

Diffusion Models for Solving Inverse Problems

Sonia Minseo Kim

EE367 / CS448I Computational Imaging • Stanford University


Motivation

Inverse problems in imaging (inpainting, deblurring) are ill-posed: the forward operator A is non-invertible, so recovering x from $y = A(x) + n$ requires a prior.

Diffusion models provide a learned image prior via iterative denoising. We implement and compare three methods for incorporating measurement consistency into the reverse diffusion process.

Background: DDPM [1]

Forward process: gradually corrupt a clean image with Gaussian noise



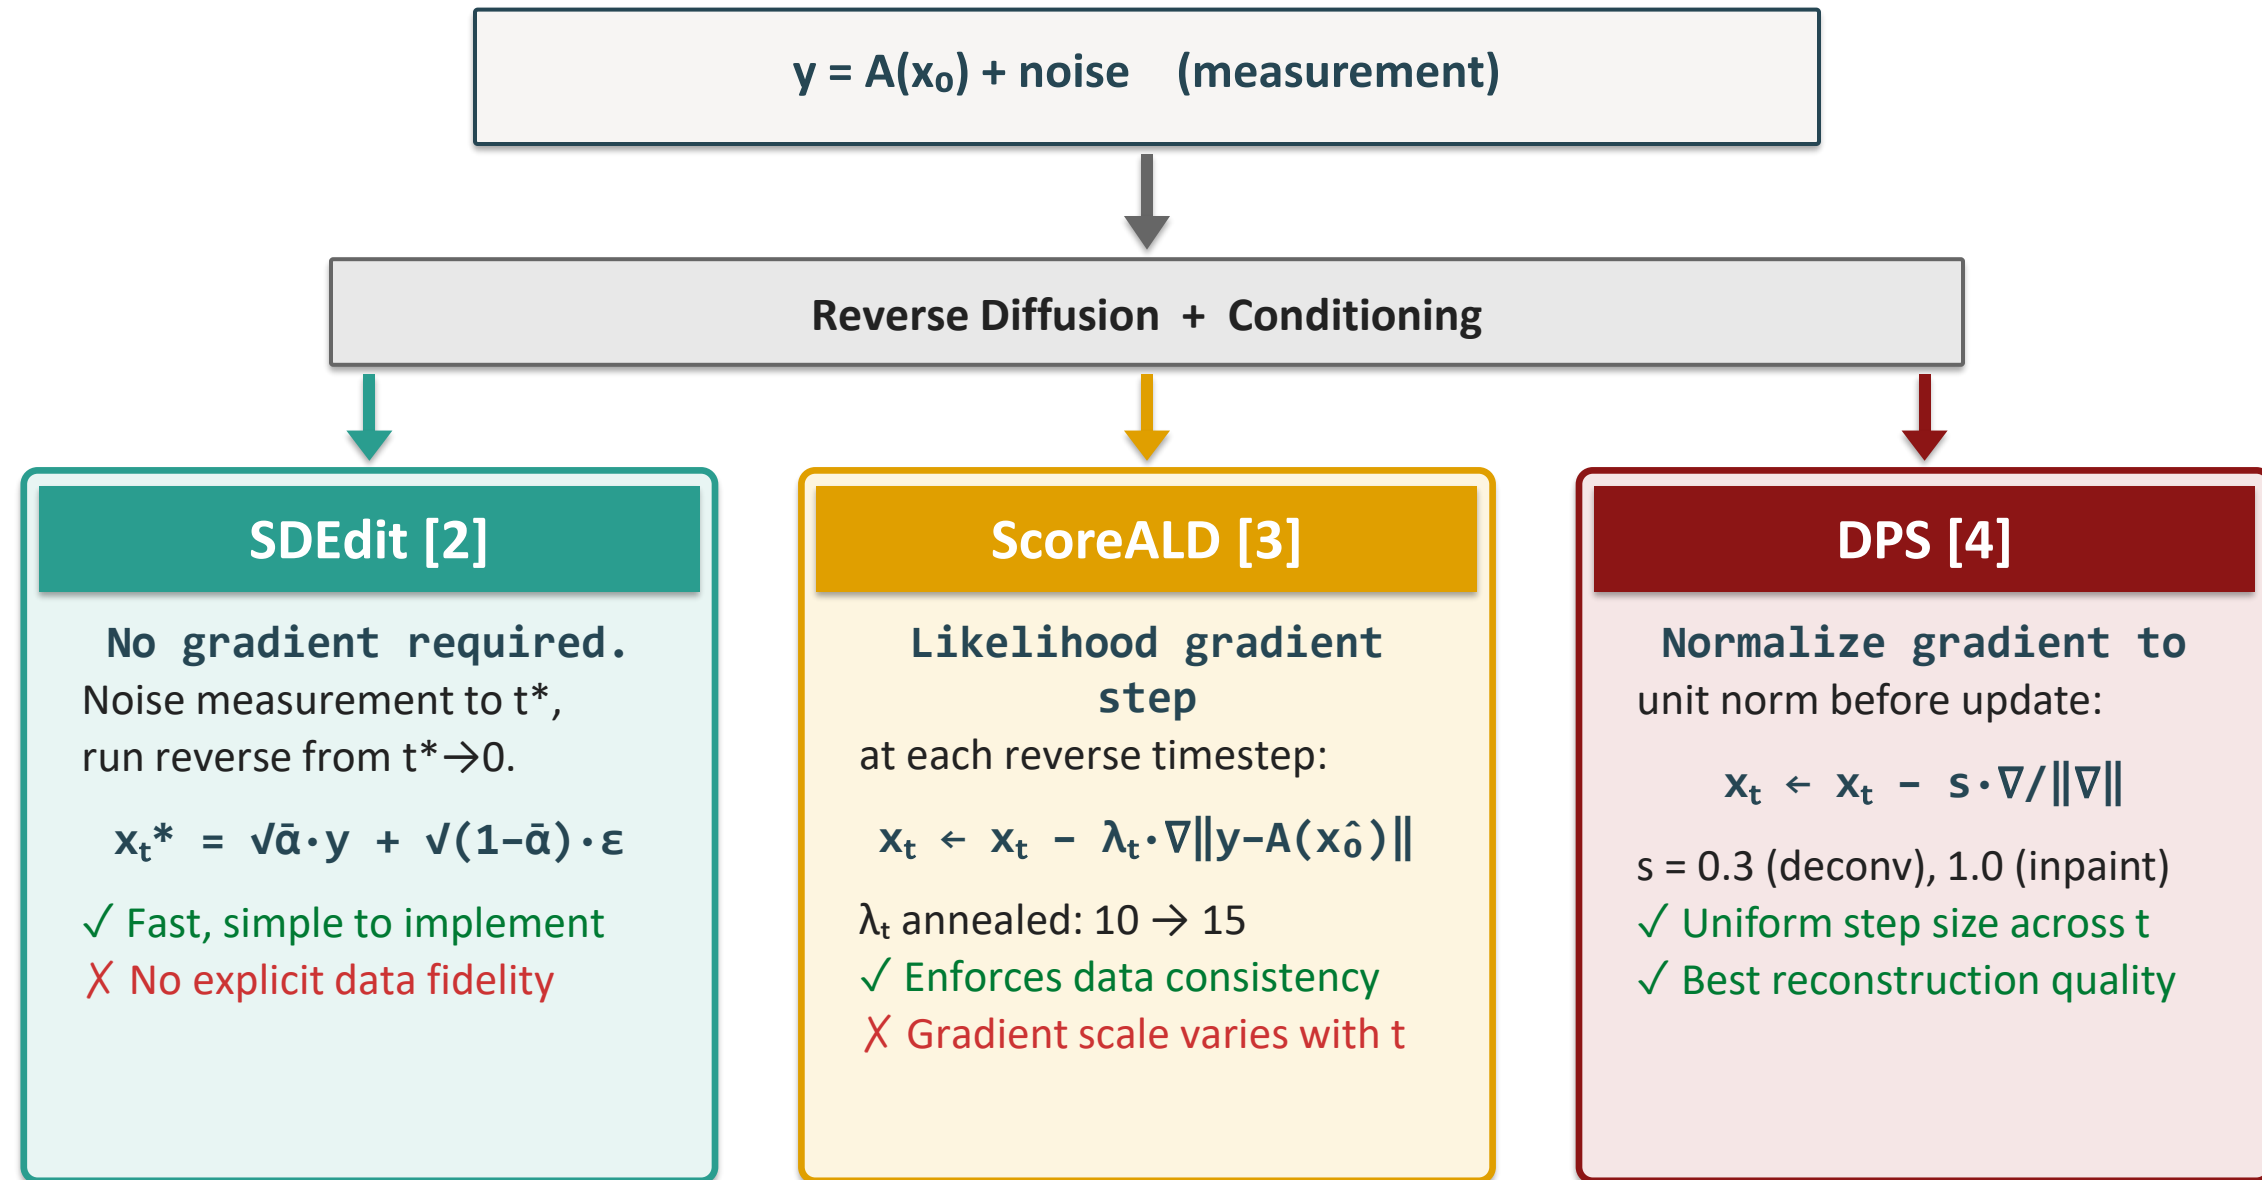
$$x_t = \sqrt{\alpha_t} \cdot x_0 + \sqrt{1 - \alpha_t} \cdot \epsilon$$

Reverse: A trained network estimates $\nabla \log p(x_t)$. Tweedie's formula yields a one-step clean image estimate \hat{x}_0 .

Single-Pass Denoising — PSNR 32.72 dB | LPIPS 0.092



Approach: Three Conditioning Strategies



Unconditional Generation

Initialize $x_t \sim N(0, I)$ and run the full 1000-step DDPM reverse process to sample from the learned face distribution (FFHQ-256).

No ground truth exists; evaluation is qualitative.



References

- [1] Ho, Jain, Abbeel. "Denoising Diffusion Probabilistic Models." NeurIPS 2020.
- [2] Meng et al. "SDEdit: Guided Image Synthesis and Editing with SDEs." 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.

Experimental Results

SDEdit — Box Inpainting ($t^* = 500$)



Masked Input

Ground Truth

SDEdit Result

PSNR: 20.18 dB

LPIPS: 0.185

No gradient computation. Noise the measurement to t^* , then run standard DDPM reverse.

t^* controls the measurement-prior tradeoff: lower t^* preserves more structure; higher t^* relies more on the learned prior.

ScoreALD vs. DPS — Gaussian Deconvolution ($\sigma = 3.0$, kernel 61x61)



Blurred Input

Ground Truth

ScoreALD

DPS

ScoreALD:
25.69 dB / 0.163


DPS:
28.76 dB / 0.055

Gradient normalization accounts for the +3.1 dB PSNR gain.

Detail Crop (eyes region):



Summary & Discussion



Key Findings

- DPS yields best reconstruction quality — 28.76 dB PSNR, 0.055 LPIPS on deconvolution.
- Gradient normalization is essential — +3.1 dB gain over ScoreALD.
- SDEdit is effective without gradients — but lacks explicit measurement fidelity.
- ScoreALD requires careful tuning — raw gradient magnitude varies across t .

Limitations & Future Work

- Results on single test image; broader evaluation needed.
- DPS extends naturally to nonlinear forward models.
- DDIM could reduce the 1000-step sampling cost.