Image Denoising and Inverse Problem Solving with Pretrained Diffusion Models

Holland Ferguson

EE367 Final Project, Stanford University, Win 2026

Motivation

Diffusion models learn to reverse a gradual noising process, enabling them to generate high-quality images from pure noise.

This powerful prior can be repurposed beyond generation: by conditioning the reverse process on degraded measurements, we can solve inverse problems like denoising, inpainting, and deblurring.

Overview

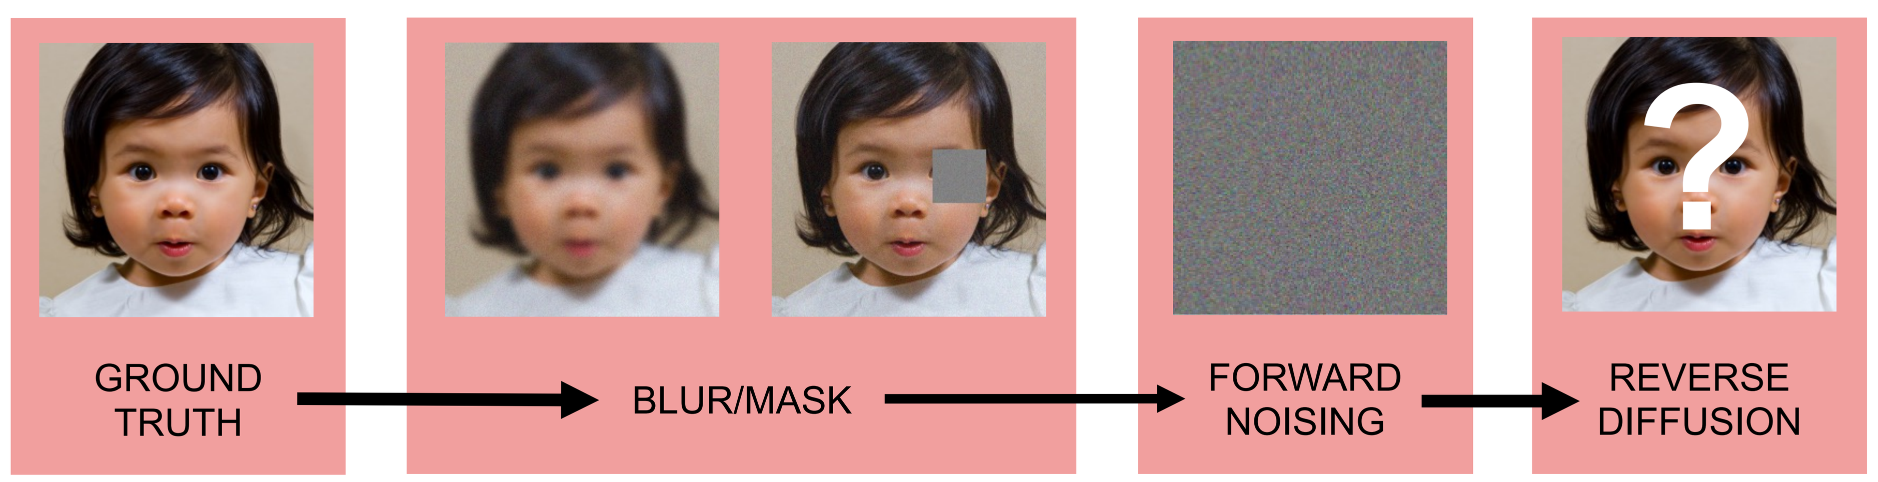

Related Work

DDPM [1]: Learns to reverse a gradual noising process for image generation, but has no mechanism for conditioning on measurements.

SDEdit [2]: Conditions by starting from a partially noised input. Simple, but t_start trades fidelity vs. realism with no measurement enforcement.

ScoreALD [3]: Adds annealed likelihood gradients during sampling. Requires careful tuning; limited to linear inverse problems..

DPS [4]: Normalizes likelihood gradients for stable posterior sampling. Handles linear and nonlinear problems with simpler tuning.

References

[1] Ho, Jain, Abbeel, "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.

Methods

1. Single-Step Denoising

Directly estimate the clean image from a noisy observation using the Tweedie formula:

$\hat{x}_0 = (1/\sqrt{\bar{a}_t}) \cdot x_t + ((1-\bar{a}_t)/\sqrt{\bar{a}_t}) \cdot \mathrm{score}(x_t, t)$

2. Unconditional Generation

Start from pure Gaussian noise and iteratively apply the DDPM posterior for 1000 timesteps:

$x_{t-1} \sim q(x_{t-1} \mid x_t, \hat{x}_0) \quad \text{for } t = T, \ldots, 1$

3. SDEdit

Noise the measurement to timestep t_start via the forward process, then run the reverse DDPM process.

- t_start controls the fidelity–realism tradeoff: lower values preserve more content, higher values allow more model creativity.

4. ScoreALD

At each reverse step, add an annealed likelihood gradient to enforce measurement consistency:

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

5. DPS (Diffusion Posterior Sampling)

Normalizes the gradient for stability using a single scale parameter ζ:

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

Experimental Results

Denoising

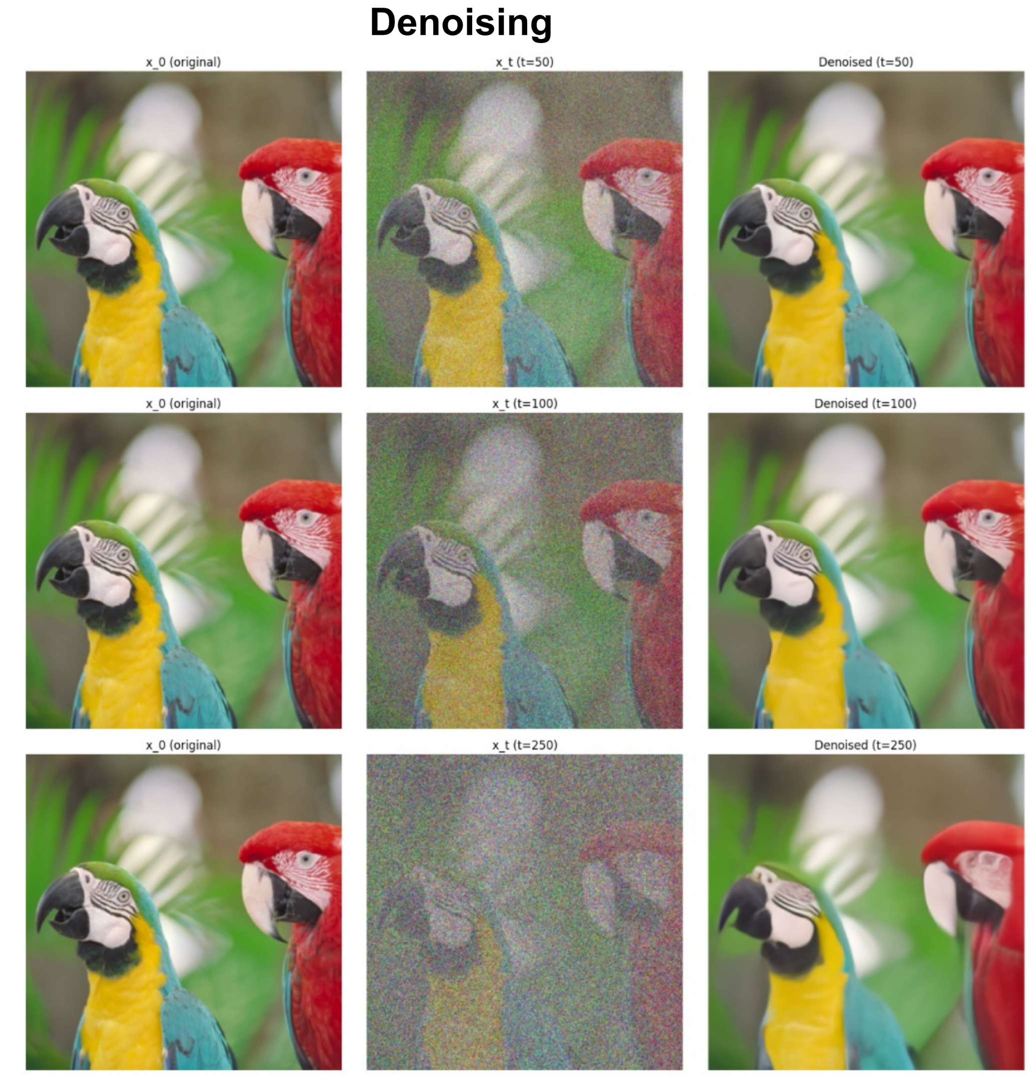

Unconditional Image Generation

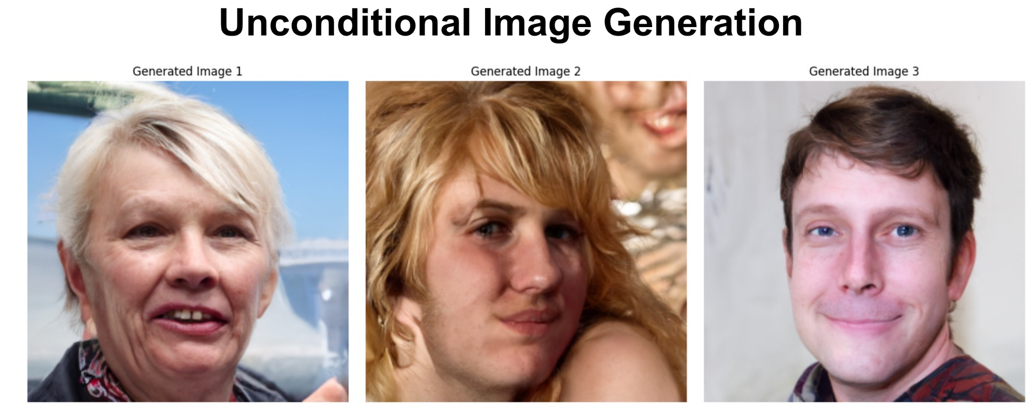

SDEdit

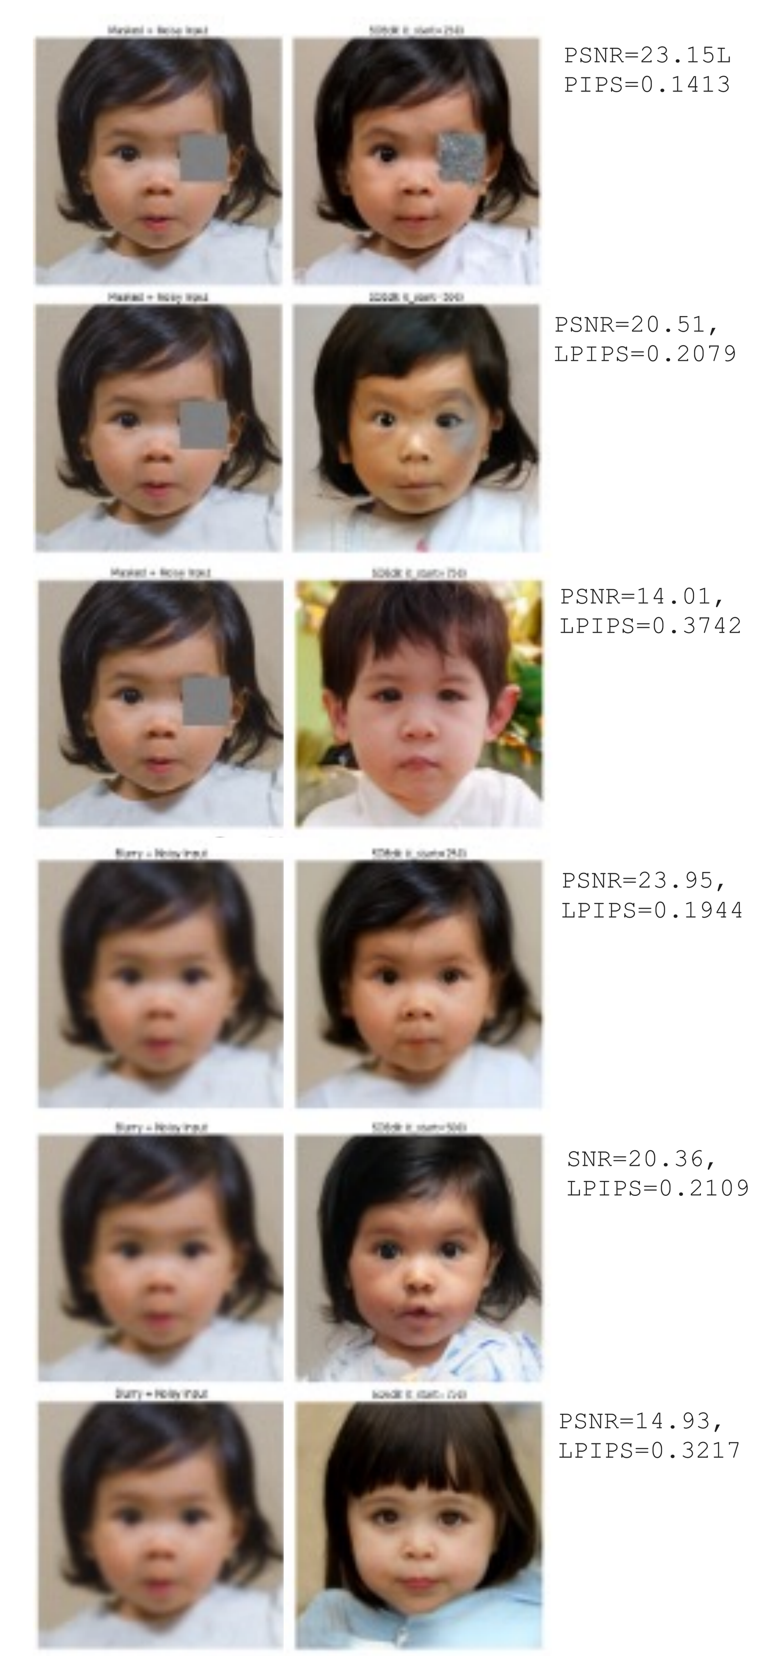

ScoreALD

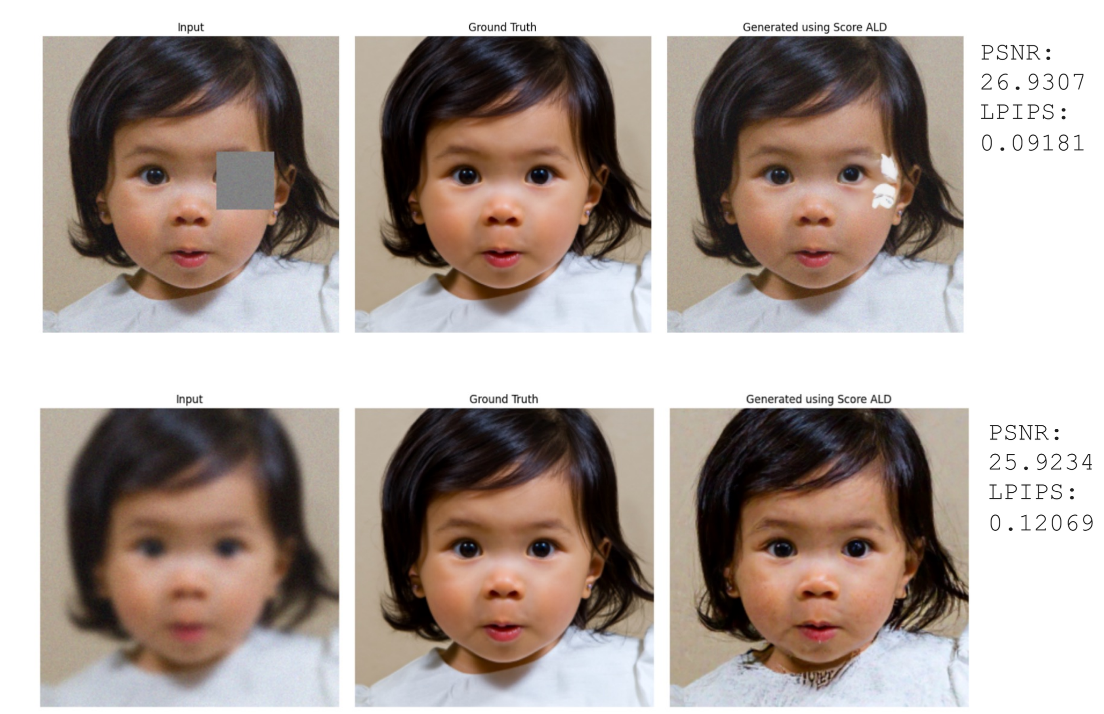

DPS