maaGMA: Modified-Adversarial-Autoencoding Generator with Multiple Adversaries
Optimizing Multiple Learning Objectives for Image Generation with GANs

Sahil Chopra, Ryan Holmdahl

CS 231n: Convolutional Neural Networks for Visual Recognition *Mentor: Shayne Longpre*

Overview

New formulations of GAN loss, e.g. WGAN, are purported to be more stable

Goal: Can we leverage these more stable loss functions to train a GAN with multiple adversaries, i.e. optimize multiple learning objectives?

Data & Evaluation

Experiment: Do we see improvement in image generation between baseline Adversarial Autoencoder and proposed maaGMA architecture on two tasks?

Generative Tasks:

1. Handwritten Digit Construction
2. Emotive Face Construction

MNIST Handwritten Digits Database

- 60,000 examples
- 28 pixel x 28 pixel (Grayscale)
- 10 Digit Classes

ICML 2013 Facial Recognition Dataset

- 32,000 examples
- 48 pixel x 48 pixel (Grayscale)
- 7 Emotion Classes

Evaluation:

- Qualitative Inspection

Loss Formulation

$$L_{dis.} = \min_{G} \max_{D \in \mathcal{D}} \mathrm{E}_{x \sim \mathrm{P}_r}[D(x)] - \mathrm{E}_{\hat{x} \sim \mathrm{P}_g}[D(\hat{x})]$$

$\mathcal{D}$ = Set of 1-Lipschitz Functions

P_G = Distr. Implicitly Defined by $\hat{x} = G(z), z \sim p(z)$

Objective: Minimize E-M Distance between P_R & P_G

$$Loss_{generator} = \lambda_1 * L_{recon.} - \lambda_2 * L_{dis(gaussian)} - \lambda_3 * L_{dis(img)}$$

Architecture

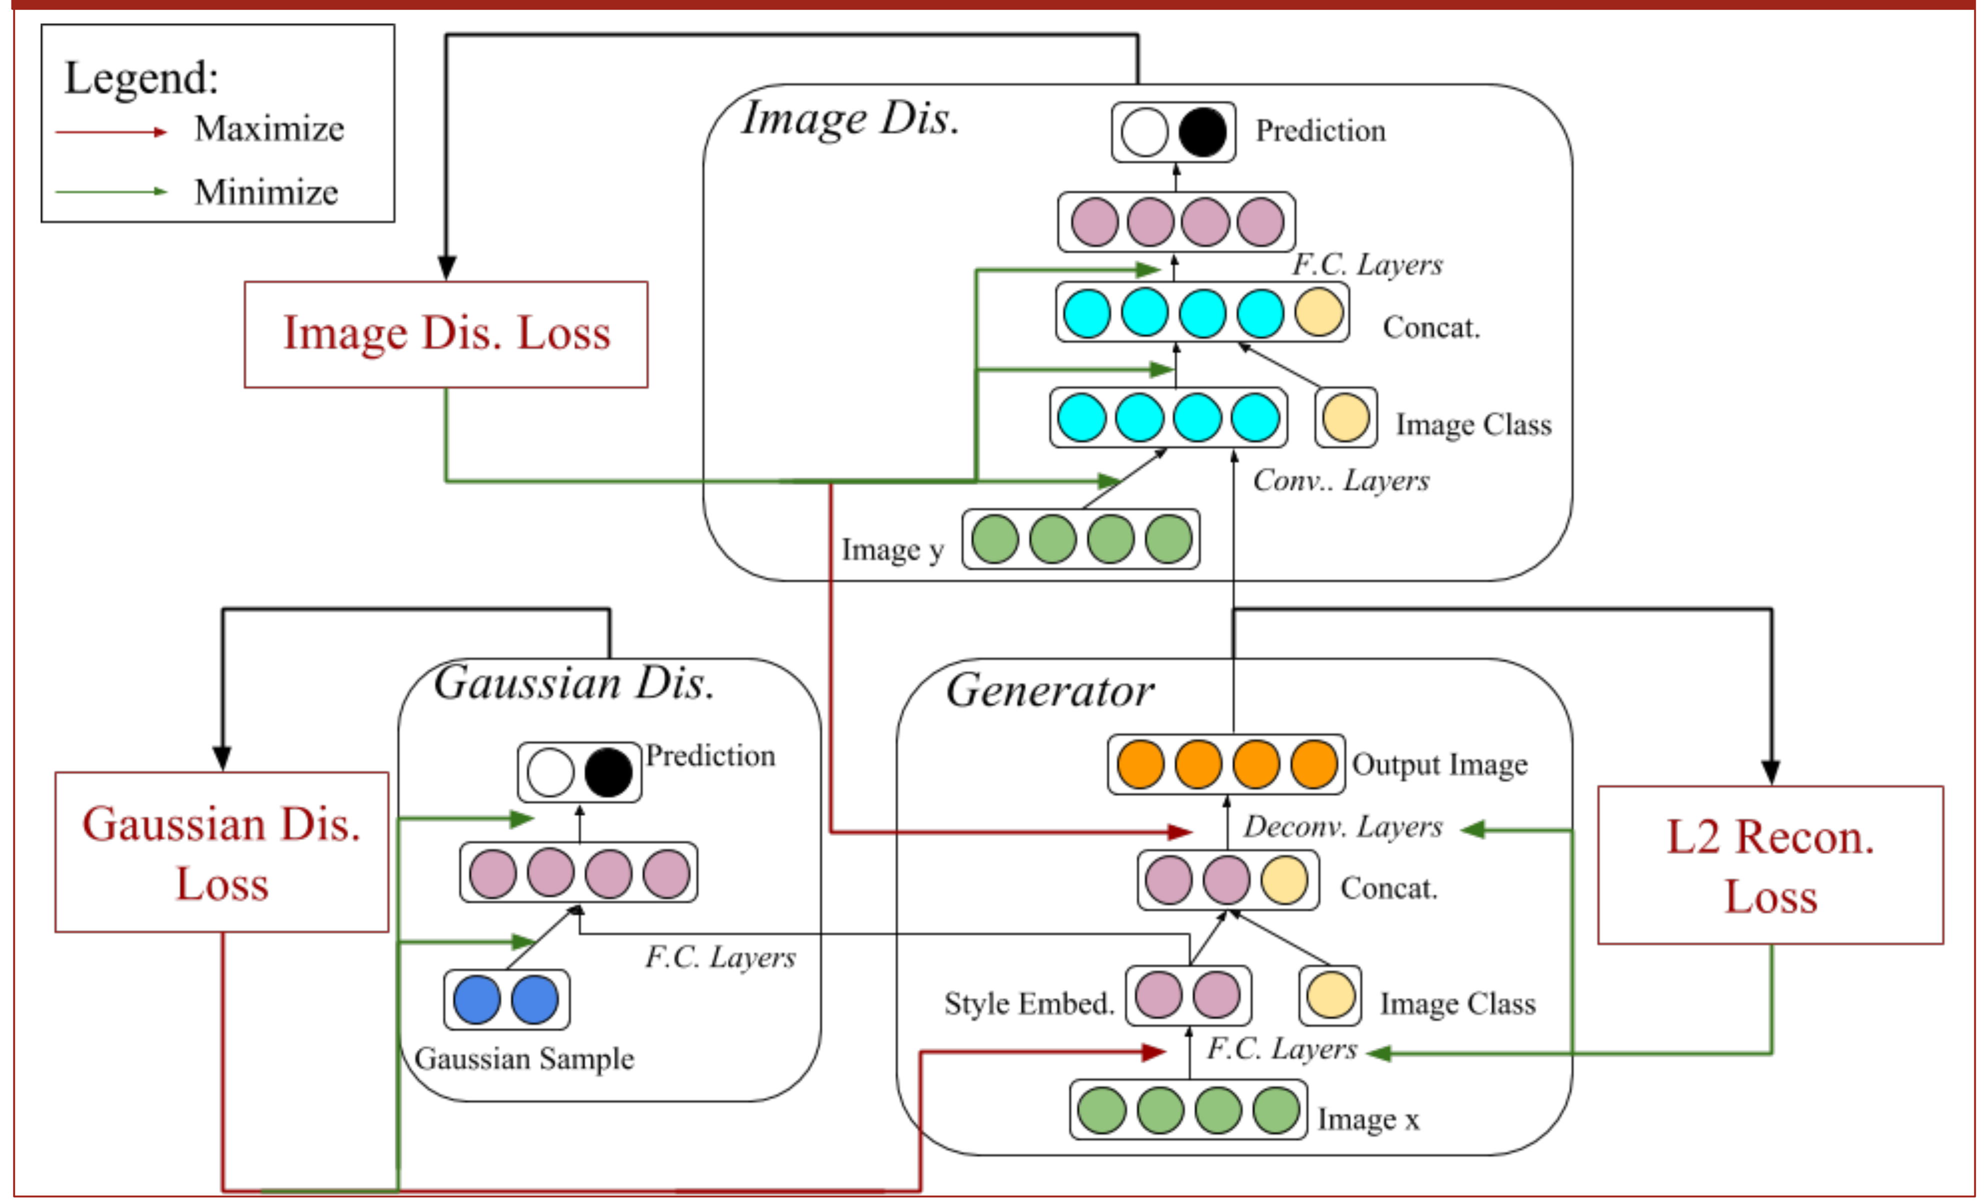

MNIST Results

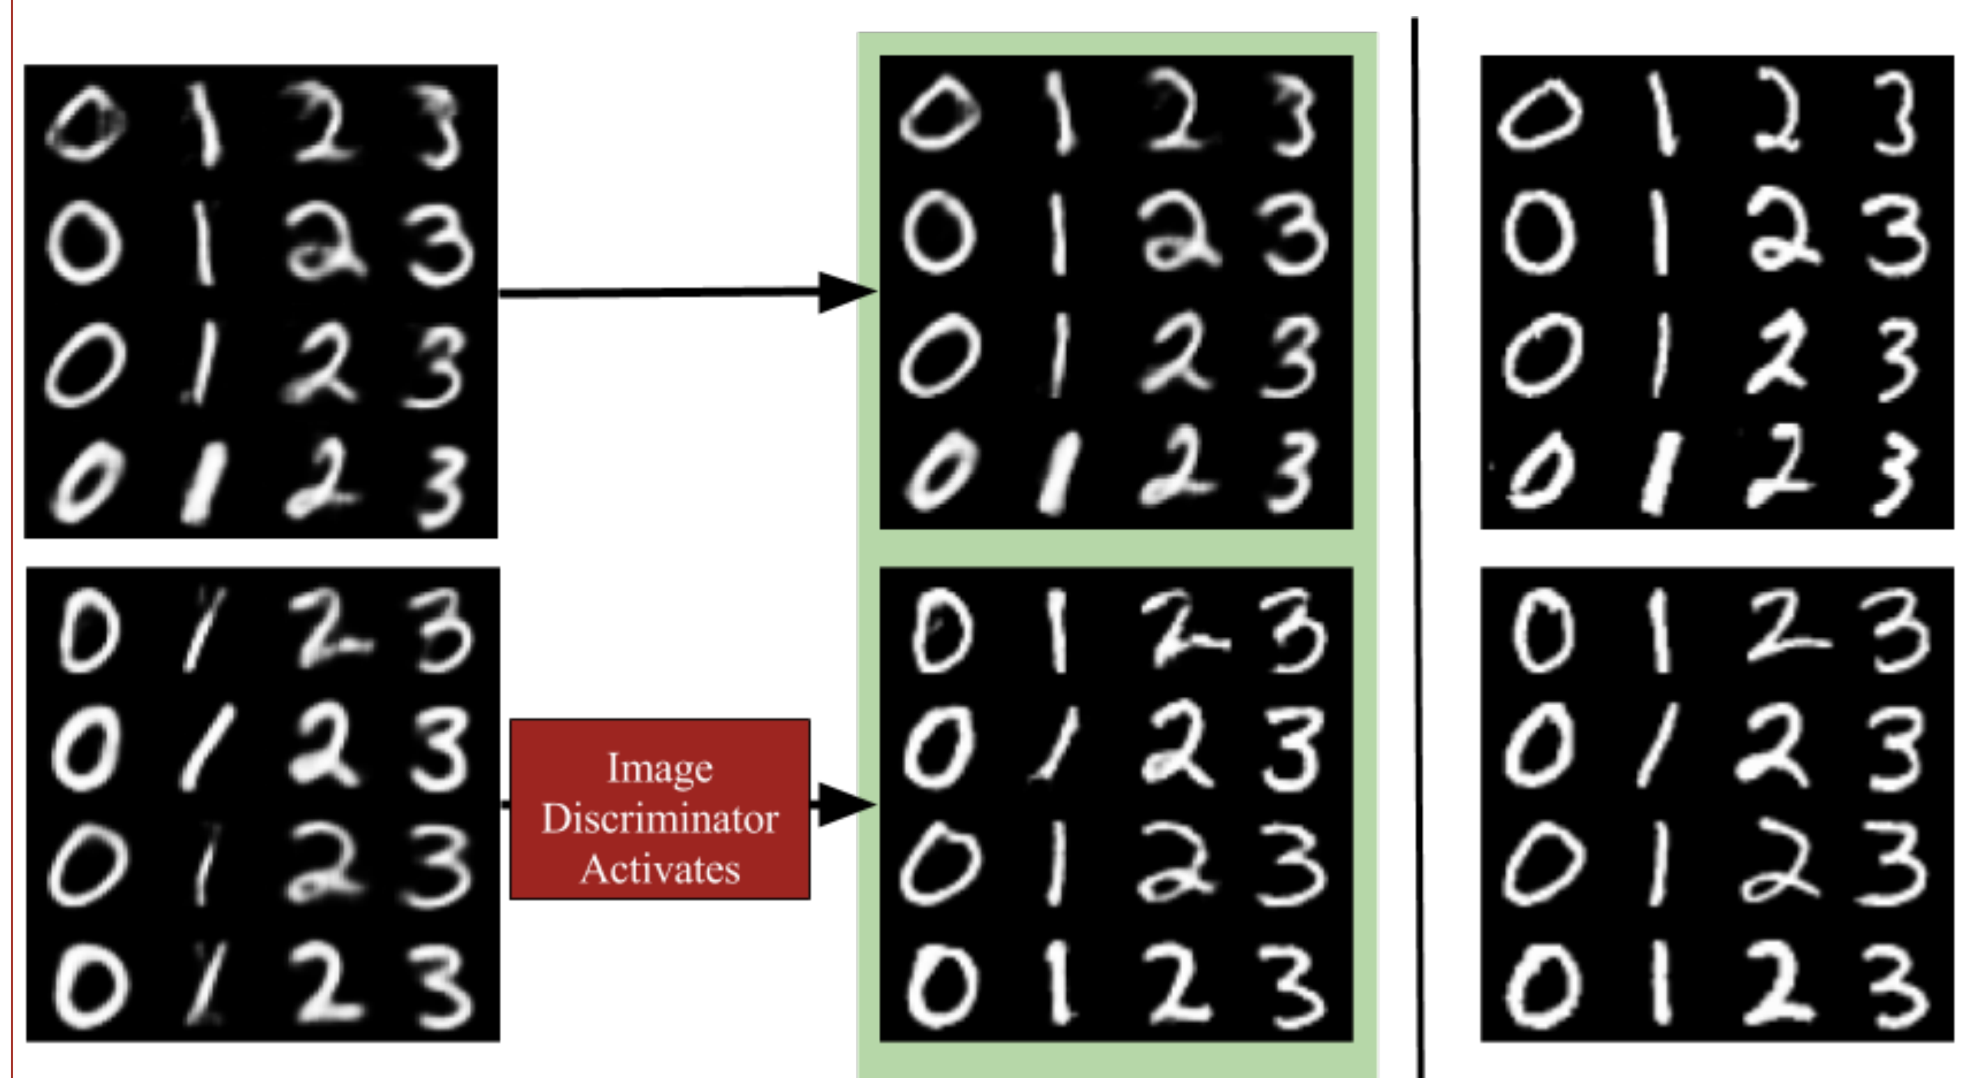

Epoch 300 Epoch 500 Nearest Train Set Samples

Top Row: Adversarial Autoencoder (Baseline)
Bottom Row: maaGMA (Proposed Architecture)

MNIST Results

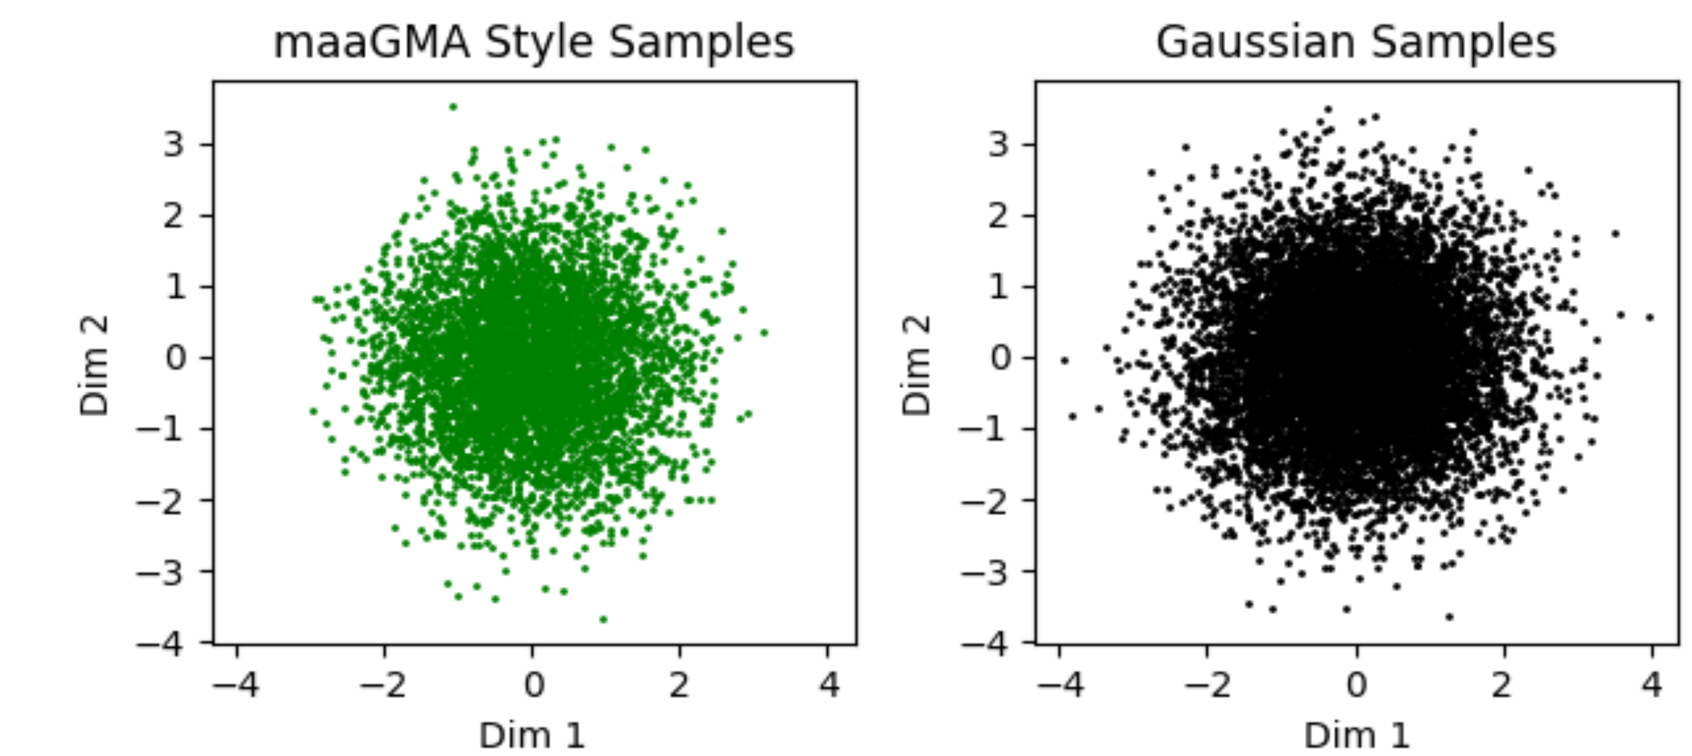

maaGMA successfully forces style embeddings to desired distribution

Preliminary ICML 2013 Results

Epoch 200 Epoch 220

Discussion & Future Work

- Image Dis. activated after *N* epochs to prevent early memorization
- maaGMA employs "indirect competition" amongst adversaries
  - Gen. Decoder confuses Image Dis.
  - Gen. Encoder. confuses Gaussian Dis.
  - Entire Gen. minimizes L2 loss, accommodating the confusions
- maaGMA qualitatively outperforms baseline, sharpening images and bypassing limitations of L2 loss
- Additional hyperparameter tuning required for quality face generation, but currently improving clarity over baseline
- Future Work: Place adversaries in direct competition over the same variables. Explore better methods of finding ideal style dim. Compute reconstruction loss with different formulation.