Diffusion Models as Generative Priors for Inverse Problems

Niharika Gunturu

Overview

Diffusion models [1] are a class of generative models that learn to synthesize data by gradually denoising a noisy input. In this work, we explore their capabilities in **image generation**, **denoising**, **inpainting**, and **deconvolution**.

Specifically, we use diffusion model-based methods for the following tasks:

- Unconditional Image Generation from Noise
- Single Step Image Denoising

We also apply diffusion based methods to inverse problems - specifically Inpainting and Deconvolution via three methods:

- **Score Distillation Editing** (SDEdit) [2]
- **Score Annealed Langevin Dynamics** (ScoreALD) [3]
- **Diffusion Posterior Sampling** (DPS) [4]

For these tasks, we use a pre-trained diffusion model[4] that is trained on the Flickr-Faces-HQ Dataset (FFHQ) dataset, and use the variance preserving formulation of diffusion models[1].

Image Generation

The sequence showing the forward (reverse) diffusion process of generating a sample by slowly adding (removing) noise.

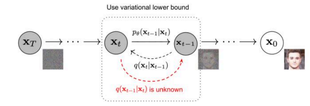

Figure from Ho et al. 2020[1], additionally annotated by Lilian Weng.

Examples of faces generated by the pretrained model starting from noise.

Single Step Image Denoising

Single Step Denoising Algorithm:

$$s \leftarrow \frac{-\epsilon}{\sqrt{1-\bar{\alpha}_t}}; \quad \hat{x}_0 \leftarrow \sqrt{\frac{1}{\bar{\alpha}_t}}(x_t + (1-\bar{\alpha}_t)s)$$

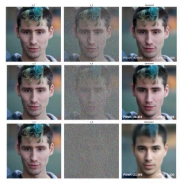

Ground Truth (left), Noisy Input (middle), Model Result (right) for noise levels t= 100, 250 and 500 (top to bottom)

Inverse Problems via SDEdit[2]

Inverse Problem:

$$y = Ax^* + w$$

Iteratively apply the SDEdit Algorithm, to calculate $\hat{x}_0(x_t) \, \forall \, t = T-1, \cdots, 0$.

SDEdit Algorithm:

$$s \leftarrow \frac{-\epsilon}{\sqrt{1-\bar{\alpha}_t}}; \quad \hat{x}_0 \leftarrow \sqrt{\frac{1}{\bar{\alpha}_t}}(x_t + (1-\bar{\alpha}_t)s)$$

$$x_{t-1} \leftarrow q(\hat{x}_{t-1}|x_t, x_0) + \sigma z \quad \text{where} \, w \sim \mathcal{N}(0,1)$$

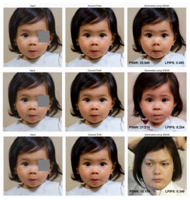

Inpainting for t=250 (top), 500 (middle) and 750 (bottom)

Deconvolution for t=250 (top), 500 (middle) and 750 (bottom)

Inverse Problems via ScoreALD[3]

Inverse Problem:

$$y = Ax^* + w$$

Sampling the posterior $\mu(x|y)$ **using Langevin Dynamics**:

$$x_{t+1} \leftarrow x_t + \eta_t \nabla_{x_t} log \, \mu(x_t|y) + \sqrt{2\eta_t}\zeta_t; \quad \zeta_t \sim \mathcal{N}(0,1)$$

ScoreALD Algorithm: Initialize $x_0 \sim \mathcal{N}_c(0,1)$ and $\forall t = 0, \cdots, T-1$:

$$x_{t+1} \leftarrow x_t + \eta_t \left(f(x_t; \beta_t) + \frac{A^H(y - Ax_t)}{\gamma_t^2 + \sigma^2} \right) + \sqrt{2\eta_t}\zeta_t; \quad \zeta_t \sim \mathcal{N}(0,1)$$

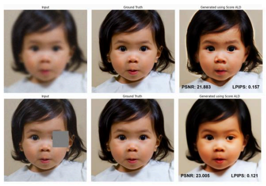

ScoreALD results for deconvolution (top) and inpainting (bottom); $\sigma = 0.05$

Inverse Problems via DPS[4]

Inverse Problem:

$$y = Ax^* + w$$

Diffusion Posterior Sampling Algorithm:

$$x_{i-1} \leftarrow x'_{i-1} - \eta_i \nabla_{x_i} ||y - \mathcal{A}(\hat{x}_0)||^2; \quad \text{where} \, \zeta_i = \frac{c}{||y - \mathcal{A}(\hat{x}_0(x_i))||^2}$$

where for $\forall i = N-1, \cdots, 0$ we have $x_N \sim \mathcal{N}_c(0, I)$ and:

$$\hat{s} \leftarrow s_\theta(x_i, i); \quad \hat{x}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}}}(x_i + (1 - \bar{\alpha}_i \hat{s})$$

$$x'_{i-1} \leftarrow \frac{\sqrt{\alpha_i}(1-\bar{\alpha}_{i-1})}{(1-\bar{\alpha}_i)} x_i + \frac{\sqrt{\bar{\alpha}_{i-1}}\beta_i}{1-\bar{\alpha}_i}\hat{x}_0 + \tilde{\sigma}_i z$$

DPS results for deconvolution (top) and inpainting (bottom)

References

[1] Pieter Abbeel Jonathan Ho Ajay Jain. "Denoising diffusion probabilistic models." In: *In NeurIPS* (2020). URL: `https://doi.org/10.48550/arXiv.2006.11239`.

[2] Yang Song Jiaming Song Jiajun Wu Jun-Yan Zhu Stefano Ermon Chenlin Meng Yutong He. "Guided image synthesis and editing with stochastic differential equations." In: (2022). URL: `https://doi.org/10.48550/arXiv.2108.01073`.

[3] Giannis Daras Eric Price Alexandros G. Dimakis-Jonathan I. Tamir Ajil Jalal Marius Arvinte. "Robust Compressed Sensing MRI with Deep Generative Priors." In: (2021). URL: `https://doi.org/10.48550/arXiv.2108.01368`.

[4] Michael T. Mccann Marc L. Klasky Jong Chul Ye Hyungjin Chung Jeongsol Kim. "Diffusion posterior sampling for general noisy inverse problems." In: *In ICLR* (2023). URL: `https://doi.org/10.48550/arXiv.2209.14687`.