Modern Diffusion Techniques for Imaging

Sylvia Chin, Benjia Zhang
Stanford University, Stanford University

Motivation

There is often a strong need for clear, high resolution image reconstruction from a corrupt source. To achieve a cleaned image from degradation without using unique trained models for each, we can use a universal, pretrained diffusion model prior. We explore three different techniques with the diffusion model: SDEdit (Score-Distillation Editing), ScoreALD, and DPS (Diffusion Posterior Sampling). Both qualitative and quantitative analysis (PSNR and Learned Perceptual Image Patch Similarity (LPIPS)) are critical to evaluating the success of these widely applicable techniques.

Related Work

Several diffusion-based methods have been proposed for solving inverse problems. **SDEdit** edits images by adding noise to an input and applying reverse diffusion, preserving structure but not strictly enforcing measurement consistency. **ScoreALD** introduces a gradient-based correction during sampling to enforce agreement with measurements, though operating on noisy intermediate states can lead to instability. **DPS** improves this by applying the constraint to the estimated clean image, resulting in more stable guidance and better reconstruction quality.

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

1. **Baseline DDPM Sampling:** Create new images from noise.
2. **SDEdit:** Solves inverse problems by adding a specific level of noise to the degraded measurement, then running the standard denoising loop.
3. **ScoreALD:** Steers the diffusion trajectory toward the measurement by applying an anneal factor and likelihood gradients to the noisy image estimates.
4. **DPS:** Eliminates generation artifacts by computing normalized likelihood gradients on the network's clean image predictions

$$y \xrightarrow{\text{Add noise}} x_t = \sqrt{\bar{\alpha}_t}\,y + \sqrt{1-\bar{\alpha}_t}\,z \xrightarrow{\text{Denoise}} x_0$$

Measurement — t is a hyperparameter

ScoreALD

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \ldots, 1$ **do**
  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
  $\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t + (1-\bar{\alpha}_t)\mathbf{s}_\theta(\mathbf{x}_t, t))$
  $\mathbf{x}_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1-\alpha_t)}{1-\bar{\alpha}_t}\hat{\mathbf{x}}_0 + \sqrt{1-\alpha_t}\mathbf{z}$
  $\mathbf{x}_{t-1} = \mathbf{x}_{t-1} - \frac{1}{2(\sigma^2+\gamma^2)}\nabla_{\mathbf{x}_t}\left\|\mathcal{A}(\mathbf{x}_t) - \mathbf{y}\right\|^2$
end for
return $\mathbf{x}_0$

DPS

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \ldots, 1$ **do**
  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
  $\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t + (1-\bar{\alpha}_t)\mathbf{s}_\theta(\mathbf{x}_t, t))$
  $\mathbf{x}'_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1-\alpha_t)}{1-\bar{\alpha}_t}\hat{\mathbf{x}}_0 + \sqrt{1-\alpha_t}\mathbf{z}$
  $\mathbf{x}_{t-1} = \mathbf{x}'_{t-1} - \zeta_t\nabla_{\mathbf{x}_t}\left\|\mathcal{A}\left(\hat{\mathbf{x}}_0\right) - \mathbf{y}\right\|^2$
end for
return $\mathbf{x}_0$

Experimental Results

Unconditional image generation

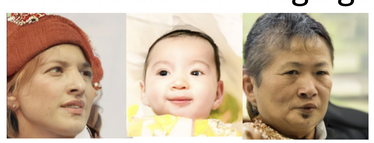

Denoising

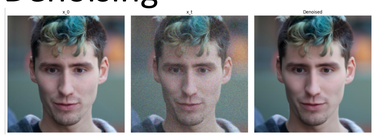

PSNR/LPIPS: 37.45/0.036

Score ALD

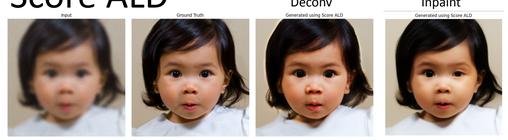

PSNR/LPIPS: 22.69/0.15 PSNR/LPIPS: 23.06/0.11

SDEdit

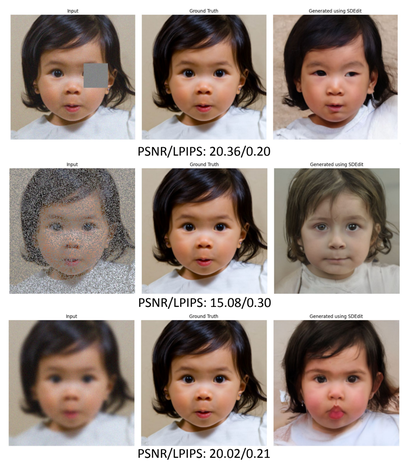

PSNR/LPIPS: 20.36/0.20
PSNR/LPIPS: 15.08/0.30
PSNR/LPIPS: 20.02/0.21

DPS

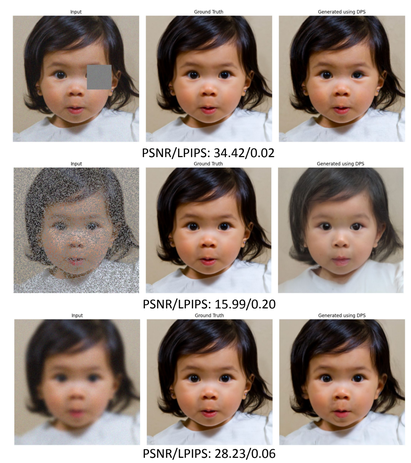

PSNR/LPIPS: 34.42/0.02
PSNR/LPIPS: 15.99/0.20
PSNR/LPIPS: 28.23/0.06