Inverse Imaging with Diffusion Models

Kien Deshpande
Department of Electrical Engineering, Stanford University

Motivation

Diffusion models generate realistic output images interpolated from training data. These models can be applied to ill-posed inverse imaging problems with algorithms such as SDEdit, ScoreALD, DPS. Diffusion model inverse imaging addresses limitations of classical methods (Wiener filter, ADMM, HQS).

Related Work

Non-Diffusion Methods

- Wiener Filter: resolution loss for high-noise inputs, unstable when transfer function near-zero
- HQS, ADMM: restricted to simple priors (Laplace, L1, TV) for regularizing term of convex opt. problem
- Hybrid methods: Wiener deblur with DnCNN denoise

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.
[5] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein "Distributed optimization and statistical learning via the alternating direction method of multipliers", Foundation and Trends in Machine Learning, 2001
[6] D. Geman and C. Yang "Nonlinear image recovery with half-quadratic regularization", IEEE Transactions on Image Processing, 1995

Approach

Noisy Input x_T
(conv, inpaint + forward process)

Iterative T-step Algorithm

- Current iteration estimate of $\widehat{x_0}$
- Calculate x_{t-1} from current t and score function
- Optional guidance factor

Denoised Output, $\widehat{x_0}$

Forward Diffusion

$$x_t = \sqrt{\overline{\alpha}_t} x_0 + \sqrt{1-\overline{\alpha}_t}\epsilon$$

$$\epsilon \sim \mathcal{N}(0, \mathrm{I})$$

Small t, close to input image
Large t, approaches rand noise

DDPM

$$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$
for $t = T, \ldots, 1$ **do**
$$\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \text{ if } t > 1, \text{ else } \mathbf{z} = \mathbf{0}$$
$$\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t + (1-\bar{\alpha}_t)s_\theta(x_t, t))$$
$$\mathbf{x}_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1-\alpha_t)}{1-\bar{\alpha}_t}\hat{\mathbf{x}}_0 + \sqrt{1-\alpha_t}\mathbf{z}$$

Optional: algorithm dependent guidance factor

SDEdit no added guidance factor

ScoreALD crude approximation

$$\nabla_x logp(\mathbf{b}|\mathbf{x_0}) \approx \nabla_x logp(\mathbf{b}|\mathbf{x}_\mathrm{t})$$
$$x_{t-1} = x_{t-1} - \frac{1}{2(\sigma^2+\gamma_t^2)}\nabla_{x_t}||\mathcal{A}(x_t) - y||^2$$

DPS approximation using Tweedie's Formula

$$\widehat{\mathbf{x_0}} = \frac{1}{\sqrt{\overline{\alpha_t}}}(x_t + (1-\overline{\alpha_t})\nabla_{x_t} logp_t(x_t))$$
$$x_{t-1} = x_{t-1} - \zeta_t \nabla_{x_t}||\mathcal{A}(\widehat{x_0}) - y||^2$$

Experimental Results

SDEdit

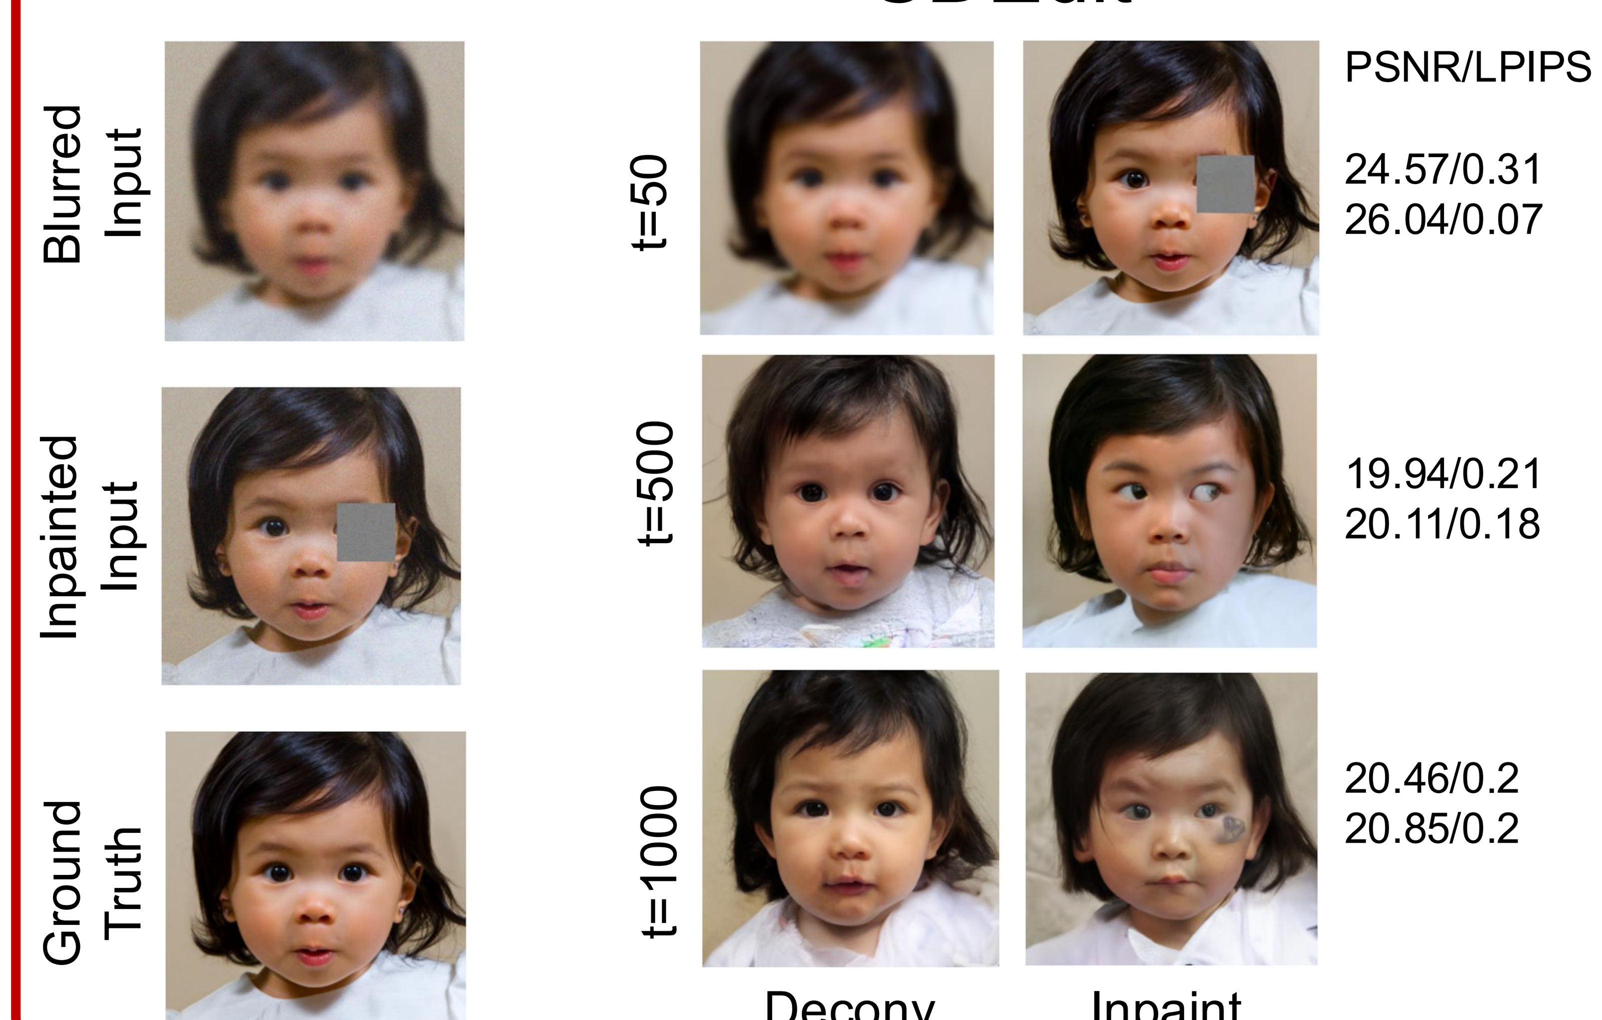

ScoreALD / DPS

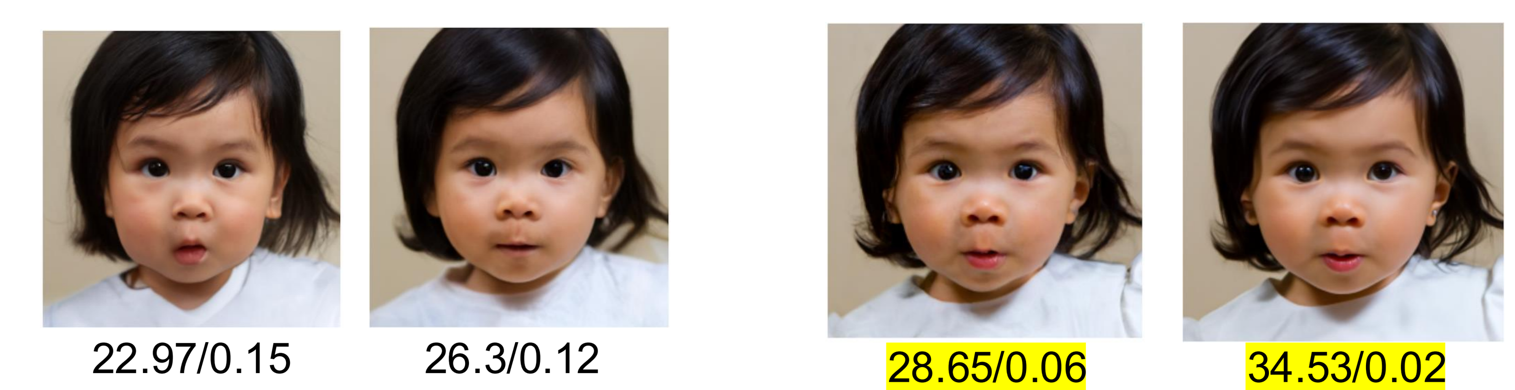

SDEdit: at low t, faithful results retain inpaint and blur artifacts (insufficient added noise). High t, results approach unconditionally generated images.

ScoreALD, DPS: DPS best results, improvement over ScoreALD by better approximation for the forward network gradient (no conditioning on noisy image).