Inverse Imaging Methods with Diffusion Model Prior

Jinhyo Huh

Stanford University

Motivation

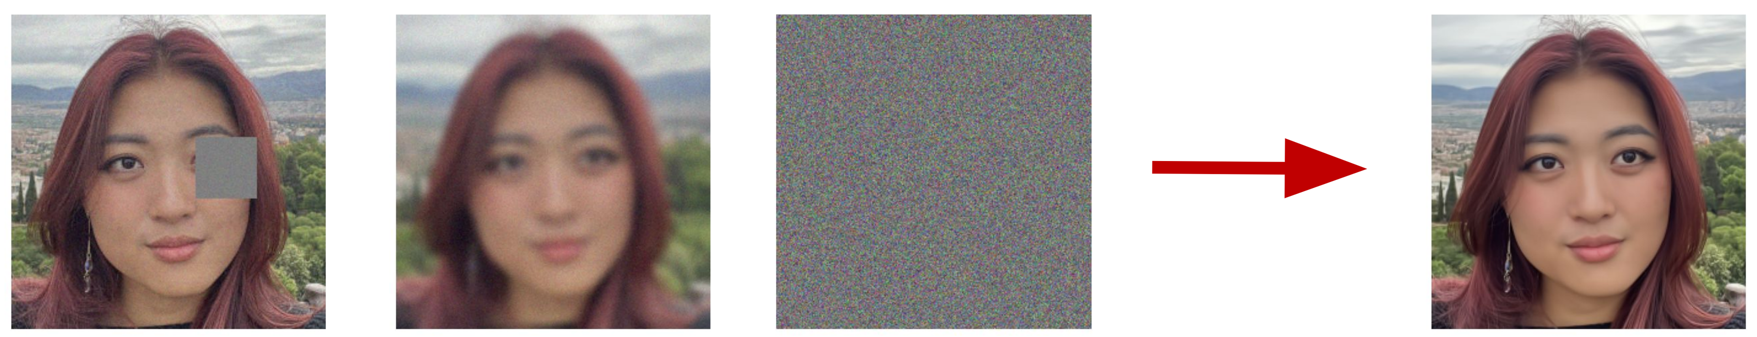

Inverse Imaging Problem: Given a blurry or corrupted image "measurement", we want to generate feasible reconstructions of the clean ground truth image.

Diffusion Model-Based Prior: Learn how images are gradually corrupted by noise and how to reverse the process to recover natural images. We use a pre-trained, score-based diffusion model for denoising, unconditional generation, inpainting, and deconvolution.

Methods

SDEdit[2]

Partially add noise then denoise using DDPM, Variance-Preserving Stochastic Differential Equation.

$$\boldsymbol{x}_t = \sqrt{\bar{\alpha}_t}\,\boldsymbol{y} + \sqrt{1-\bar{\alpha}_t}\,\boldsymbol{z}$$

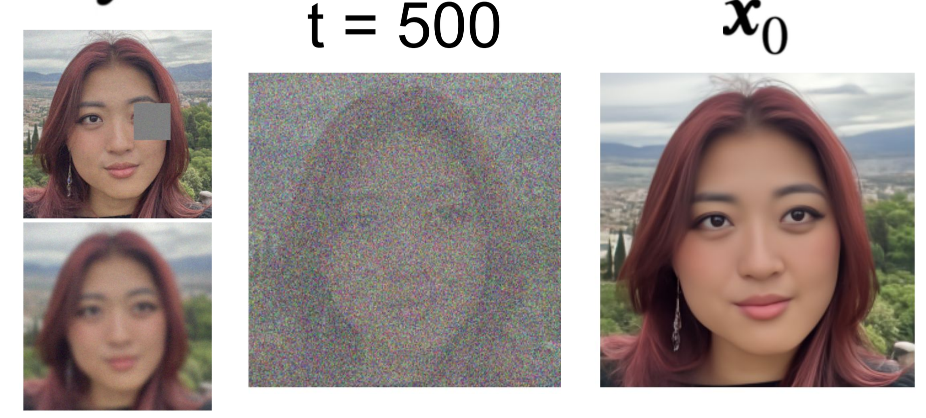

ScoreALD[3]

Langevin Likelihood-Guided Diffusion

Annealing factor γ determines guidance strength of measurement-consistency gradient.

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

Naive Approximation

$$\nabla_x \log p(\mathbf{b}|\mathbf{x_0}) \approx \nabla_x \log p_t(\mathbf{b}|\mathbf{x_t})$$

Diffusion Posterior Sampling (DPS)[4]

Posterior-Guided Reverse Diffusion

$$\boldsymbol{x}_{t-1} = \boldsymbol{x}'_{t-1} - \zeta_t \nabla_{\boldsymbol{x}_t}\left\|\mathscr{A}\left(\hat{\boldsymbol{x}}_0\right) - \boldsymbol{y}\right\|^2$$

Uses diffusion model estimate of the clean image $\hat{\boldsymbol{x}}_0$ as guidance.

Better Approximation

$$\nabla_x \log p(\mathbf{b}|\mathbf{x_0}) \approx \nabla_x \log p_t(\mathbf{b}|\mathbf{x_0} = \mathbb{E}[\mathbf{x_0}|\mathbf{x_t}])$$

Related Work

DDPM (Denoising Diffusion Probabilistic Models)[1]:

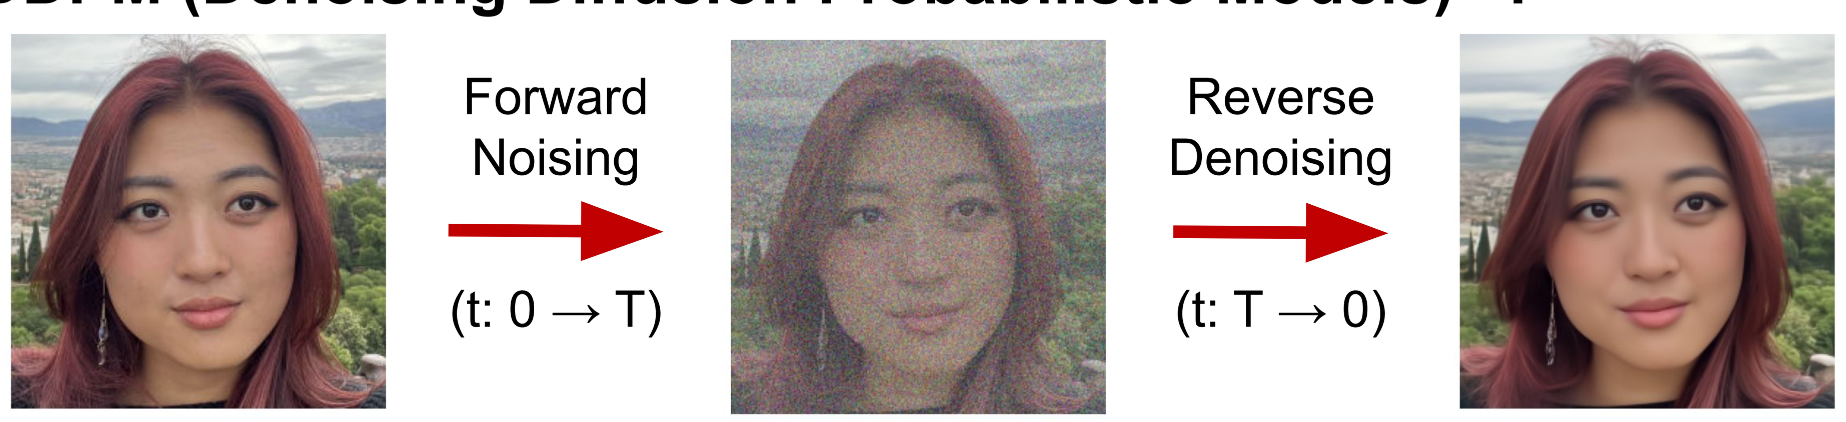

Score Function:

$\nabla_{x_t} \log p_t(x_t) = \boldsymbol{s}_\theta(x_t, t)$

PSNR (Peak Signal-to-Noise Ratio): Higher = Better

LPIPS (Learned Perceptual Image Patch Similarity): Lower = Better

Experimental Results

Ground Truth

DDPM Denoising

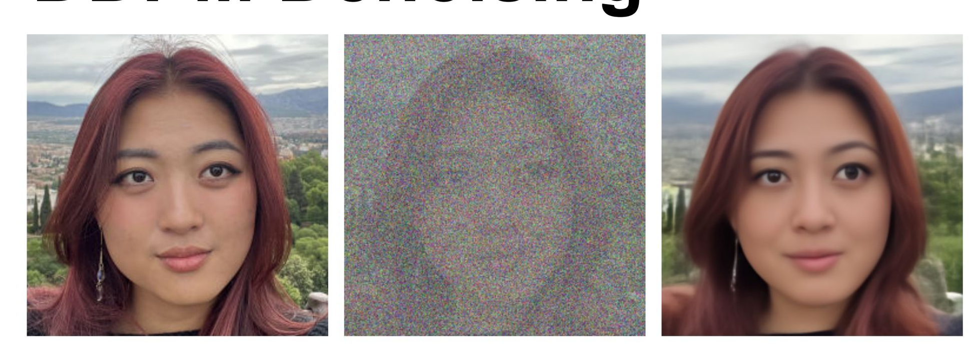

t = 300, PSNR: 25.80, LPIPS: 0.208

Unconditional Generation

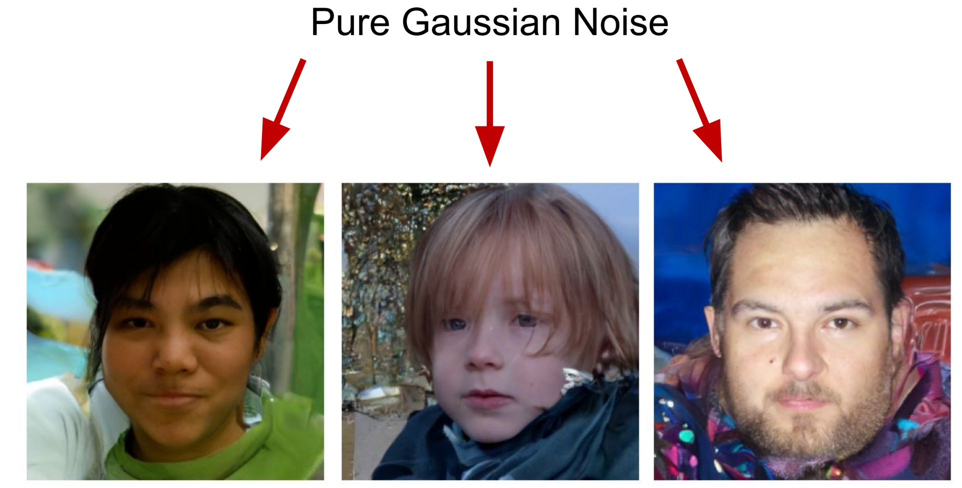

SDEdit

t = 500, PSNR: 19.04, LPIPS: 0.234

t = 500, PSNR: 20.15, LPIPS: 0.277

ScoreALD

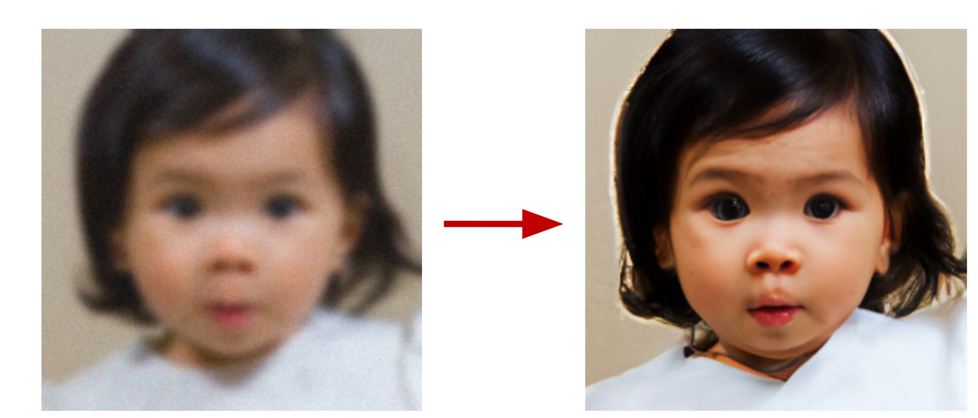

γ = [10, 15]

PSNR: 21.95

LPIPS: 0.163

Diffusion Posterior Sampling

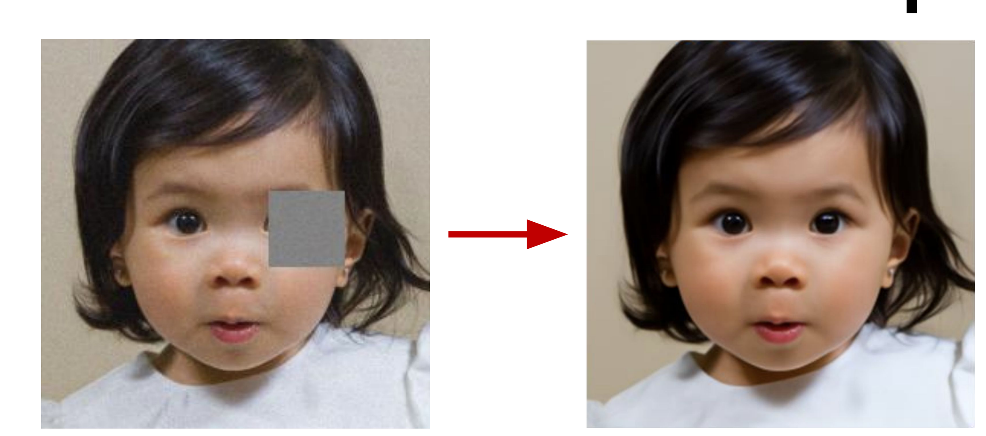

ζ = 0.3

PSNR: 33.65

LPIPS: 0.031

References

[1] Ho et al., Denoising diffusion probabilistic models, NeurIPS, 2020.
[2] Meng et al., SDEdit: guided image synthesis and editing with stochastic differential equations, ICLR, 2022.
[3] Jalal et al., Robust compressed sensing MRI with deep generative priors, NeurIPS, 2021.
[4] Chung et al., Diffusion posterior sampling for general noisy inverse problems, ICLR, 2023.