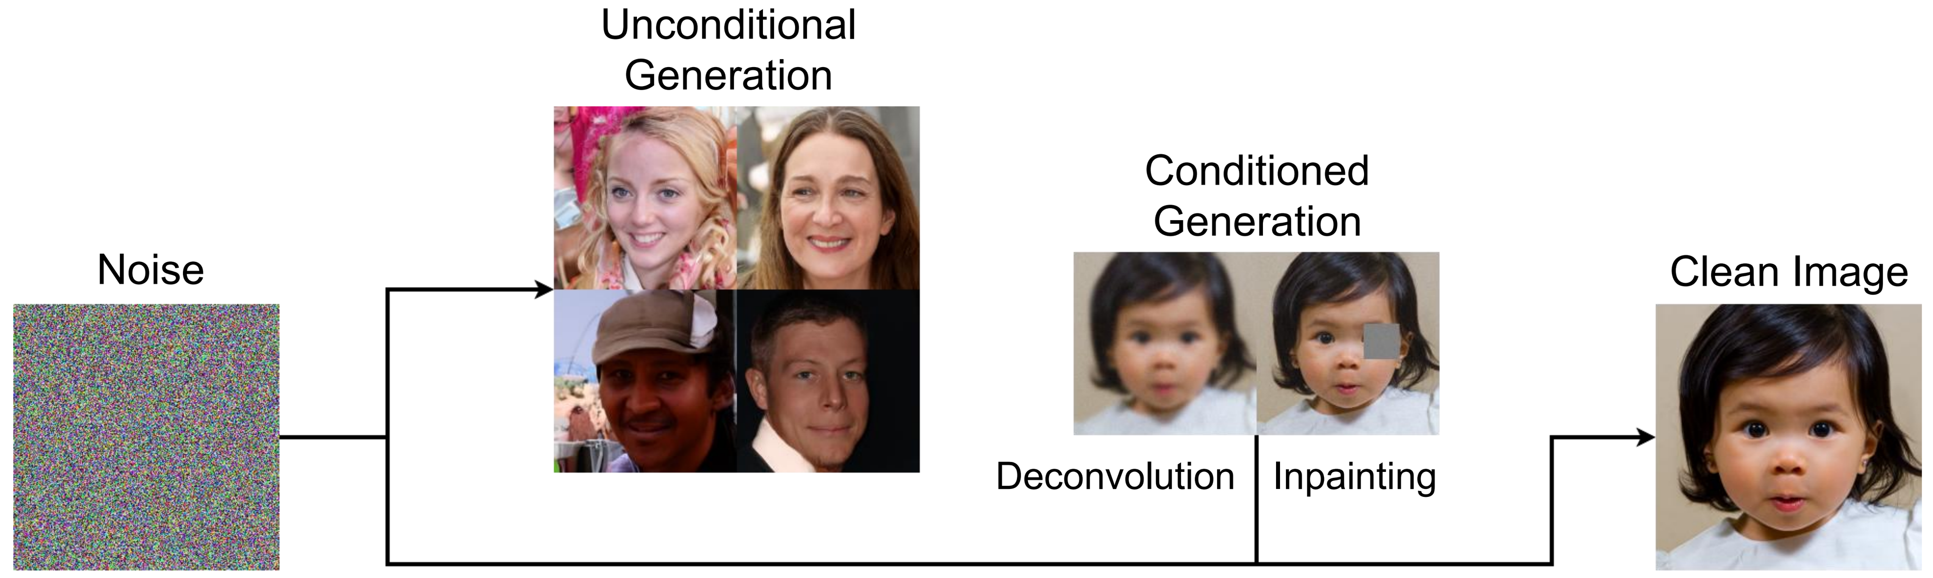

Related Work

- For generating images: **VAE**, **GAN**
  - Generates blurry images, loss of fine details (VAE)
  - Training instability (GAN)
- For inverse problems: **HQS**, **ADMM** (Optimization with regularization)
  - Gives a "single solution"
  - Slow convergence

References

[1] H. Chung, J. Kim, M. Mccann, M. Klasky, J. Ye, "Diffusion Posterior Sampling for General Noisy Inverse Problems", ICLR 2023
[2] J. Ho, A. Jain, P. Abbeel, "Denoising Diffusion Probabilistic Models", NeurIPS 2020
[3] A. Jalal, M. Arvinte, G. Daras, E. Price, A. Dimakis, J. Tamir, "Robust Compressed Sensing MRI with Deep Generative Priors", NeurIPS 2021
[4] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.Y. Zhu, S. Ermon, "SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations", ICLR 2022
[5] S. Boyd, N. Parikh, E. chu, B. Peleato, J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," Foundation and Trends in Machine Learning, 2001

Method

$$d\boldsymbol{x} = f(\boldsymbol{x}, t)dt + g(t)dw$$

$\boldsymbol{x}_0 \longrightarrow \boldsymbol{x}_t$

$\boldsymbol{x}_0 \longleftarrow \boldsymbol{x}_t$

$$d\boldsymbol{x} = [f(\boldsymbol{x}, t) - g^2(t)\nabla_x \log p_t(\boldsymbol{x})]dt + g(t)d\tilde{\boldsymbol{w}}$$

$$p_t(\boldsymbol{x} \mid \boldsymbol{b}) \propto p_t(\boldsymbol{b} \mid \boldsymbol{x})p_t(\boldsymbol{x})$$

$$\nabla_x \log p_t(\boldsymbol{x}_t \mid \boldsymbol{b}) = \nabla_x \log p_t(\boldsymbol{x}_t) + \nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_t)$$

However, ... $\nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_t) \neq \nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_0)$

Approximating $\nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_t)$

SDEdit

Apply partial forward diffusion,

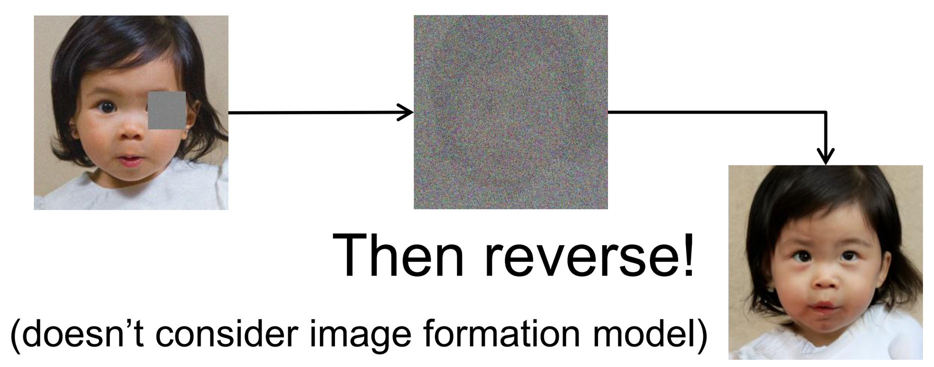

ScoreALD

$$\nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_t) \approx \nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_0)$$

$$\approx -\frac{1}{\sigma^2 + \gamma_t^2}(\boldsymbol{A}^T(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}))$$

$$\boldsymbol{x}_{t-1} = \boldsymbol{x}_{t-1} - \frac{1}{\sigma^2 + \gamma_t^2}\nabla_{x_t}\|\mathcal{A}(\boldsymbol{x}_t) - \boldsymbol{y}\|^2$$

DPS

$$\nabla_x \log p_t(\boldsymbol{b} \mid \boldsymbol{x}_t) \approx \nabla_x \log p_t(\boldsymbol{b} \mid \hat{\boldsymbol{x}}_0)$$

$$\hat{\boldsymbol{x}}_0 = \mathbb{E}[\boldsymbol{x}_0 \mid \boldsymbol{x}_t]$$

$$= \frac{1}{\sqrt{\bar{\alpha}_t}}(\boldsymbol{x}_t + (1 - \bar{\alpha}_t)\nabla_x \log p_t(\boldsymbol{x}_t))$$

$$\boldsymbol{x}_{t-1} = \boldsymbol{x}_{t-1} - \frac{\zeta}{\|\nabla_{x_t}\|\mathcal{A}(\hat{\boldsymbol{x}}_0)) - \boldsymbol{y}\|^2\|}\nabla_{x_t}\|\mathcal{A}(\hat{\boldsymbol{x}}_0) - \boldsymbol{y}\|^2$$

Experimental Results

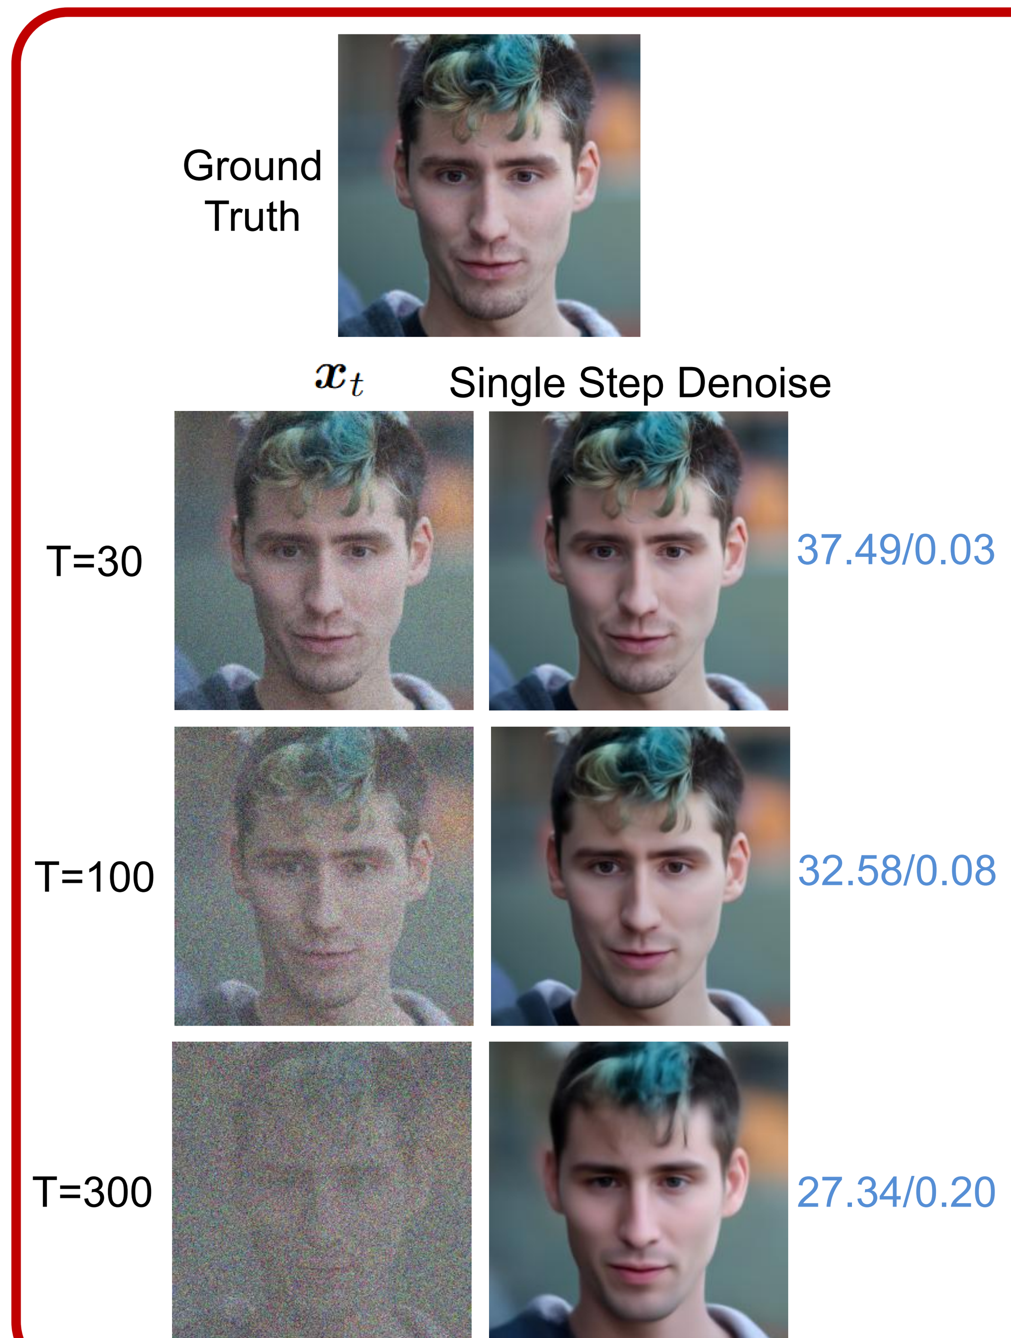

Unconditional Generation: