Diffusion Models for Image Generation and Inverse Problems

Olivia Doman
Stanford University

Motivation

- Diffusion models are powerful generative models that learn to reconstruct images by reversing a gradual noise process.
- These can be used image generation but also to solve inverse problems like image denoising, inpainting, and deblurring.
- In this project we explore how pretrained diffusion models are able to perform image denoising, unconditional image generation, and image restoration with Diffusion Posterior Sampling

Related Work

- Diffusion Models by Ho produce high-quality images through iterative denoising, but standard diffusion sampling can be computationally expensive [1].
- Score based generative models extend diffusion methods using stochastic differential equations [2].
- Diffusion Posterior Sampling by Chung incorporates measurement constraints to solve inverse problems such as inpainting and deblurring [3].

References

[1] Ho et al., Denoising Diffusion Probabilistic Models, NuerIPS 2020
[2] Song et al., Score based Generative Modeling through Stochastic Differential Equations, arXiv, 2020
[3] Chung et al, Diffusion Posterior Sampling for General Noisy Inverse Problems, arXiv, 2022

New Technique

- Approach: Diffusion models generate images by learning to reverse a gradual Gaussian noise In the forward process, noise is added until the entire image is pure noise. A neural network is trained to predict the noise at each step and reconstruct the clean image. During sampling, the model starts from random noise and iteratively removes noises to reconstruct the image. For inpainting and deblurring, we apply Diffusion Posterior Sampling (DPS) to enforce consistency with corrupted measurements.
- Difference: Unlike GANs that generate images in one step, diffusion models generate images through iterative denoising. DPS extends diffusion models to inverse problems by incorporating measurement constraints.

Input Image/ Noise → Add Noise (Forward Diffusion → Predict Noise with Model → Remove Noise → Repeat until Clean Image

Experimental Results

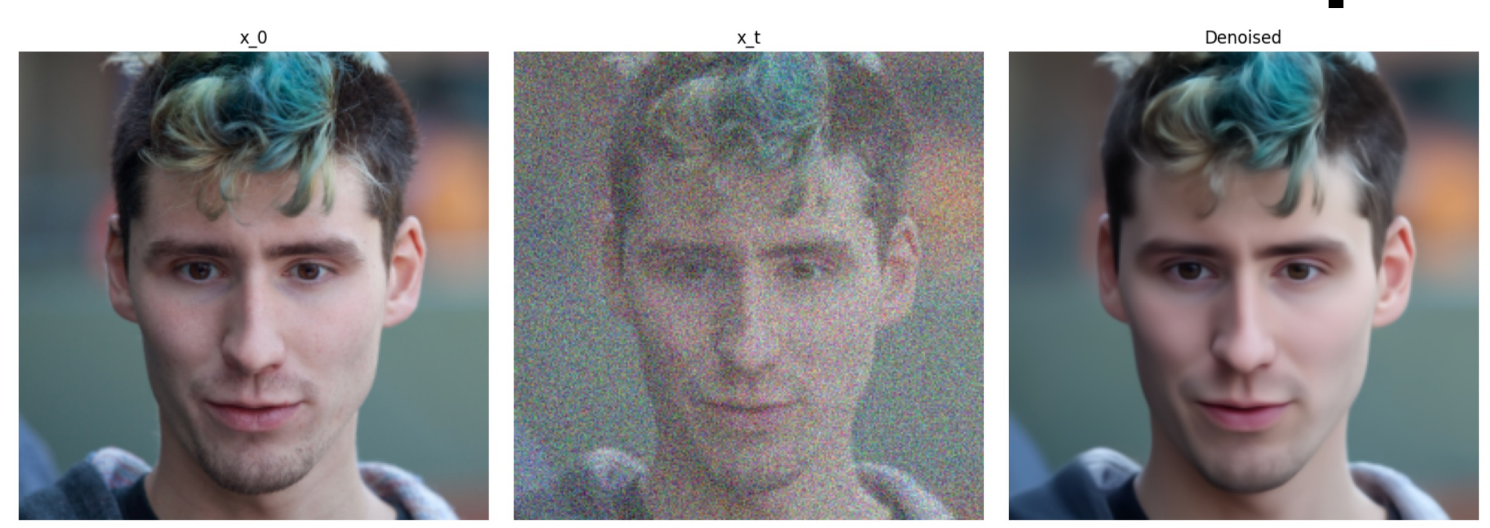

Figure 1: Single Step Denoising with pretrained DDPM.
PSNR: 32.55
LPIPS: 0.093

Figure 2: 1000 step sample from pure Gaussian noise produces Realistic faces.

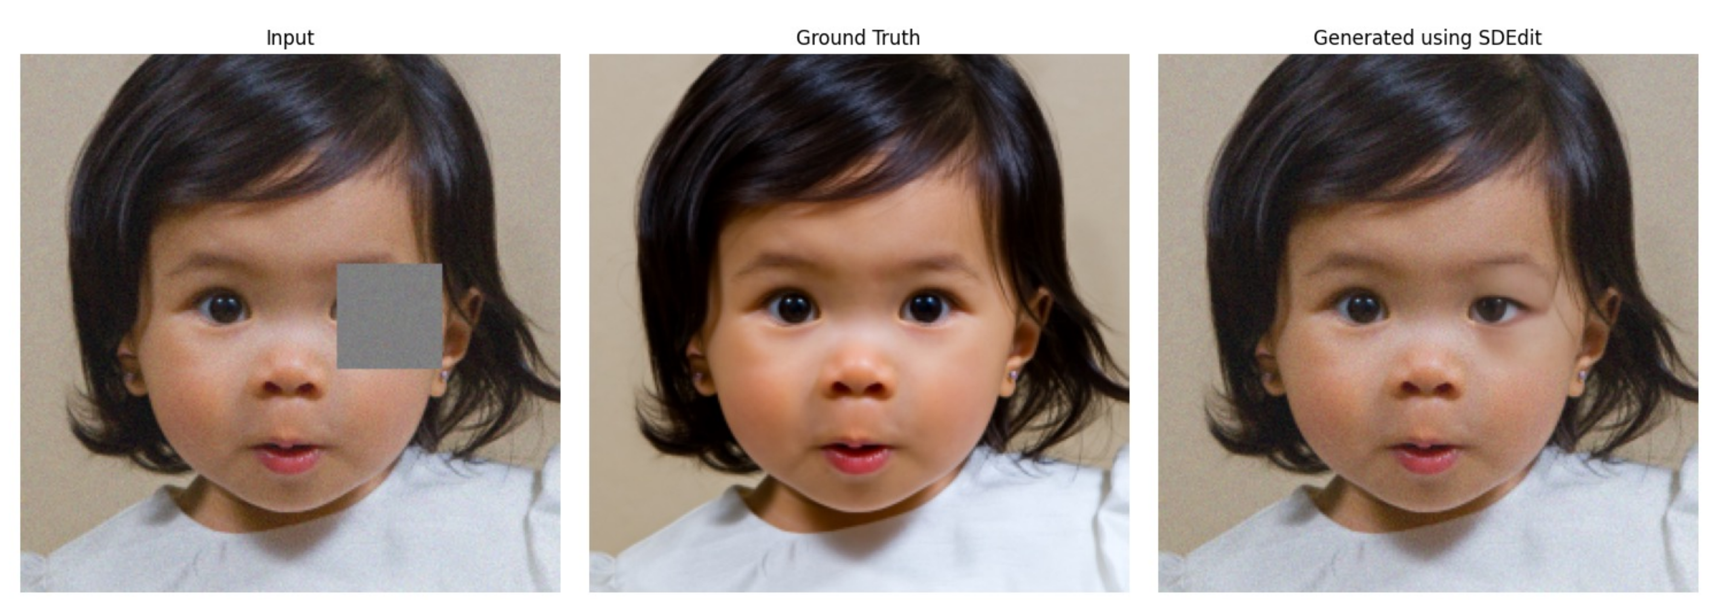

Figure 3: SDEdit
Inpainting: box region inpainted with realistic skin and eye

	300	500	700
PSNR	25.72	29.5776	30.9135
LPIPS	0.122	0.0872	0.08769

Figure 4: Deconvolution (Gaussian Blur): Core Facial Structure recovered

	SDEdit t=500	ScoreALD [10, 20]	ScoreALD [20, 10]
PSNR	20.1440	23.1282	19.4937
LPIPS	0.4452	0.3541	0.3408