Image Denoising with Generative Adversarial Networks

NYU CDS Research Conference
December 11, 2020

Michael Stanley
Chuan Chen

Advisors:
Carlos Fernandez-Granda
Sreyas Mohan

Problem Description

Denoising:

- Recovering the underlying, clean image from a noisy image
- Most approaches to denoising today use mean squared error (MSE) as a loss function
- MSE has known issues in denoising: blurry images, “artifacts”

Our approach:

- Train a GAN that generate realistic clean microscope images
- Find the latent vector that generates a clean image which resembles the noisy image

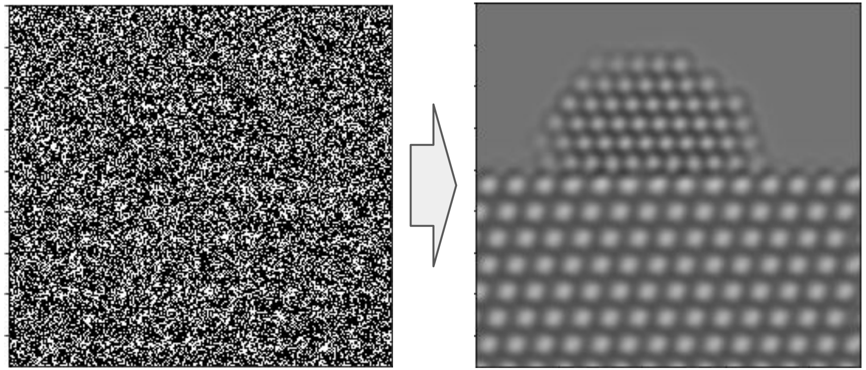

[1]Yang, Qingsong, et al. "Low-dose CT image denoising using a generative adversarial network with Wasserstein distance and perceptual loss." *IEEE transactions on medical imaging* 37.6 (2018): 1348-1357.

Generative Adversarial Network

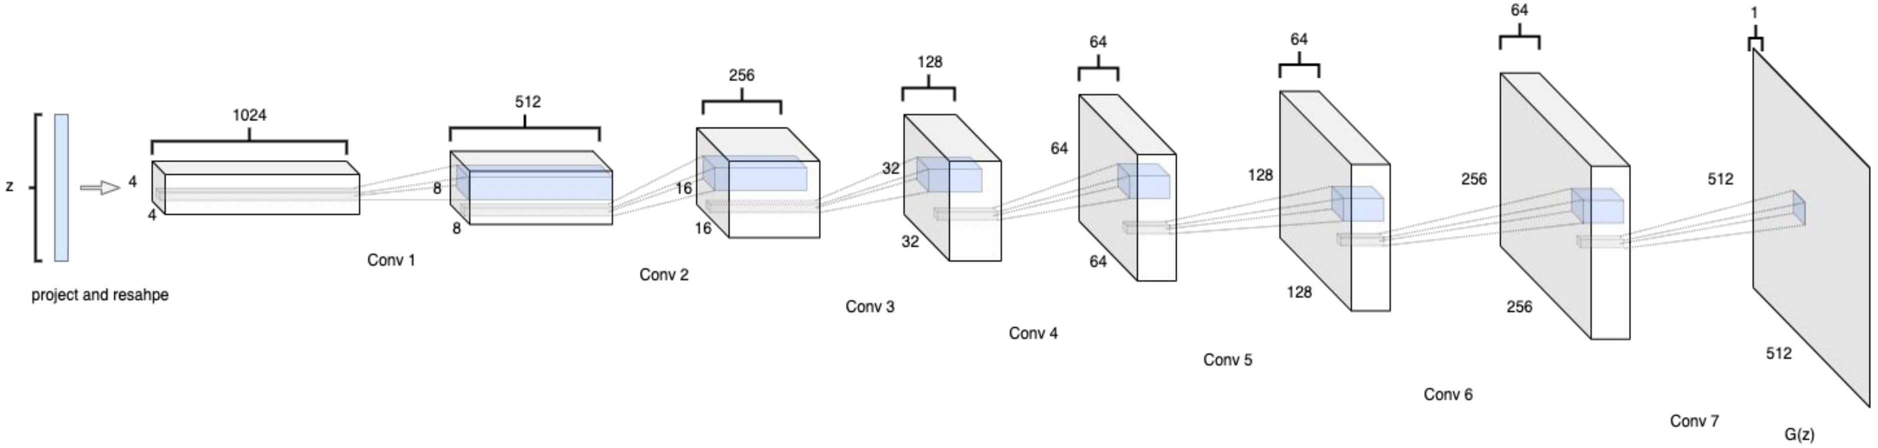

- 3 different latent space: latent vectors of dimension 50, 100, and 150
- all model weights are randomly initialized from a Normal distribution with mean=0, stdev=0.02

Training GAN

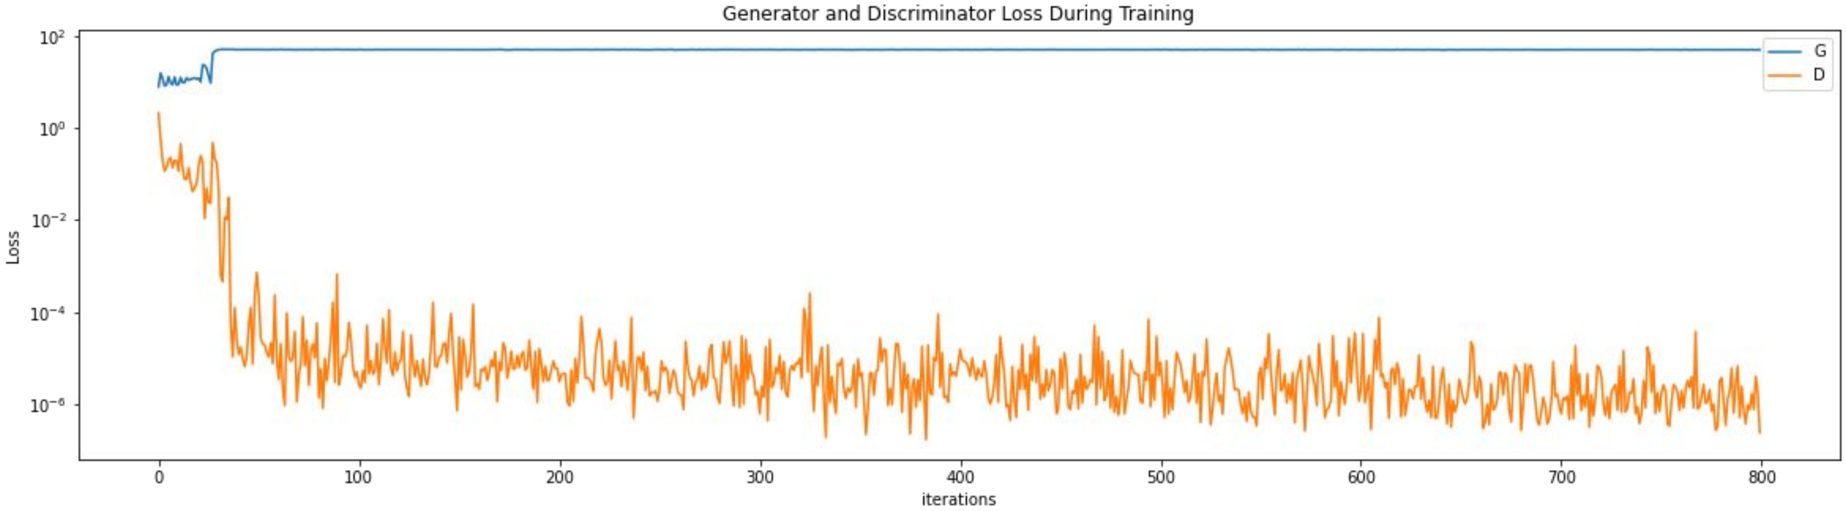

Failure Mode

Modifications to help with training

- Symmetric Architecture on the Generator and Discriminator
- Normalize input to range [-1, 1]
- Removing the sigmoid() layer and using BCELogitLoss
    - Taking advantage of the log-sum-exp trick for numerical stability
- Label smoothing
    - using targets for real examples in the discriminator with a value of 0.9
- Two Time-Scale Update Rule
    - using different learning rates to converge to the Nash Equilibrium

Updated GAN

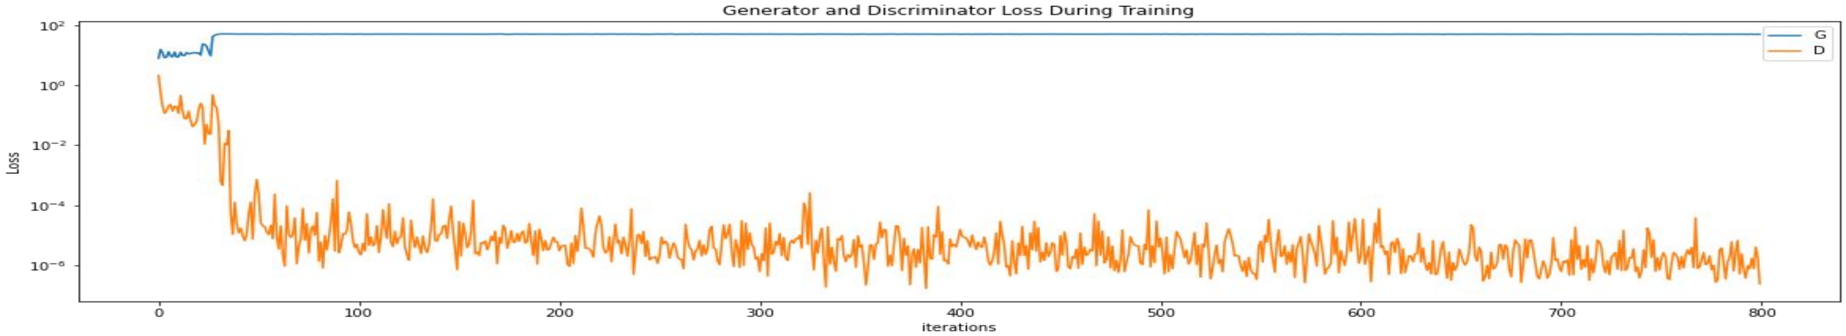
Generator and Discriminator Loss During Training
G
D
Loss
iterations

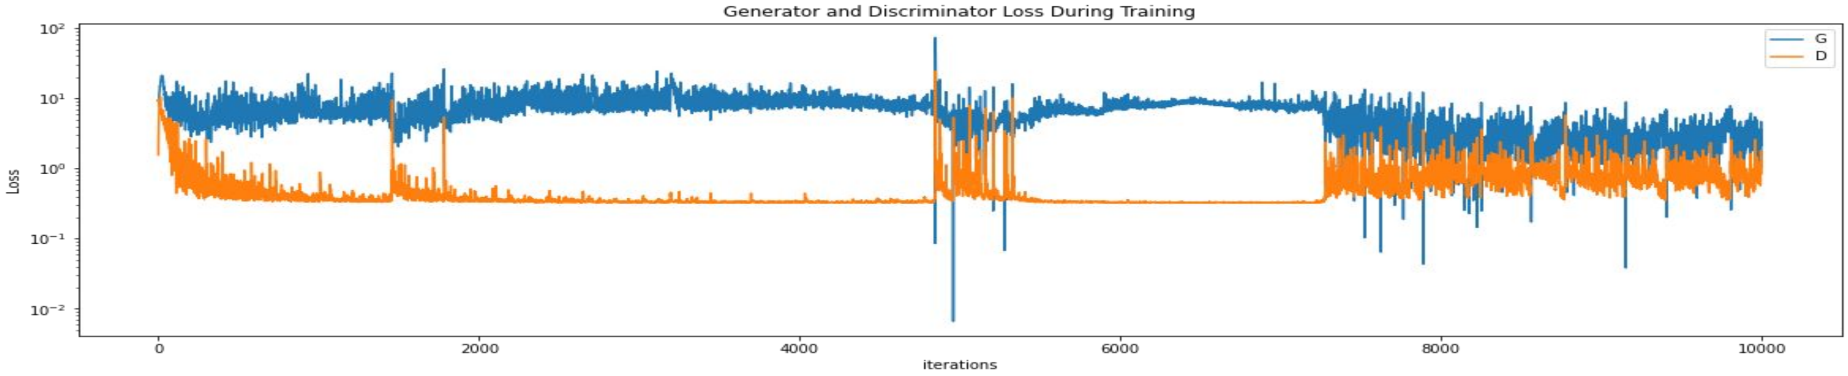
Generator and Discriminator Loss During Training
G
D
Loss
iterations

Generated Images

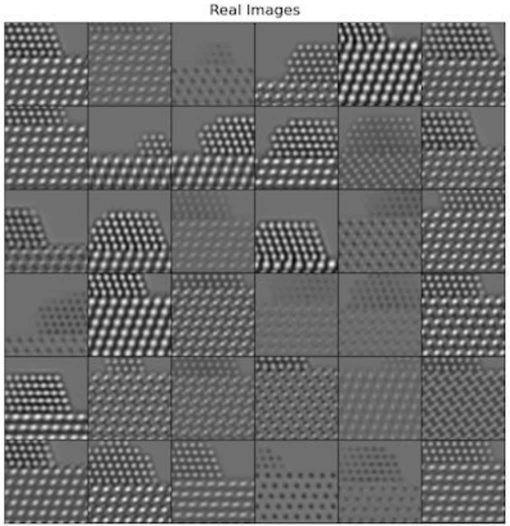

Real images

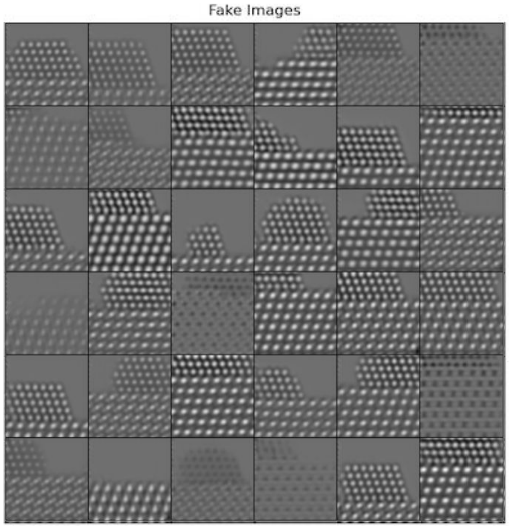

Fake images

G’s Progression

Denoising by optimization in GAN latent space

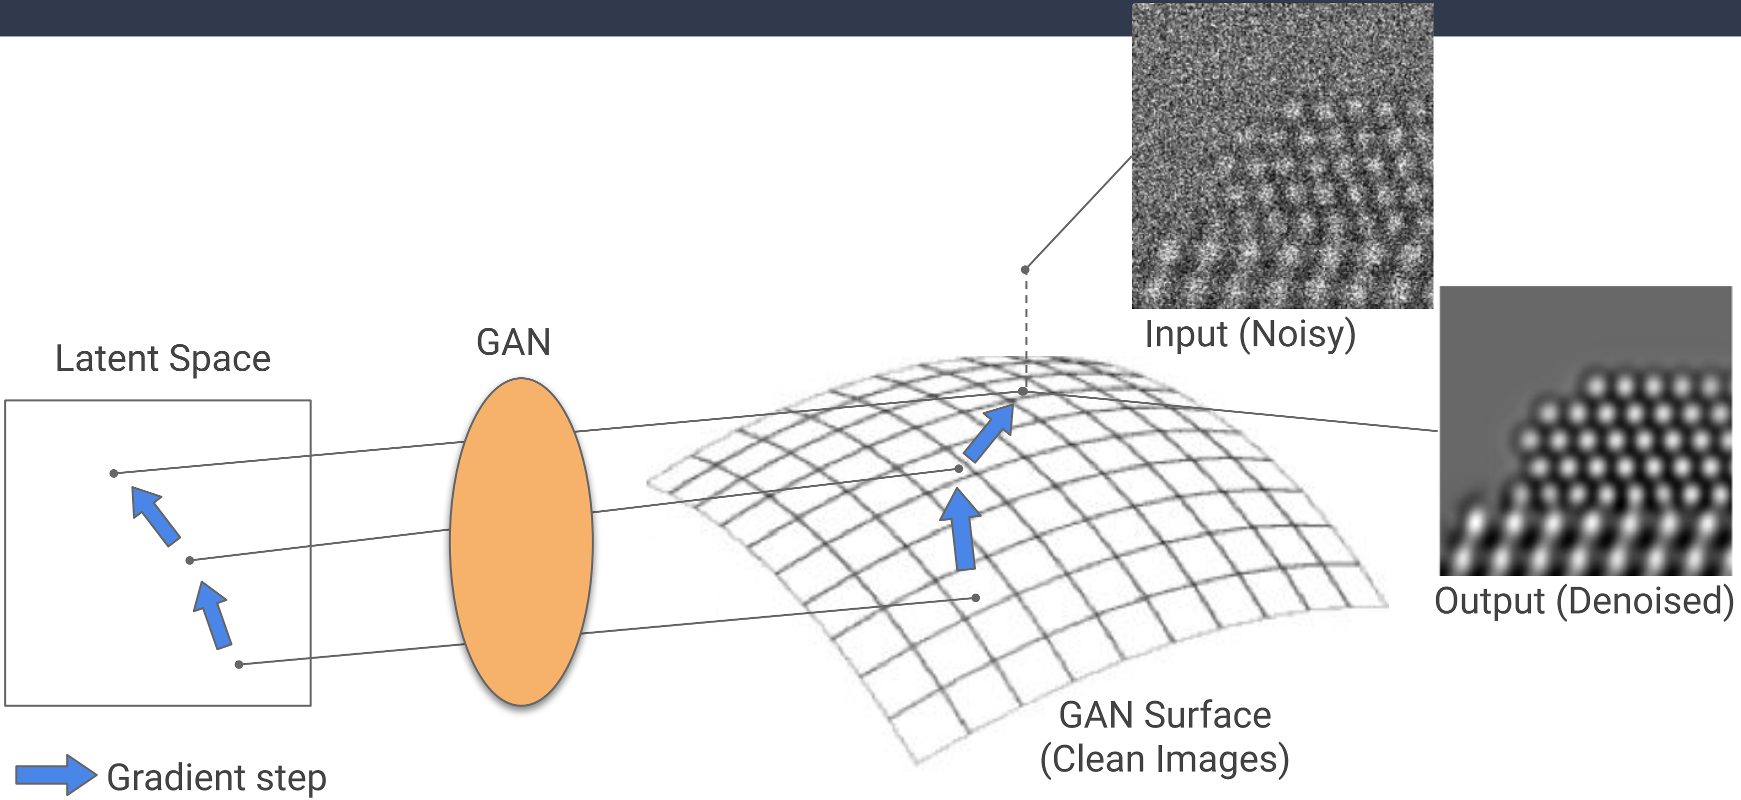

Optimizer algorithm

- To avoid local minima, the optimizer
  - samples the latent space at random
  - performs brief gradient optimization
  - optimizes from the best sample

- Fairly similar results across iterations

- Optimizer typically converges within 20,000 epochs

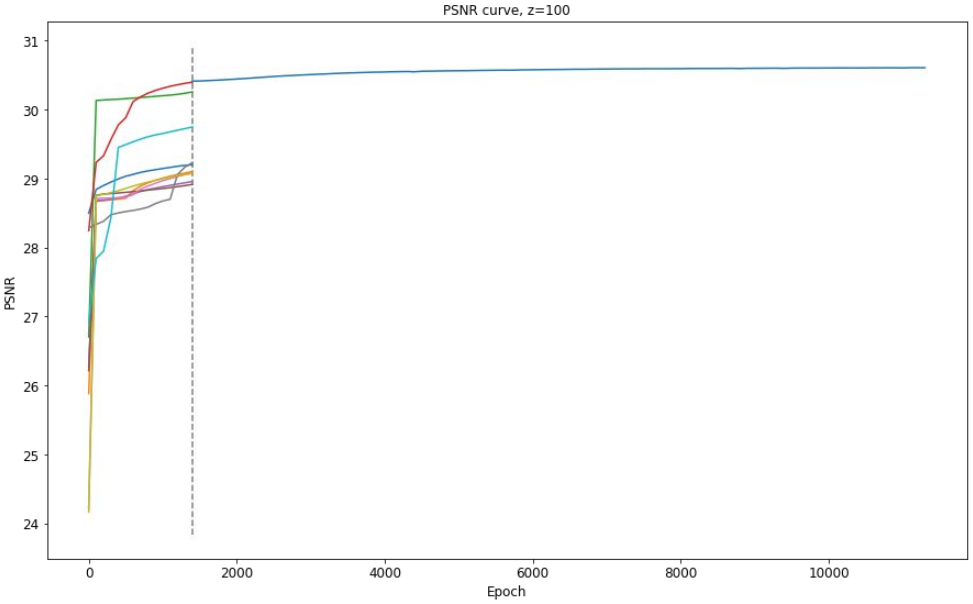

Denoising Results

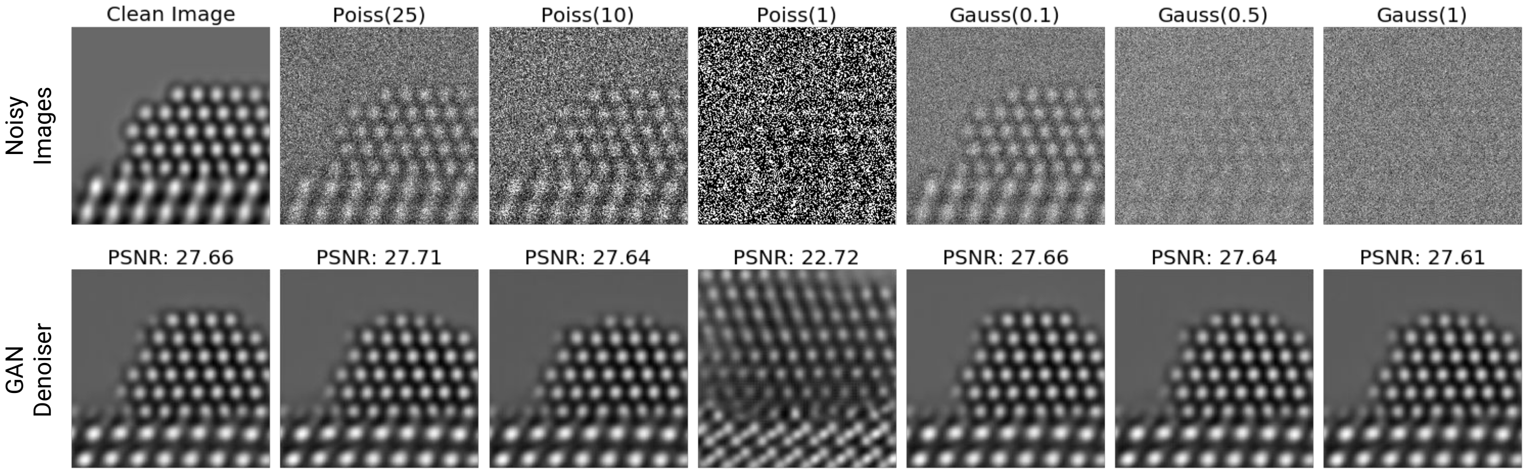

- **Benchmark**: The best supervised MSE denoisers can achieve **PSNR > 30** for Poiss(1)
  - Does not generalize across noise levels
- GAN denoiser works very well, except in the most extreme noise setting
  - **Generalizes across noise levels** without prior knowledge (unlike MSE denoising)
  - Finds **very similar images** regardless of **noise level**

Denoising real microscope data

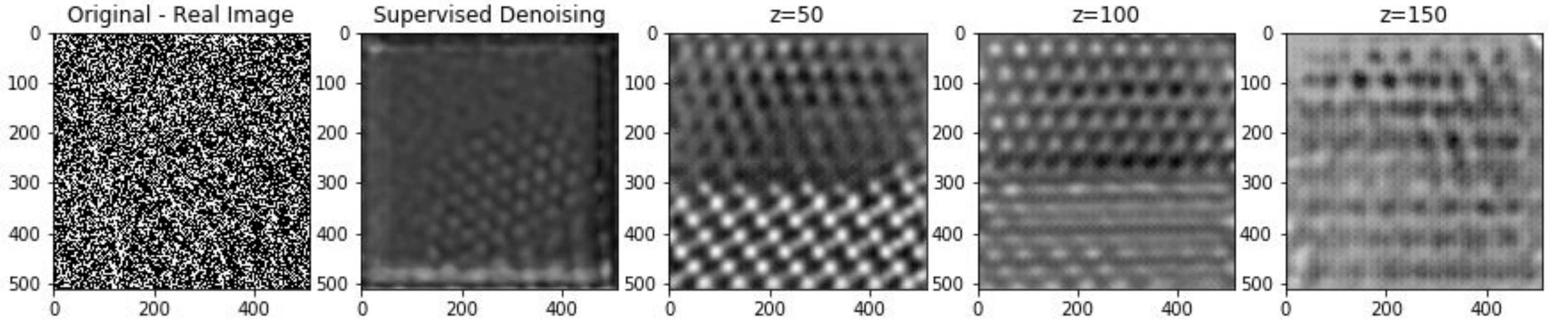

- No benchmark, no ground truth!
- Noise levels are extremely high
- Believed to be Poisson distributed (dependent on pixel values)

Appendix

Denoising Results (incl MSE denoiser)

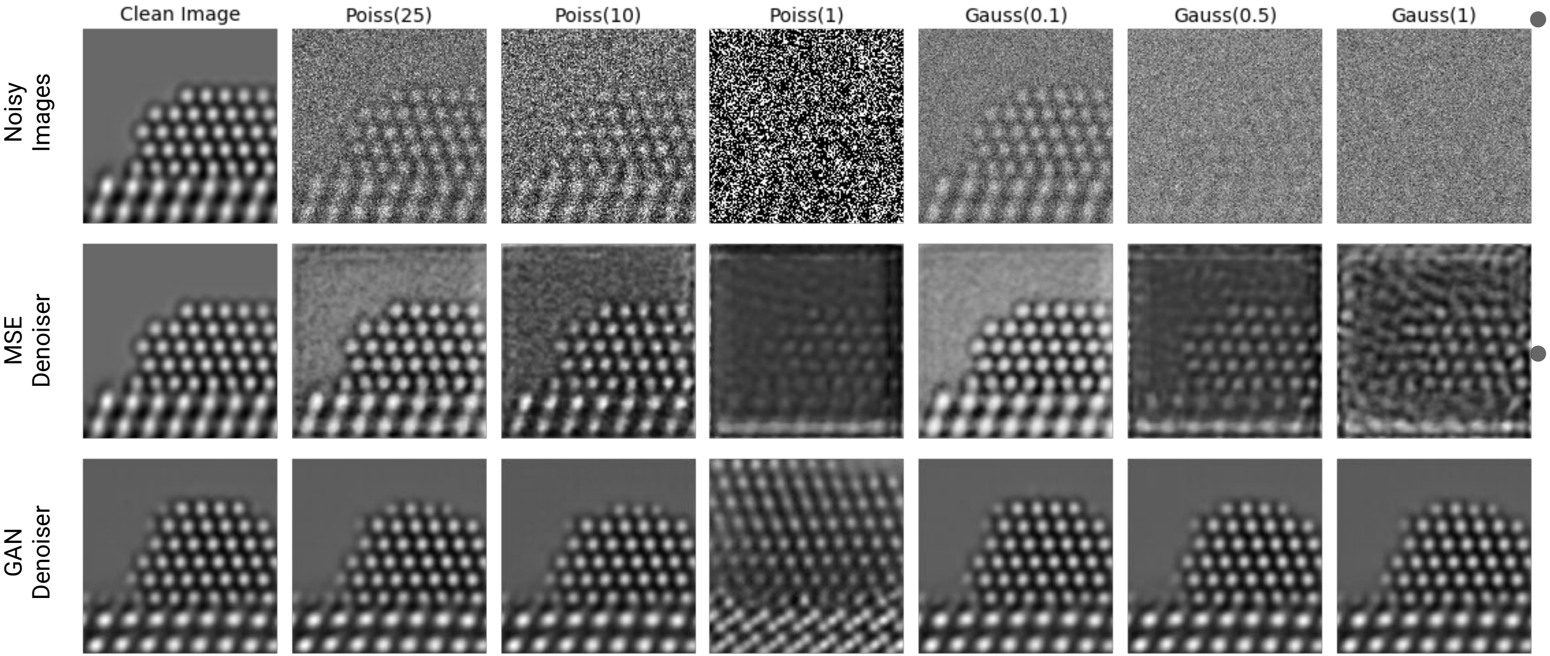

- **Benchmark**: The best supervised MSE denoisers can achieve **PSNR > 30** for Poiss(1)
- GAN denoiser works very well, except in the most extreme noise setting

Electron Microscope Data

Electron Microscope images

- 40 actual images (1200x1200 pixels)
- 20,000 simulated images (850x850 pixels)

Broader project context

- Microscope imaging seeks to capture catalytic processes
- Images are of 3D structures, goal is to predict 3D structure from 2D images

Electron microscope Images

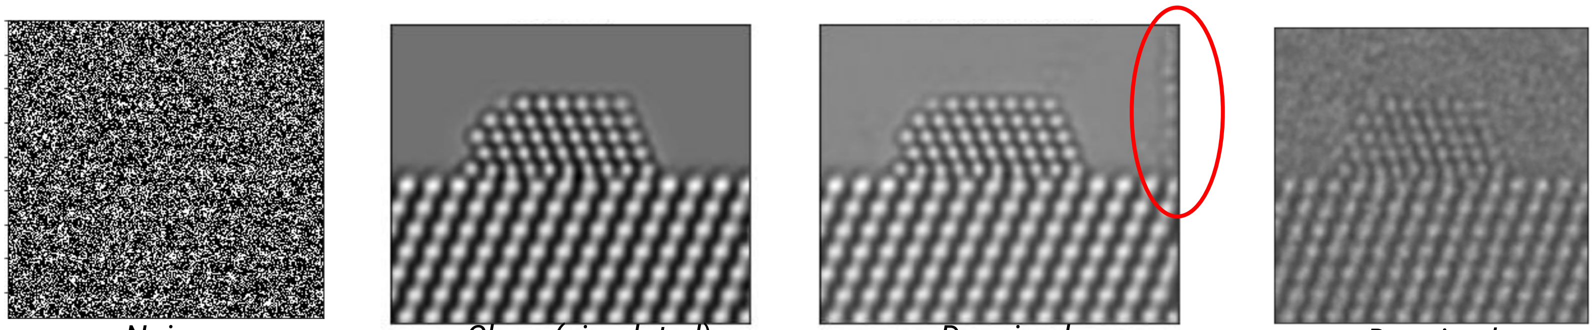

Noisy *Clean (simulated)* *Denoised* *Denoised*

GAN optimizer approach

- To validate that the optimization approach works, we first tested on MNIST and CIFAR100
  - 32x32 images, gray and RGB
  - Trained a DCGAN for 500 epochs

- We evaluated a range of **optimizers**, **learning rates**, and **schedulers**
  - Adam with learning rate 1e-2 consistently worked best

- Two optimizer tests
  - **Test 1:** Generate an image with the GAN, try to recover the **same latent vector** from a random initialization
  - **Test 2:** Select random image NOT in GAN, try to recover the latent vector mapped to the **nearest image** in the range of GAN

CIFAR 100 with color

Test 1: Recover image generated by GAN

True image

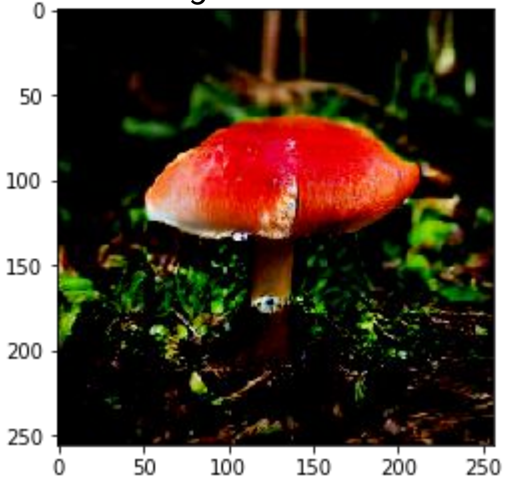

Initialized image

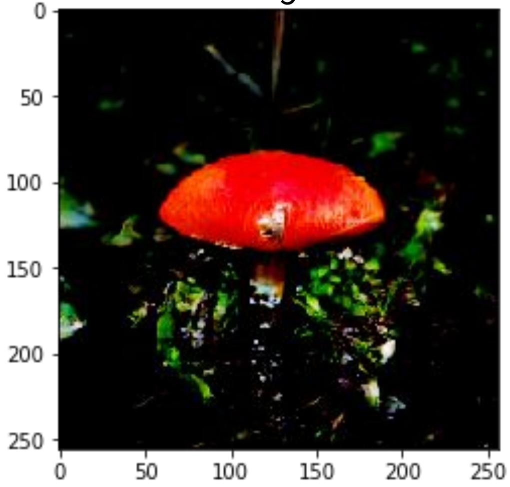

Initial Reconstruction Loss: ~1100

Predicted image

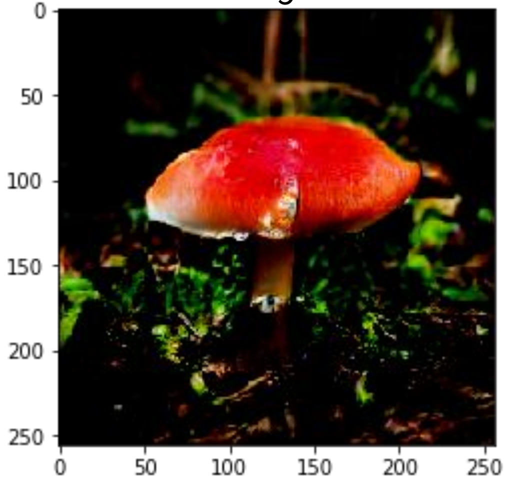

Reconstruction Loss: ~3.4

CIFAR 100 with color

Test 2: Approximate image not generated by GAN

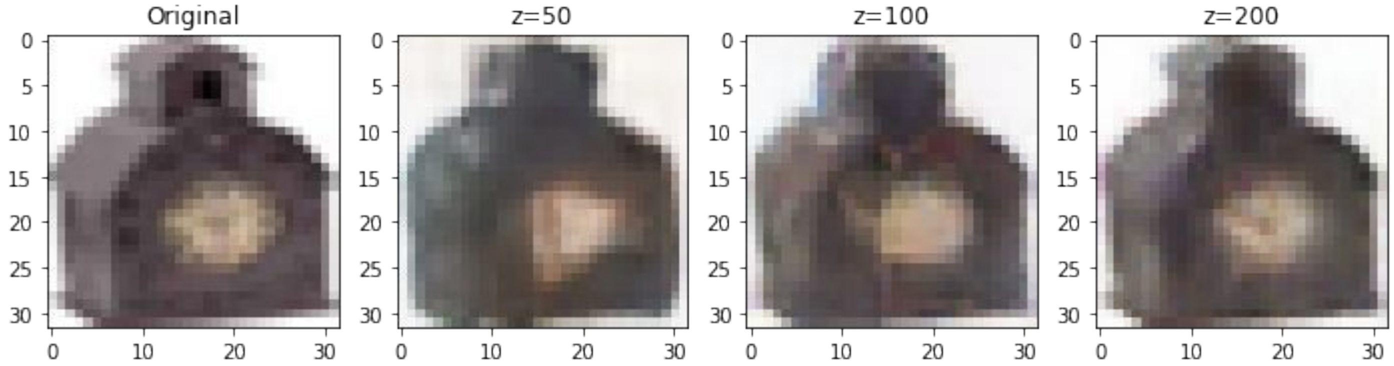

MNIST 32x32 Gray

Test 2: Approximate image not generated by GAN

True image

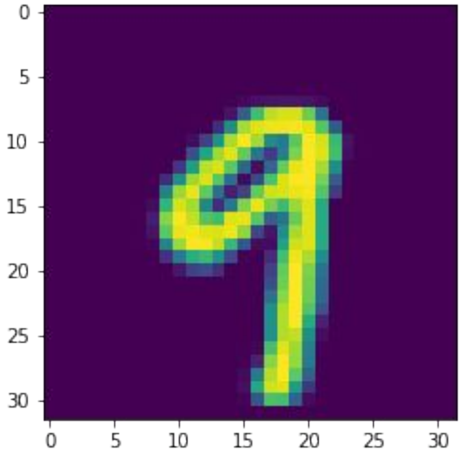

Initialized image

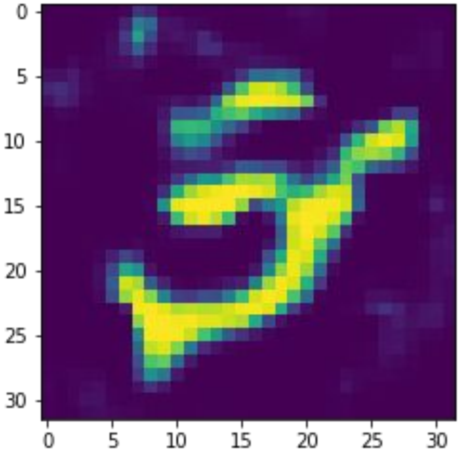

Predicted image

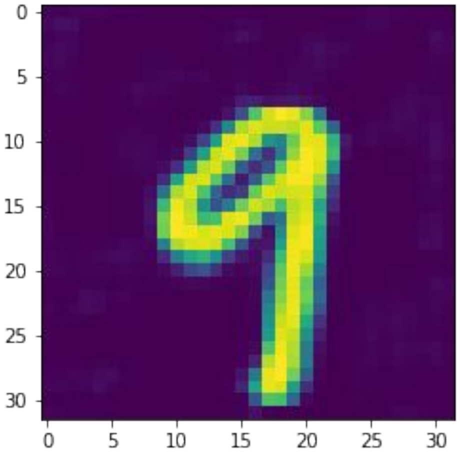

Reconstruction Loss < 1e-2

CIFAR100 32x32 Gray

Test 2: Approximate image not generated by GAN

True image

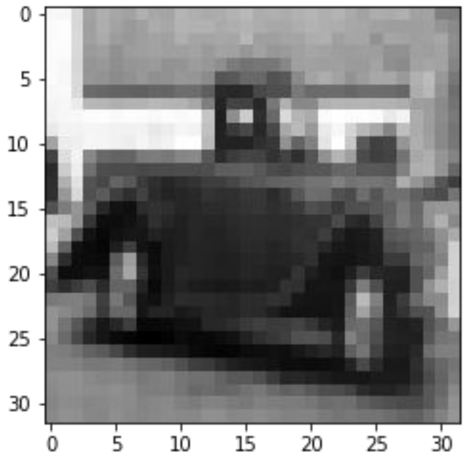

Initialized image

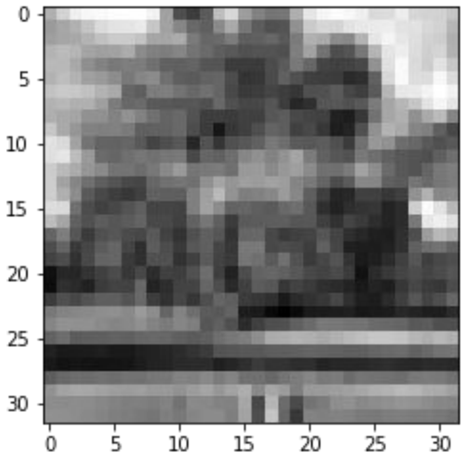

Predicted image

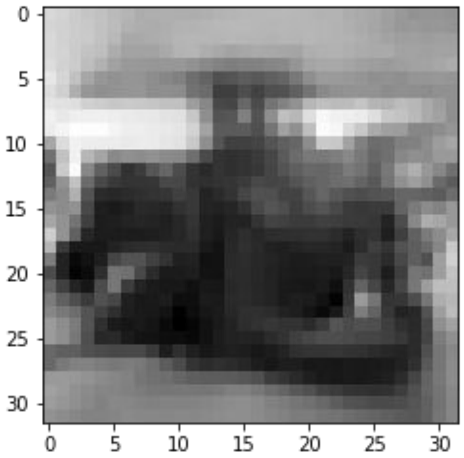

Reconstruction Loss ~25