Solving Inverse Imaging Problems with Generative Diffusion Models

Kory Yang, Stanford University

Motivation

- Common imaging inverse problems, such as inpainting and deconvolution, are fundamentally ill-posed.
- While traditional methods rely on simple mathematical priors, the goal of this project is to leverage pre-trained diffusion models to generate priors, allowing us to reconstruct high-fidelity images.

Related Work

- Approaches to inverse problems using "hand-crafted" priors, such as Total Variation, often fail to capture the complexity of natural data.
- Recent advances have shifted toward generative priors. This project builds on the DDPM framework to explore score-based methods like ScoreALD and DPS.

References

[1] Chung, H., Kim, J., Mccann, M. T., Klasky, M. L., and Ye, J. C. (2023). Diffusion posterior sampling for general noisy inverse problems. In ICLR.
[2] Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. In NeurIPS.
[3] Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A. G., and Tamir, J. (2021). Robust compressed sensing mri with deep generative priors.
[4] Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.-Y., and Ermon, S. (2022). Sdedit: Guided image synthesis and editing with stochastic differential equations.

Methods

DDPM

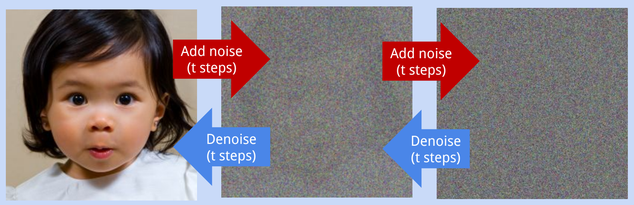

SDEdit

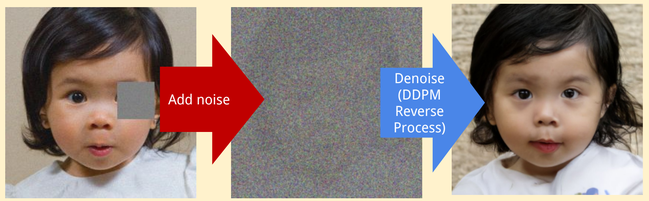

Posterior Sampling (ScoreALD, DPS)

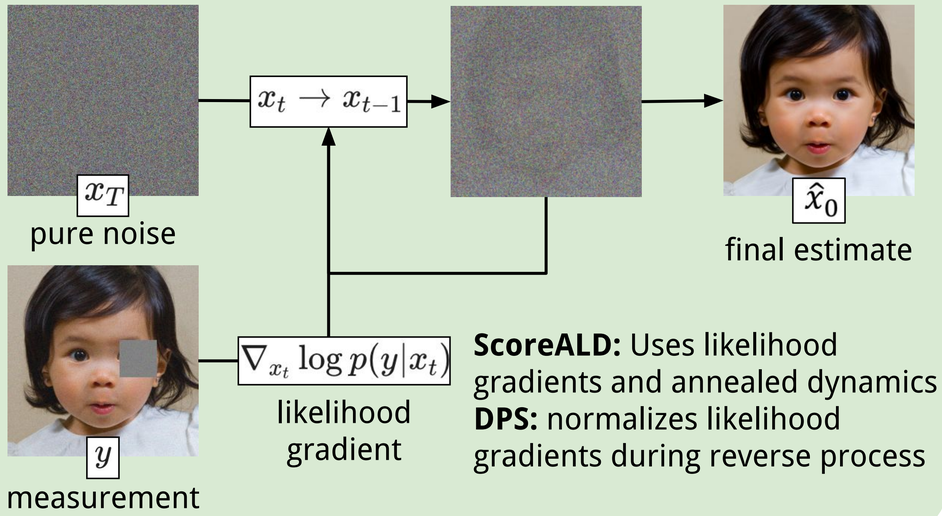

ScoreALD: Uses likelihood gradients and annealed dynamics
DPS: normalizes likelihood gradients during reverse process

Results

DDPM Generation

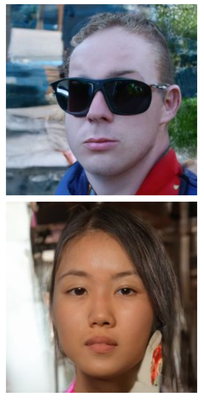

SDEdit (PSNR / LPIPS)