Inverse problems with Pretrained Diffusion models

Venkat S. Rao
Stanford University

Motivation

- Gain better understanding of diffusion models.
- Use pretrained diffusion model as a prior to solve two inverse problems:
  - inpainting
  - deconvolution
- Test multiple state of the art methods

Related Work

- DDPM
- SDEdit
- ScoreALD
- Diffusion Posterior Sampling

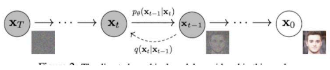

References

[1] Meng, Chenlin, et al. "Sdedit: Guided image synthesis and editing with stochastic differential equations." *arXiv preprint arXiv:2108.01073* (2021).
[2] A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir, "Robust compressed sensing mri with deep generative priors," Advances in Neural Information Processing Systems, vol. 34, pp. 14 938–14 954, 2021.
[3] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," in The Eleventh International Conference on Learning Representations, 2023.
[4] Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." *Advances in neural information processing systems* 33 (2020): 6840-6851.

5 Test Methods

Single-Pass Image Denoising

- Single step to add and remove noise
- Uses a Unet model trained on a Faces dataset to denoise.

Unconditional Generation with DDPM

- DDPM adds noise in with a fixed markov chain.

Algorithm 2 Sampling

1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
2: **for** $t = T, \ldots, 1$ **do**
3: $\quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
4: $\quad \mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)\right) + \sigma_t \mathbf{z}$
5: **end for**
6: **return** $\mathbf{x}_0$

* From [4]

SDEdit

Partially add noise and then Denoise

$$\underset{\text{Measurement}}{\boldsymbol{y}} \xrightarrow{\text{Add noise}} \boldsymbol{x}_t = \underbrace{\sqrt{\bar{\alpha}_t}\boldsymbol{y} + \sqrt{1-\bar{\alpha}_t}\boldsymbol{z}}_{t \text{ is a hyperparameter}} \xrightarrow{\text{Denoise}} \boldsymbol{x}_0$$

Score ALD

Add a condition to DDPM algorithm

$$\boldsymbol{x}_{t-1} = \boldsymbol{x}_{t-1} - \frac{1}{\sigma^2+\gamma_t^2}\nabla_{\boldsymbol{x}_t}\|\mathcal{A}(\boldsymbol{x}_t) - \boldsymbol{y}\|^2$$

Diffusion Posterior Sampling(DPS)

Simplifies condition

$$\boldsymbol{x}_{t-1} = \boldsymbol{x}'_{t-1} - \frac{\zeta_t}{2\sigma^2}\nabla_{\boldsymbol{x}_t}\|\mathcal{A}(\hat{\boldsymbol{x}}_0) - \boldsymbol{y}\|^2$$

Experimental Results

Single-Pass Image Denoising

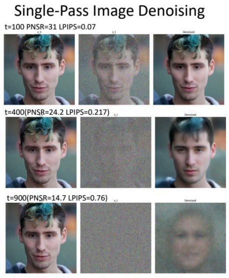

Unconditional Generation

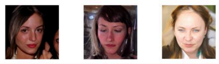

SDEdit

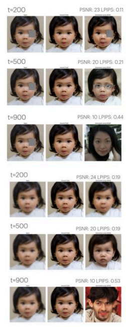

ScoreALD vs DPS