Analysis of Modern Diffusion Techniques for Image Reconstruction

EE 367 - Computational Imaging
Stephen Zhu

Motivation

Ill-posed inverse problem: given a noisy, blurred, incomplete, etc. image, how can we recover the original image?

The method used in this project relies on diffusion models and varying sampling algorithms to attempt reconstruction. Peak Signal to Noise Ratio (PSNR) and Learned Perceptual Image Patch Similarity (LPIPS) are the main metrics for determining success.

Challenges & Limitations

- Hallucinates high-frequency details
- Struggled on non-human subjects

References

1. Chung, Hyungjin, et al. "Diffusion posterior sampling for general noisy inverse problems." arXiv preprint arXiv:2209.14687 (2022).
2. Jalal, Ajil, et al. "Robust compressed sensing mri with deep generative priors." Advances in Neural Information Processing Systems 34 (2021): 14938-14954.
3. Meng, Chenlin, et al. "Sdedit: Guided image synthesis and editing with stochastic differential equations." *arXiv preprint arXiv:2108.01073* (2021).
4. Zhang, Richard, et al. "The unreasonable effectiveness of deep features as a perceptual metric." Proceedings of the IEEE conference on computer vision and pattern recognition. 2018.

Overview

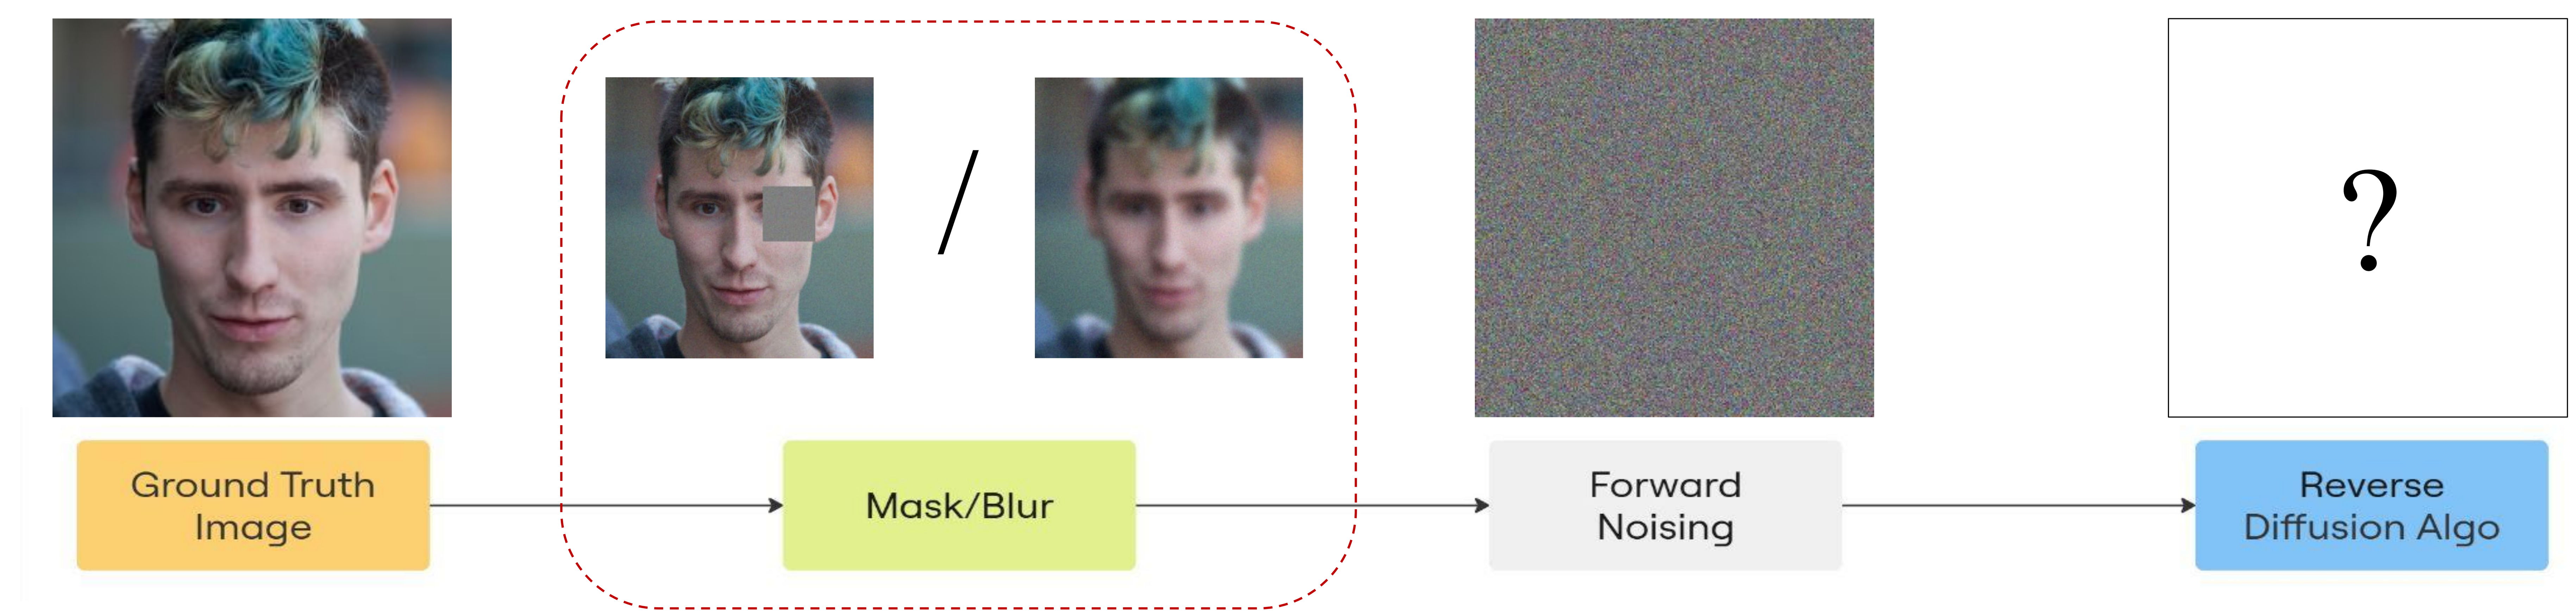

Score-Distillation Editing (SDEdit)

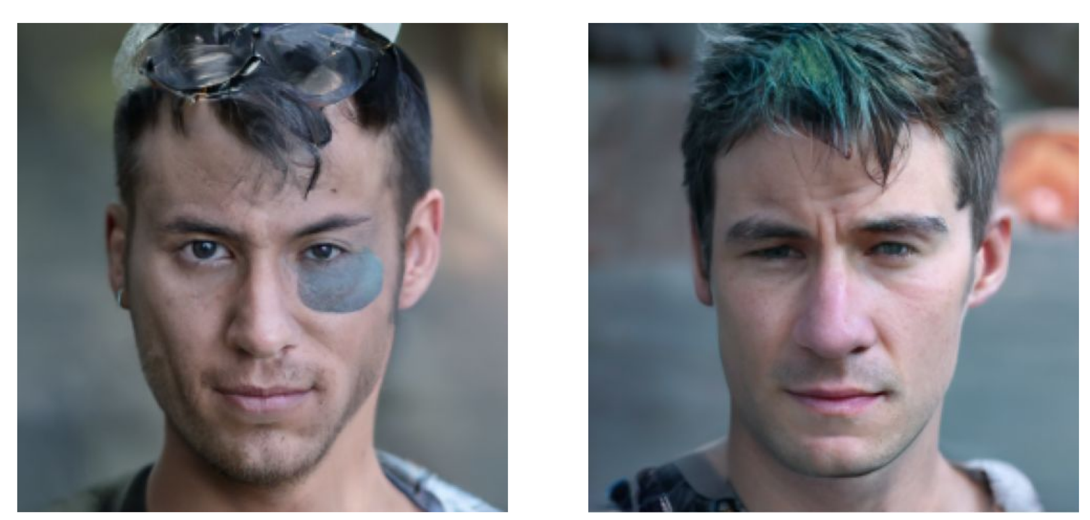

- Partially noise, then reverse diffusion
  - Used 500 / 1000 noising steps
- PSNR = 20
- LPIPS = 0.20

Score-Based Annealed Langevin Dynamics (ScoreALD)

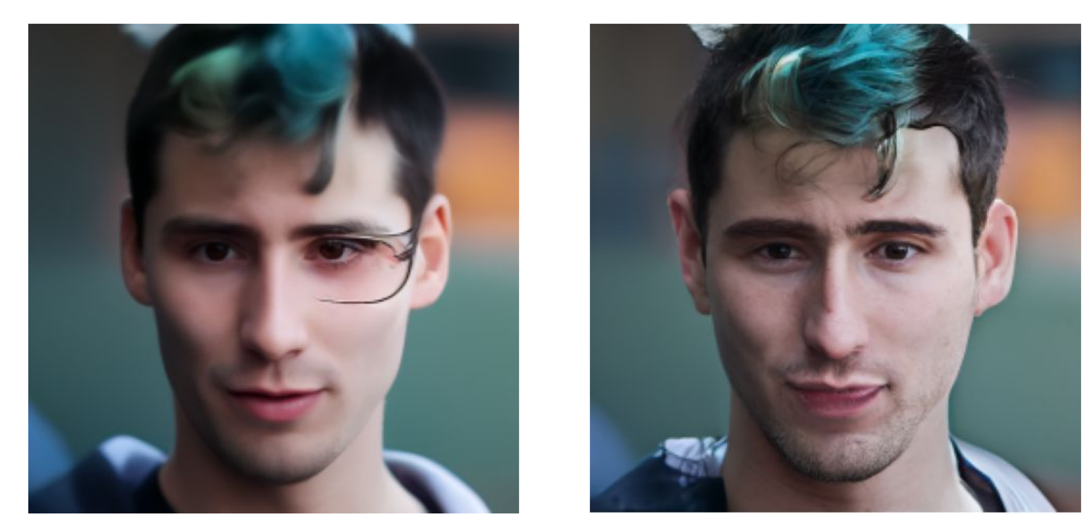

- Utilizes gradient of log likelihood
- Anneal gradient for better convergence
- PSNR = 22
- LPIPS = 0.15

Diffusion Posterior Sampling (DPS)

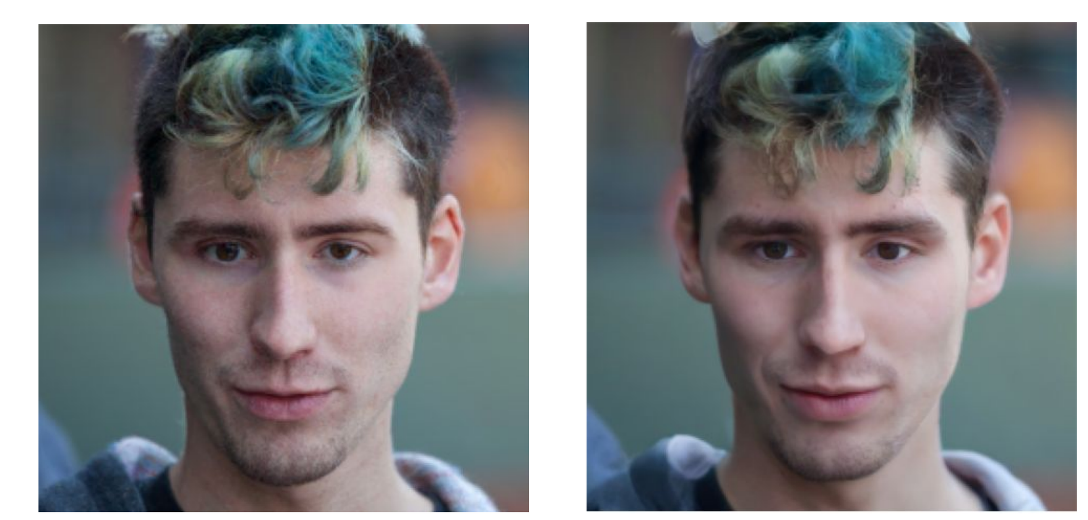

- Takes into account estimate of original image in gradient
- No annealing, but normalize gradients
- PSNR = 30
- LPIPS = 0.05