

GENERATIVE ADVERSARIAL NETWORKS


Rajeswari Sivakumar
Deep Learning @UGA
2/2/17

MOTIVATIONS


- Many applications.
- Reinforcement Learning
 - Learn different generative models to simulate environment and future actions of an agent
- Improve training in semi-supervised learning
- Multi-modal outputs
- Realistic generation
 - Super-resolution of images
 - Video frame prediction

GENERATIVE MODELS

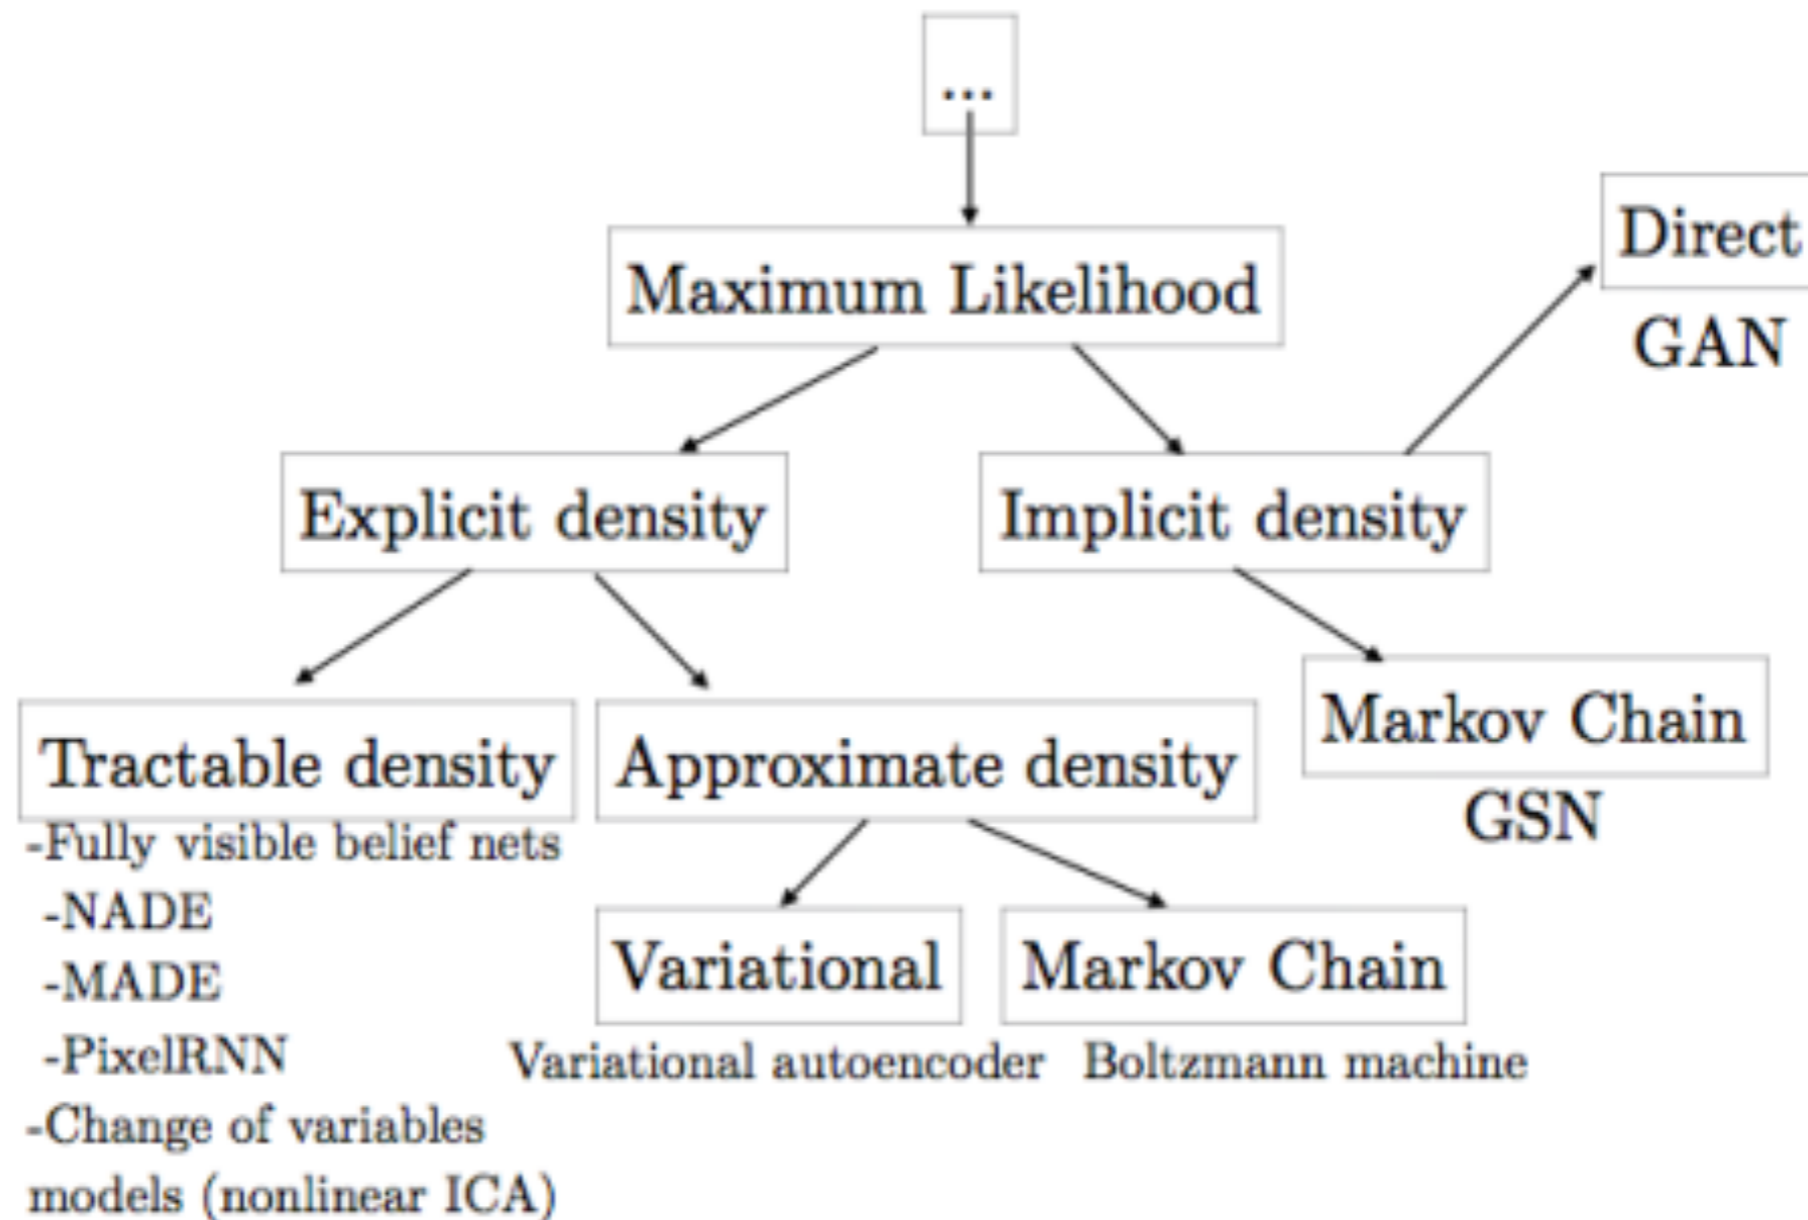
- Traditionally trying to learn distributions
 - i.e. clustering algorithms
- Alternatively can be used to generate samples that resemble real data.



$$\theta^* = \arg \max_{\theta} \mathbb{E}_{x \sim p_{\text{data}}} \log p_{\text{model}}(x | \theta)$$

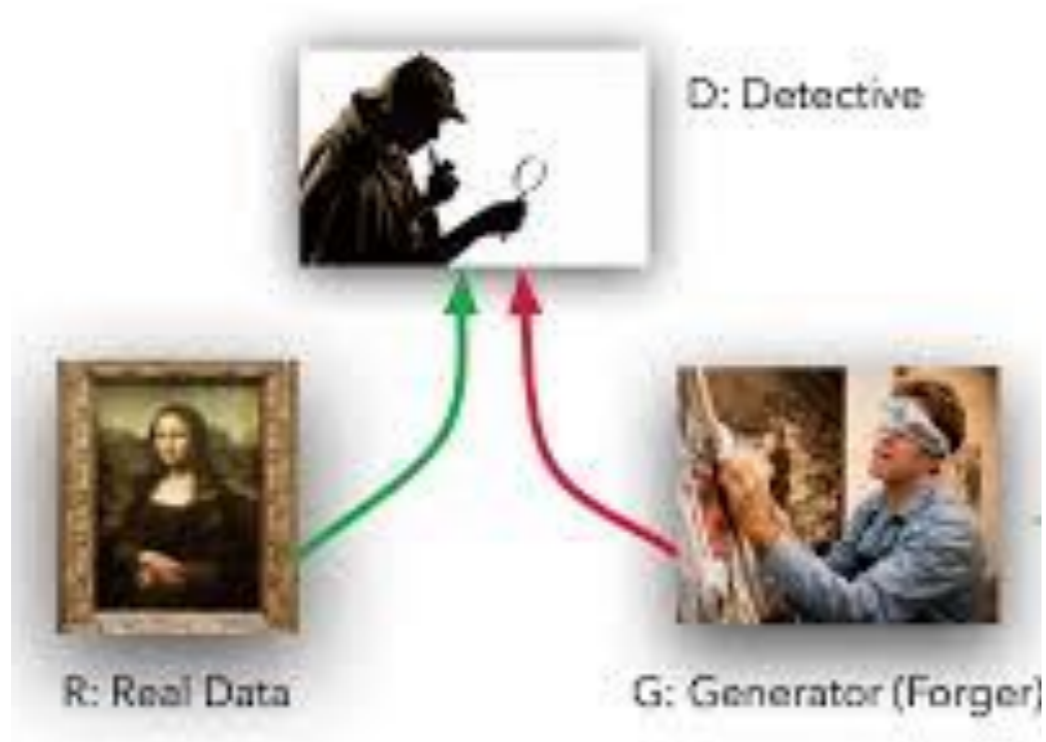


FINDING MAXIMUM LIKELIHOOD



WHAT ARE GANS?

- Two part generative network
- Generative
 - Tries to correctly represent distribution of features from data you are training on
- Discriminative
 - Tries to correctly differentiate an example from the generator from an example from the data



THE COUNTERFEITER AND THE DETECTIVE

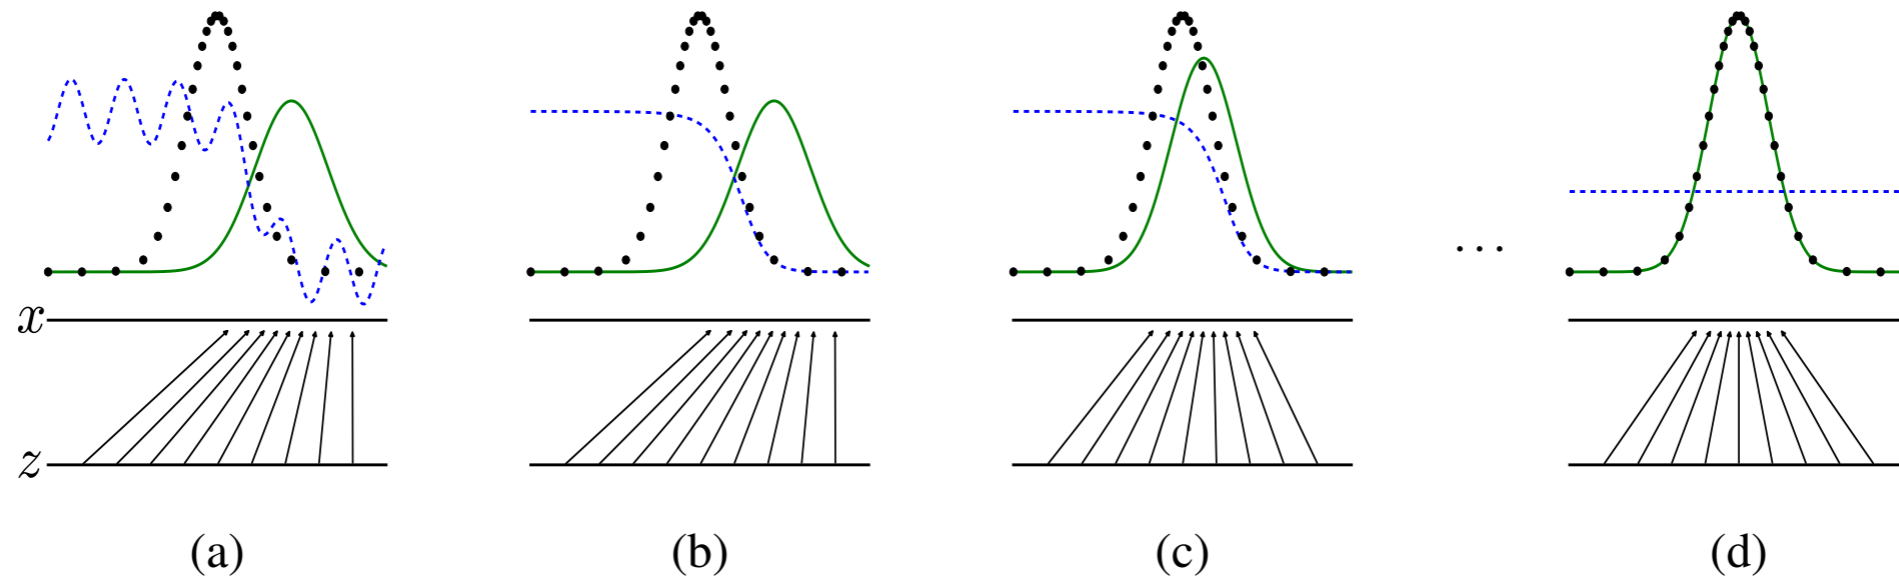
.....

- Think of generator as a counterfeiter and discriminator
- Generator makes fakes to pass off as real
- Discriminator has to tell the two apart
- Both are always trying to one-up each other.



OPTIMIZING A VALUE FUNCTION

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$



THEORETICAL

Algorithm 1 Minibatch stochastic gradient descent training of generative adversarial nets. The number of steps to apply to the discriminator, k , is a hyperparameter. We used $k = 1$, the least expensive option, in our experiments.

for number of training iterations **do**

for k steps **do**

- Sample minibatch of m noise samples $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(m)}\}$ from noise prior $p_g(\mathbf{z})$.
- Sample minibatch of m examples $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$ from data generating distribution $p_{\text{data}}(\mathbf{x})$.
- Update the discriminator by ascending its stochastic gradient:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m \left[\log D(\mathbf{x}^{(i)}) + \log \left(1 - D(G(\mathbf{z}^{(i)})) \right) \right].$$

end for

- Sample minibatch of m noise samples $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(m)}\}$ from noise prior $p_g(\mathbf{z})$.
- Update the generator by descending its stochastic gradient:

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log \left(1 - D(G(\mathbf{z}^{(i)})) \right).$$

end for

The gradient-based updates can use any standard gradient-based learning rule. We used momentum in our experiments.

OPTIMAL SOLUTIONS

$$D_G^*(\mathbf{x}) = \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + p_g(\mathbf{x})}$$

$$\begin{aligned} C(G) &= \max_D V(G, D) \\ &= \mathbb{E}_{\mathbf{x} \sim p_{data}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - D_G^*(G(\mathbf{z})))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{data}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_g} [\log(1 - D_G^*(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{data}} \left[\log \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_{data}(\mathbf{x}) + p_g(\mathbf{x})} \right] \end{aligned}$$

$$p_g = p_{data}, D_G^*(\mathbf{x}) = \frac{1}{2},$$

$$C(G) = -\log(4) + KL\left(p_{data} \parallel \frac{p_{data} + p_g}{2}\right) + KL\left(p_g \parallel \frac{p_{data} + p_g}{2}\right)$$

$$C^* = -\log(4)$$

- We can define an optimal discriminator for a given generator
- Can think of the solution to the as maximizing the value function
- For optimal generator, the probability distribution matches the data
- Thus a discriminator should be equally likely to say that a given x is data or generated



EXPERIMENTS

-
- ▶ Trained on MNIST, Toronto Face Database, and CIFAR-10

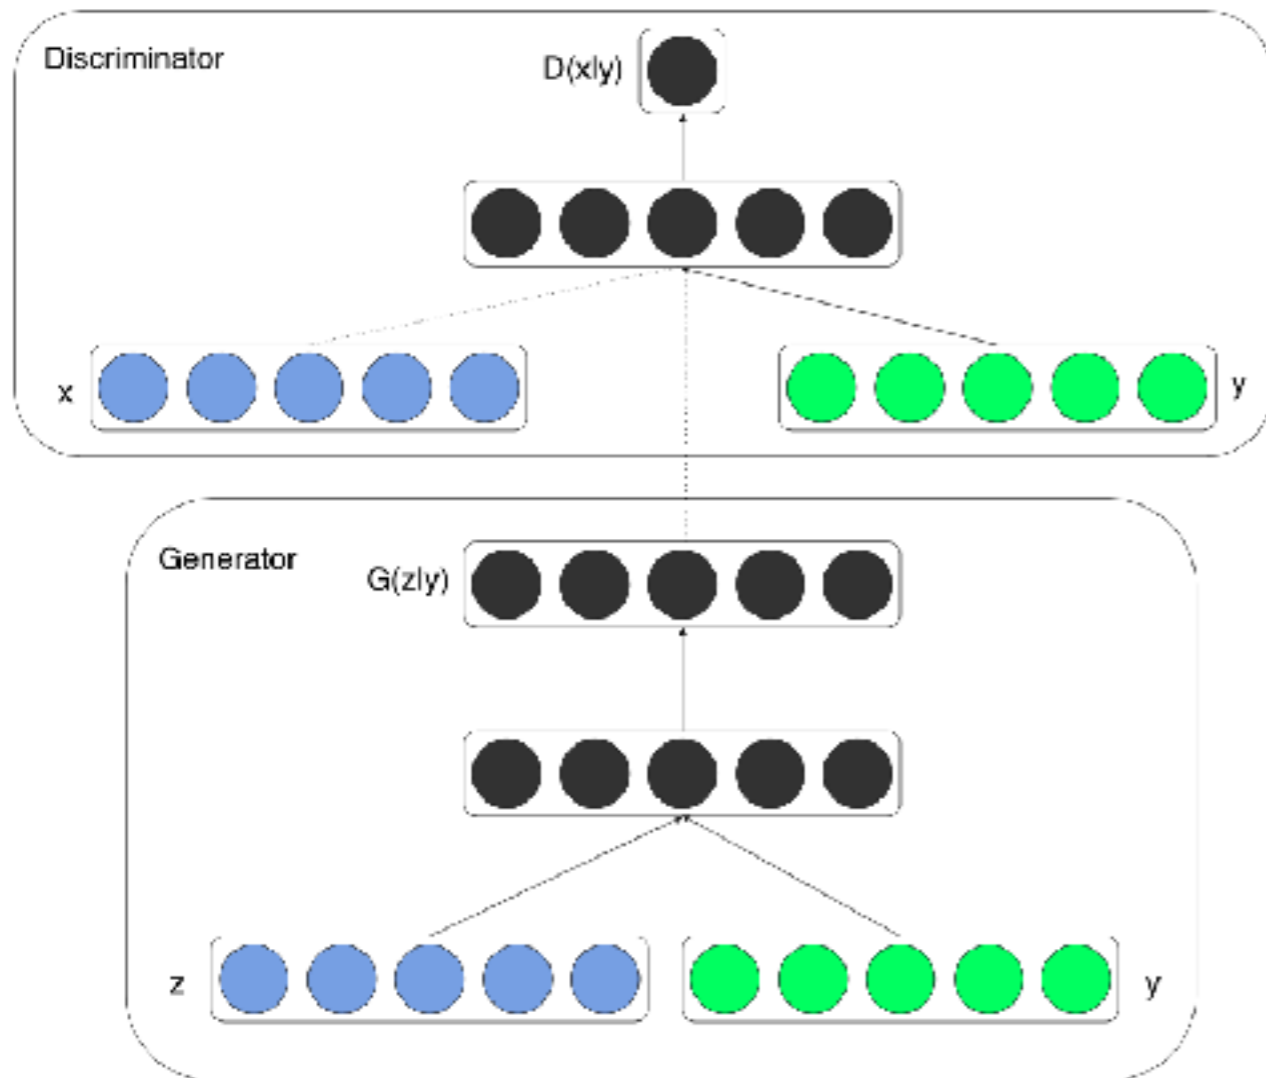
Model	MNIST	TFD
DBN [3]	138 ± 2	1909 ± 66
Stacked CAE [3]	121 ± 1.6	2110 ± 50
Deep GSN [6]	214 ± 1.1	1890 ± 29
Adversarial nets	225 ± 2	2057 ± 26

APPLICATIONS AND IMPROVEMENTS

- Conditional generative model
 - Conditional Generative Adversarial Networks (CGAN)
- Semi-supervised learning
 - Improving training efficiency

C-GAN

- Mimic GAN, but with added complexity of a given prior event.
- Image tagging
- Avoiding mode collapse



C-GAN

.....

- Mimic GAN, but with added complexity of a given prior event.
- Image tagging
- Avoiding mode collapse

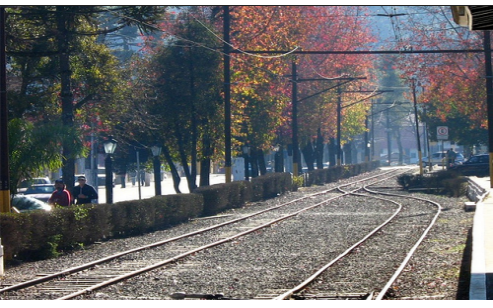

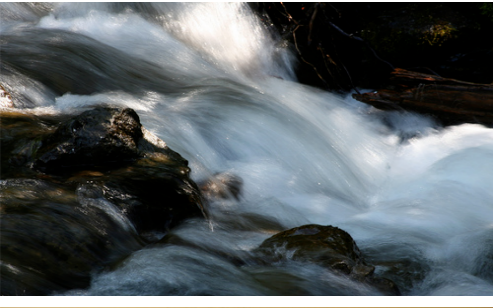

	User tags + annotations	Generated tags
	<p>montanha, trem, inverno, frio, people, male, plant life, tree, structures, transport, car</p>	<p>taxi, passenger, line, transportation, railway station, passengers, railways, signals, rail, rails</p>
	<p>food, raspberry, delicious, homemade</p>	<p>chicken, fattening, cooked, peanut, cream, cookie, house made, bread, biscuit, bakes</p>
	<p>water, river</p>	<p>creek, lake, along, near, river, rocky, treeline, valley, woods, waters</p>
	<p>people, portrait, female, baby, indoor</p>	<p>love, people, posing, girl, young, strangers, pretty, women, happy, life</p>

Table 2: Samples of generated tags

C-GAN

.....

- Mimic GAN, but with added complexity of a given prior event.
- Image tagging
- Avoiding mode collapse




Figure 2: Generated MNIST digits, each row conditioned on one label

SEMI-SUPERVISED GAN

- ▶ Train generative network and classifier at the same time
- ▶ By doing so:
 - ▶ Cut down on training time and improve accuracy
- ▶ Useful when there isn't a lot of data for training.



Figure 1. Output samples from SGAN and GAN after 2 MNIST epochs. SGAN is on the left and GAN is on the right.

Table 1. Classifier Accuracy

EXAMPLES	CNN	SGAN
1000	0.965	0.964
100	0.895	0.928
50	0.859	0.883
25	0.750	0.802

QUESTIONS?