Solving Inverse Problems with Diffusion Model

Fangjun Zhou
Computer Science Department, Stanford University

Motivation

- Inverse problem such as image inpainting and deconvolution are ill-posed.
- Solving these inverse problem with diffusion models leverages the natural image prior to produce higher quality result.

Related Works

- Image inpainting: pattern matching & texture synthesis
  - Limitations: only works for existing texture
- Image deconvolution: Wiener filtering
  - Limitations: simple image formation model

Image Denoising with Diffusion Model

- Forward diffusion process:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}z, \quad t = 1,2,...,T$$

- Backward diffusion process:

$$\hat{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t + (1-\bar{\alpha}_t)s_\theta(x_t, t))$$

$$x_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1-\alpha_t)}{1-\bar{\alpha}_t}\hat{x}_0$$

References

[1] Kundur, et al., Blind image deconvolution, IEEE Signal Processing Magazine, 1996
[2] Efros, et al., Texture synthesis by non-parametric sampling, IEEE International Conference on Computer Vision, 1999
[3] Ho, et al., Denoising Diffusion Probabilistic Models, 2020
[4] Meng, et al., SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations, 2022
[5] Chung, et al., Diffusion Posterior Sampling for General Noisy Inverse Problems, 2024

Method

- Unconditional Image Generation
  - Start from random noise -> denoise according to the backward diffusion model at each time step
- SDEdit
  - Start from the target image -> add noise -> run backward diffusion process
- ScoreALD
  - Start from random noise -> denoise according to the backward diffusion model at each time step -> gradient descent the noisy image on the image formation model
- DPS
  - Start from random noise -> denoise according to the backward diffusion model at each time step -> gradient descent the predicted clean image on the image formation model

Unconditional Image Generation

SDEdit

Score ALD

- Perform well on deconvolution
- Perform worse on inpainting
  - Gradient descent the noisy image
  - Inpainted area doesn't have gradient

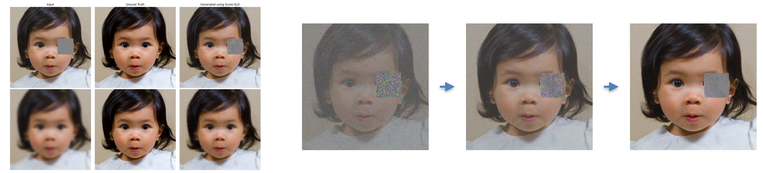

DPS

- Perform well on both inpainting and deconvolution
- Guide the denoising steps using the correct gradient

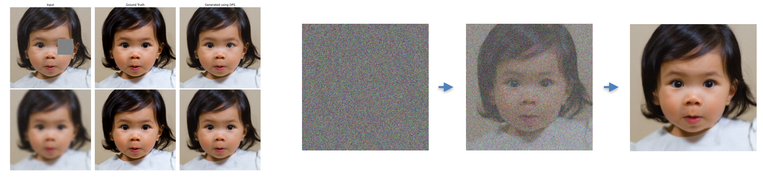

Results

- Unconditional Image Generation
  - Start from random noise -> denoise according to the backward diffusion model at each time step

	PSNR (inpaint)	LPIPS (inpaint)	PSNR (deconvolve)	LPIPS (deconvolve)
SDEdit	22.39	0.32	22.59	0.39
ScoreALD	23.39	0.17	26.84	0.18
DPS	28.51	0.23	26.88	0.25