

CreativeAI: Deep Learning for Computer Graphics

Unsupervised Learning in CG

Niloy Mitra

UCL/Adobe

Iasonas Kokkinos

UCL/Ariel AI

Paul Guerrero

UCL/Adobe

Vladimir Kim

Adobe

Nils Thuerey

TUM

Leonidas Guibas

Stanford/FAIR



Timetable

		Niloy	Iasonas	Paul	Nils	Leonidas
Introduction	9:00	X				
Neural Network Basics	~9:15		X			
Supervised Learning in CG	~9:50	X				
Unsupervised Learning in CG	~10:20			X		
Learning on Unstructured Data	~10:55					X
Learning for Simulation/Animation	~11:35				X	
Discussion	12:05	X	X	X	X	X



Unsupervised Learning

There is no direct ground truth for the quantity of interest


Focus on **generative** models:

- Variational Autoencoders (VAEs)
- Normalizing Flows
- Autoregressive Models (slides only)
- Generative Adversarial Networks (GANs)



Generative Models

- Assumption: the dataset are samples from an unknown distribution $p_{\text{data}}(x)$
- Goal: create a new sample from $p_{\text{data}}(x)$ that is not in the dataset



Dataset

Generated



Generative Models

- Assumption: the dataset are samples from an unknown distribution $p_{\text{data}}(x)$
- Goal: create a new sample from $p_{\text{data}}(x)$ that is not in the dataset

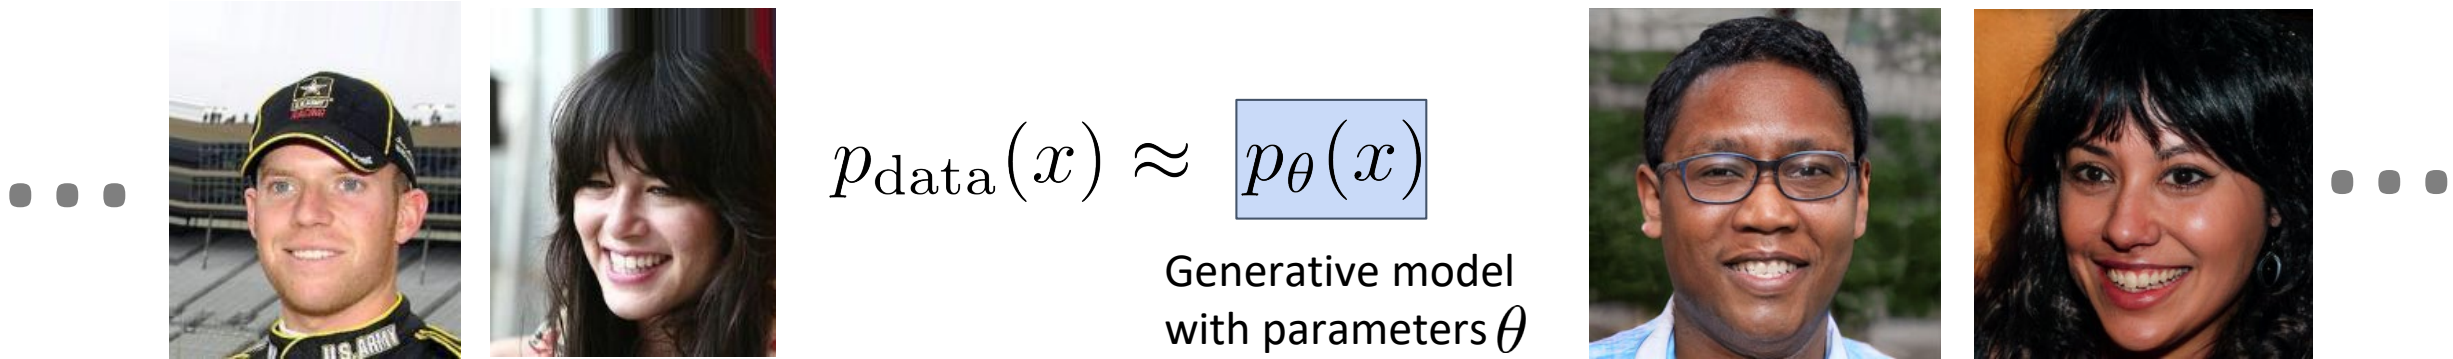


Dataset

Generated



Generative Models



How do we measure the similarity of $p_{\theta}(x)$ and $p_{\text{data}}(x)$?

Which model?



Generative Models

How do we measure the similarity of $p_{\theta}(x)$ and $p_{\text{data}}(x)$?

1) Likelihood of data
samples in $p_{\theta}(x)$

2) Adversarial game

Variational Autoencoders (VAEs)

**Generative Adversarial
Networks (GANs)**

Normalizing Flows


Autoregressive Models




Generative Models

How do we measure the similarity of $p_{\theta}(x)$ and $p_{\text{data}}(x)$?

1) Likelihood of data samples in $p_{\theta}(x)$




$p_{\text{data}}(x)$



$\approx KL(p_{\text{data}} \parallel p_{\theta})$

2) Adversarial game



$\approx JS(p_{\text{data}} \parallel p_{\theta})$



Likelihood-Based Models: Two Goals

Sample?
Evaluate?

- 1) Sample from the generative model
- 2) Evaluate the likelihood of a given sample in the model



$$x \sim p_{\theta}$$

Generative model
with parameters θ

$$p_{\theta}(x)$$

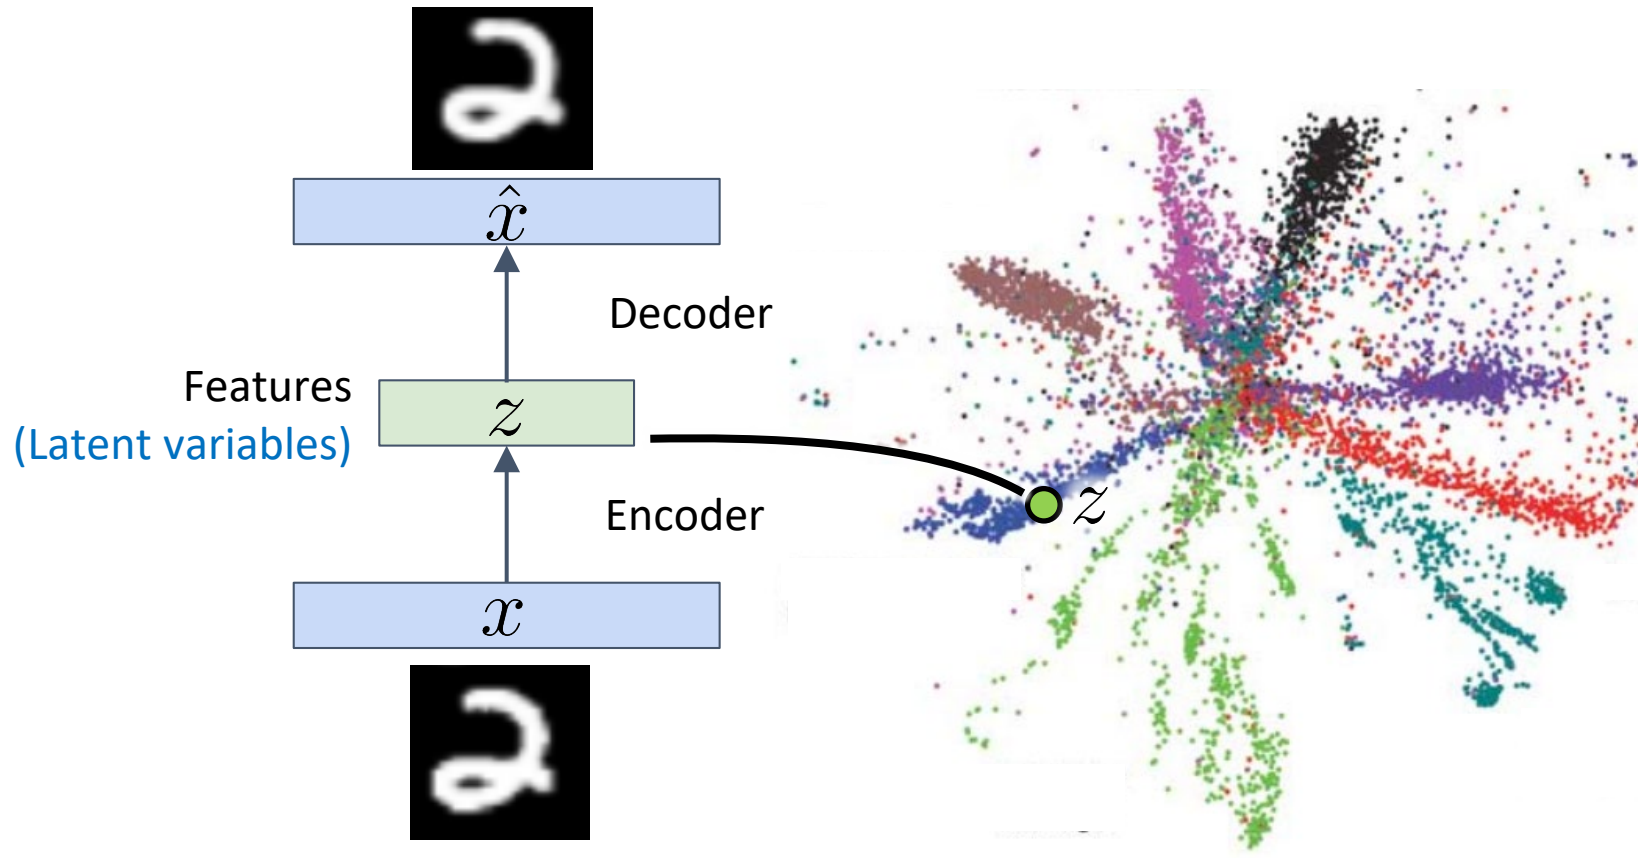
Generative model
with parameters θ

$$x$$



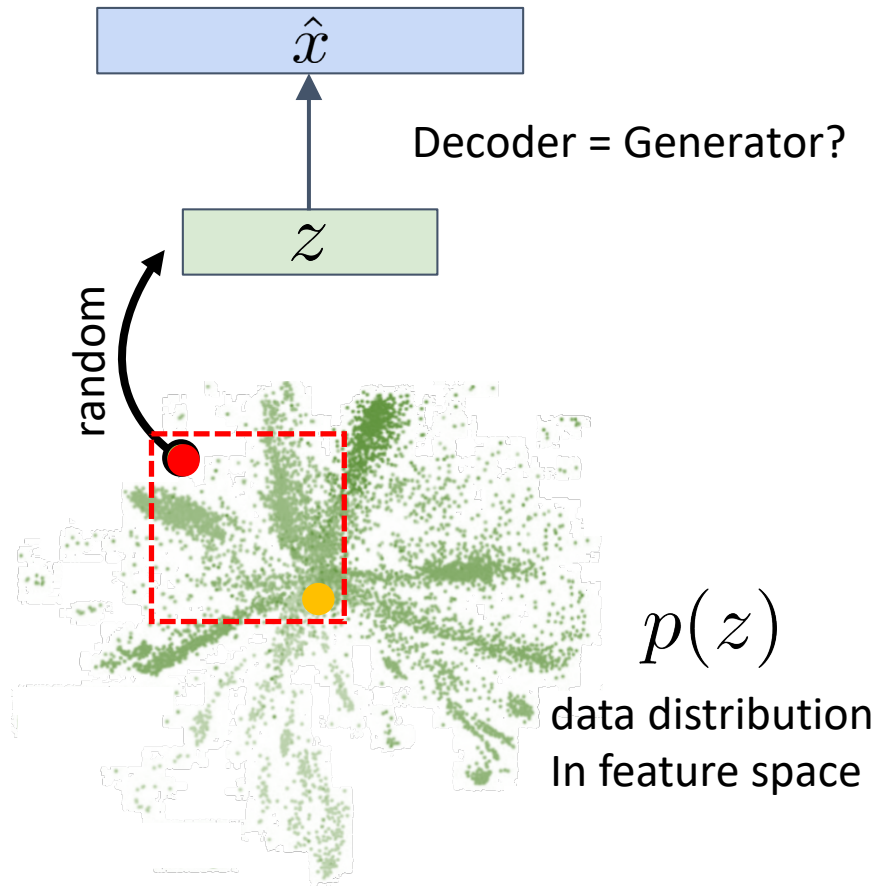

The Feature Space

Data distribution in 2D feature space (colors are class labels)



Autoencoders as Generative Models?

Sample?
Evaluate?




- Is a trained decoder a generative model?

- Can we generate a new sample $x \sim p_\theta$?



sample grid

random samples



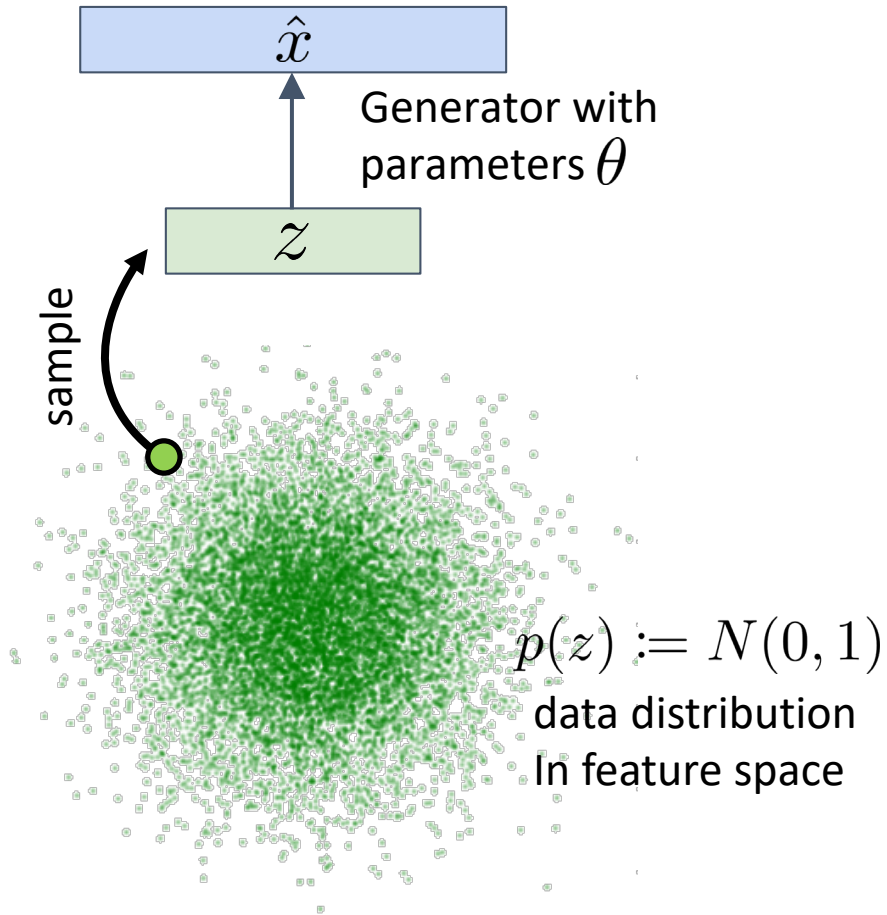
We do not know the distribution $p(z)$
in feature space

Latent Variable Model

Sample?
Evaluate?

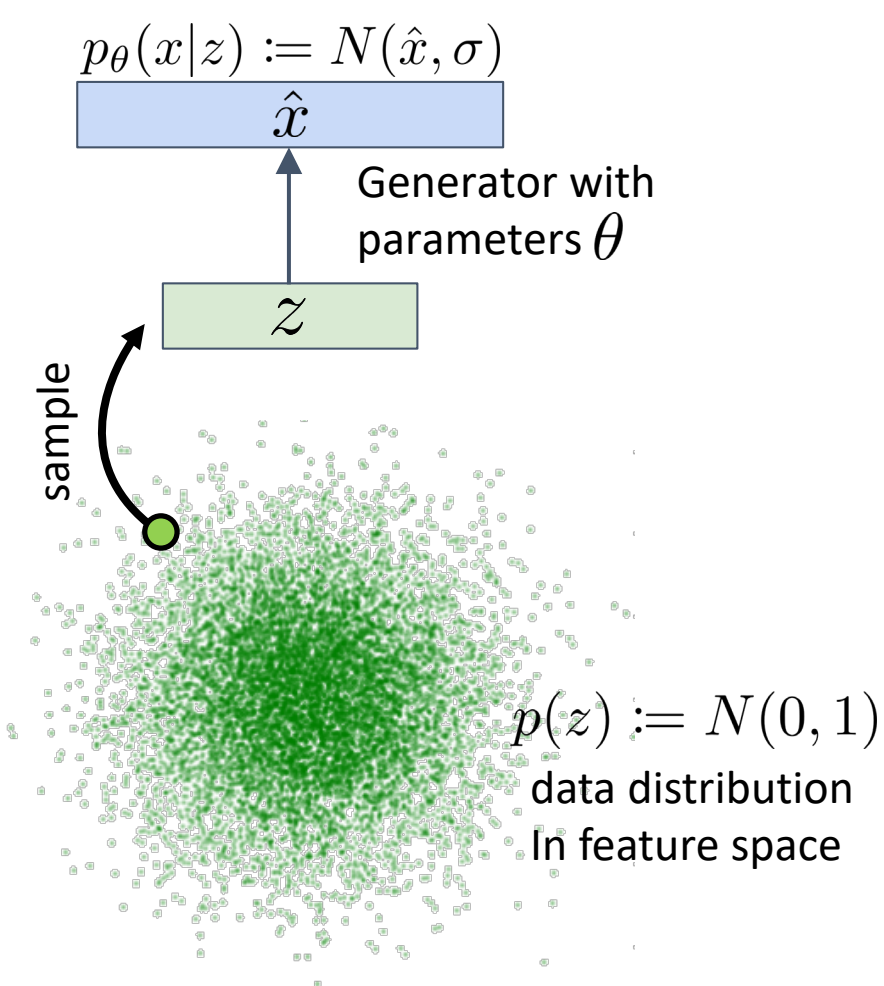


- Define $p(z)$ as a known distribution



Latent Variable Model

Sample?
Evaluate?



(Negative log-likelihood)

- Train generator with NLL of data as loss

$$-\log \sum_{x_i \in \text{data}} p_{\theta}(x_i)$$

- Can we compute the likelihood $p_{\theta}(x)$?

$$p_{\theta}(x) = \int p_{\theta}(x|z) p(z) dz$$



Latent Variable Model

Sample?
Evaluate?



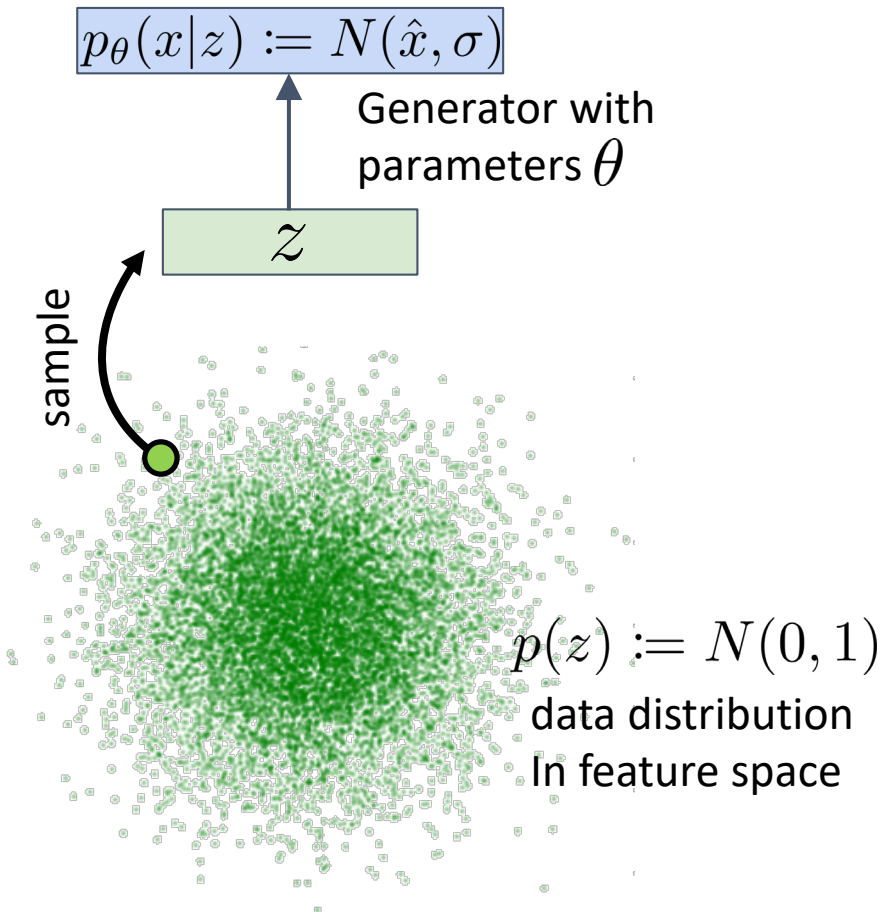
(Negative log-likelihood)

- Train generator with NLL of data as loss

$$-\log \sum_{x_i \in \text{data}} p_{\theta}(x_i)$$

- Can we compute the likelihood $p_{\theta}(x)$?

$$p_{\theta}(x) = \int p_{\theta}(x|z) p(z) dz$$



Latent Variable Model: Monte-Carlo

Sample?
Evaluate?

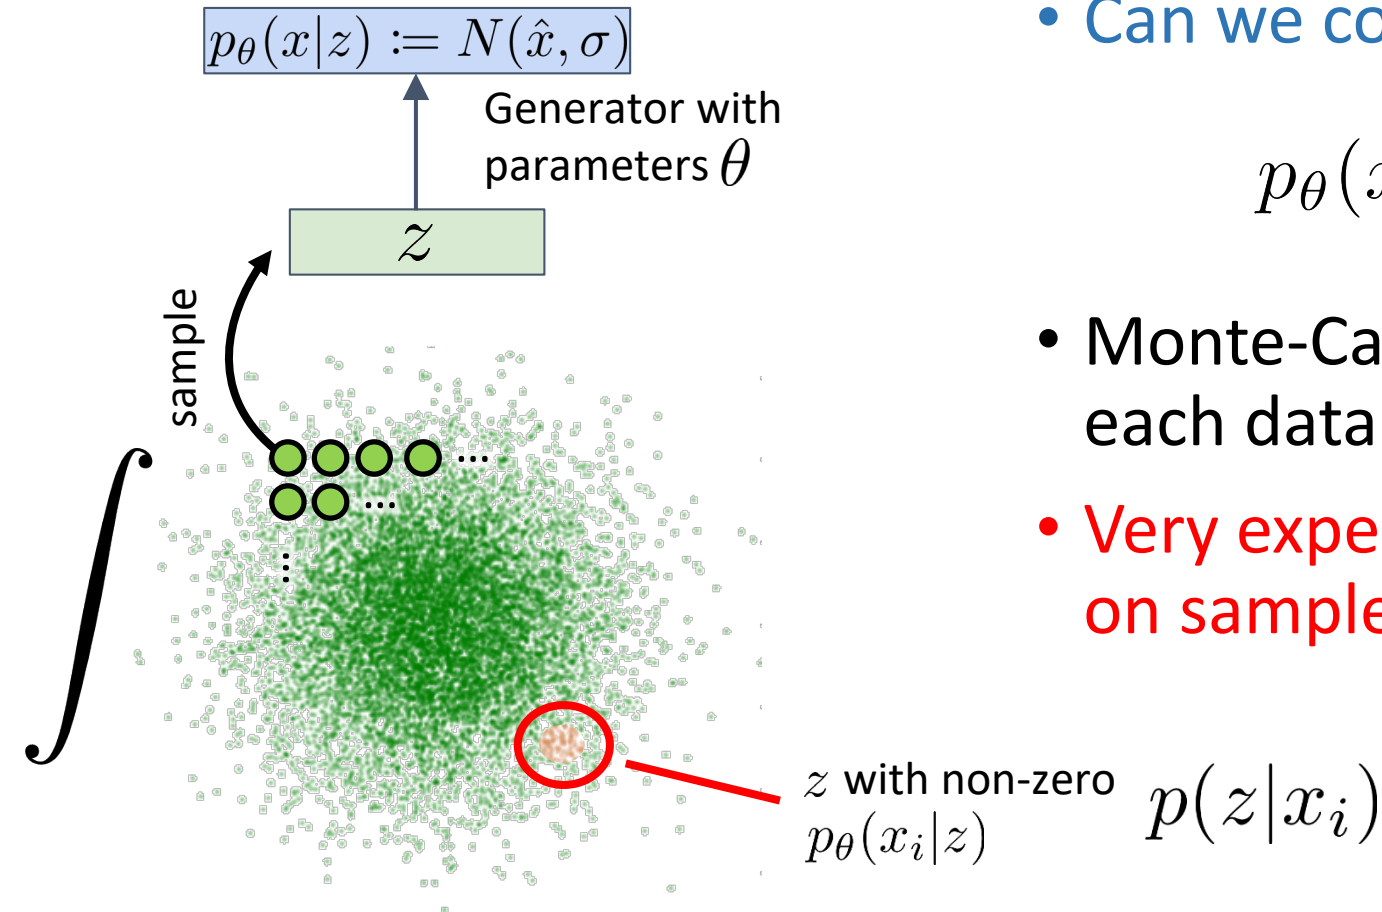


(Negative log-likelihood)

- Can we compute the likelihood $p_{\theta}(x)$?

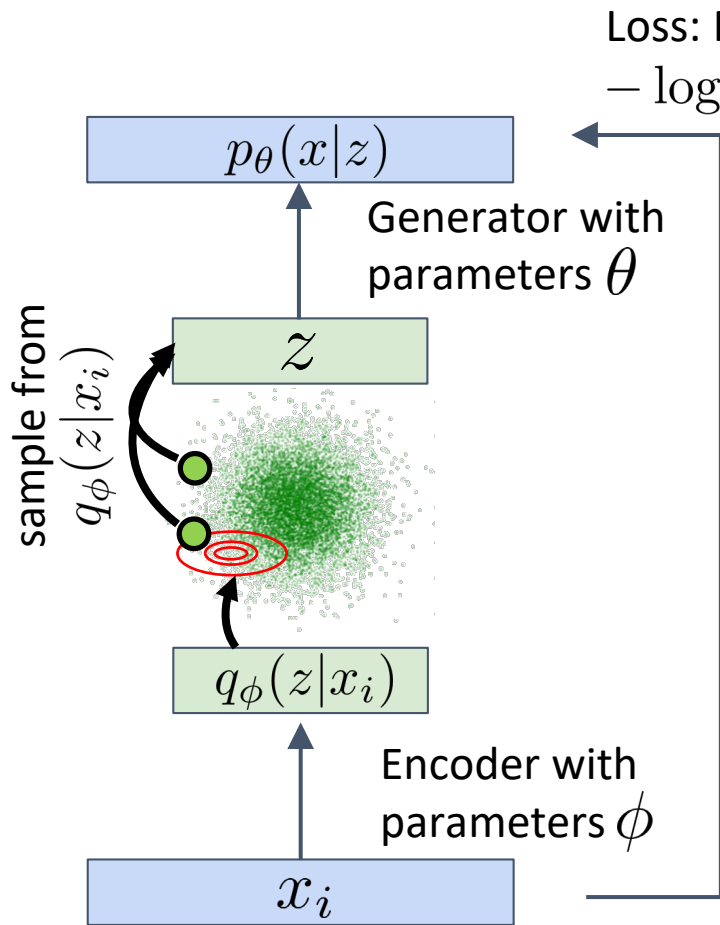
$$p_{\theta}(x) = \int p_{\theta}(x|z) p(z) dz$$

- Monte-Carlo integration to solve integral for each data sample
- Very expensive, or very inaccurate (depending on sample count)



Variational Autoencoders (VAEs): The Encoder

Sample?
Evaluate?



Loss: NLL of data

$$-\log p_\theta(x_i|z) \text{ with } z \sim q_\phi(z|x_i)$$

- During training, another network can learn to approximate the distribution $p(z|x_i)$
- $q_\phi(z|x_i)$ should be much smaller than $p(z)$
- Makes the computing the integral tractable

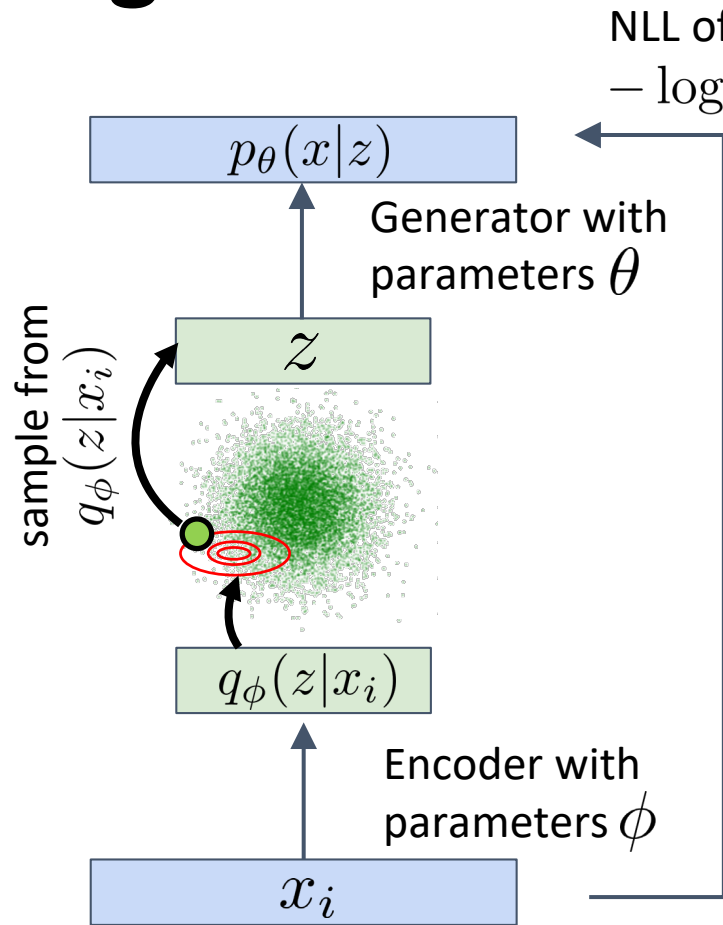
$$p_\theta(x) = \int p_\theta(x|z) p(z) dz$$

- Instead of integrating over all $p(z)$, integrate over $q_\phi(z|x_i)$ only
- A single random sample from $q_\phi(z|x_i)$ per iteration is usually enough



Variational Autoencoders (VAEs): Regularization

Sample?
Evaluate?




NLL of data as loss

$$-\log p_\theta(x_i|z) \text{ with } z \sim q_\phi(z|x_i) + KL(q_\phi(z|x_i) \parallel p(z))$$

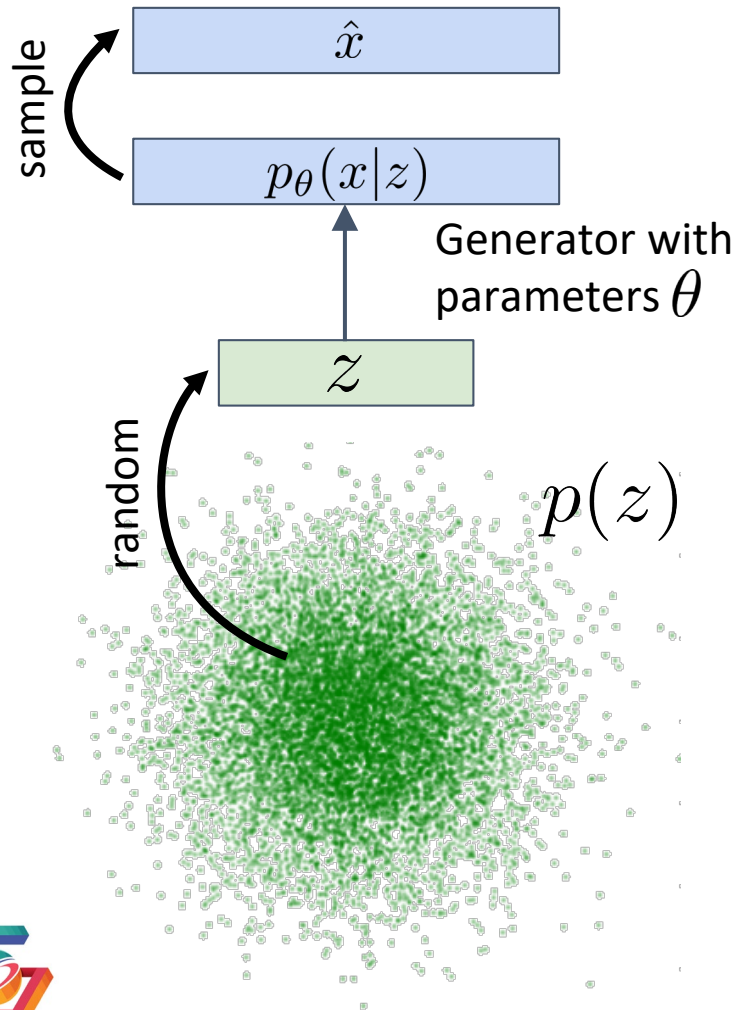
- The network can choose $q_\phi(z|x_i)$ freely
- But it is **regularized** it to approximate $p(z)$
- Neg. loss is a lower bound for the data's likelihood in the generated distribution



Generating Data



Generating Data



VAE




Autoencoder




Feature Space of Autoencoders vs. VAEs

Autoencoder



VAE



Summary: Variational Autoencoders (VAEs)

Positives

- Creates a feature space
- Relatively stable to train


Negatives

- Likelihood evaluation can only be approximated
- Projection of sample into feature space can only be approximated
- Regularization makes the results a bit blurry



Normalizing Flows

Sample?
Evaluate?



Variational Autoencoders (VAEs):

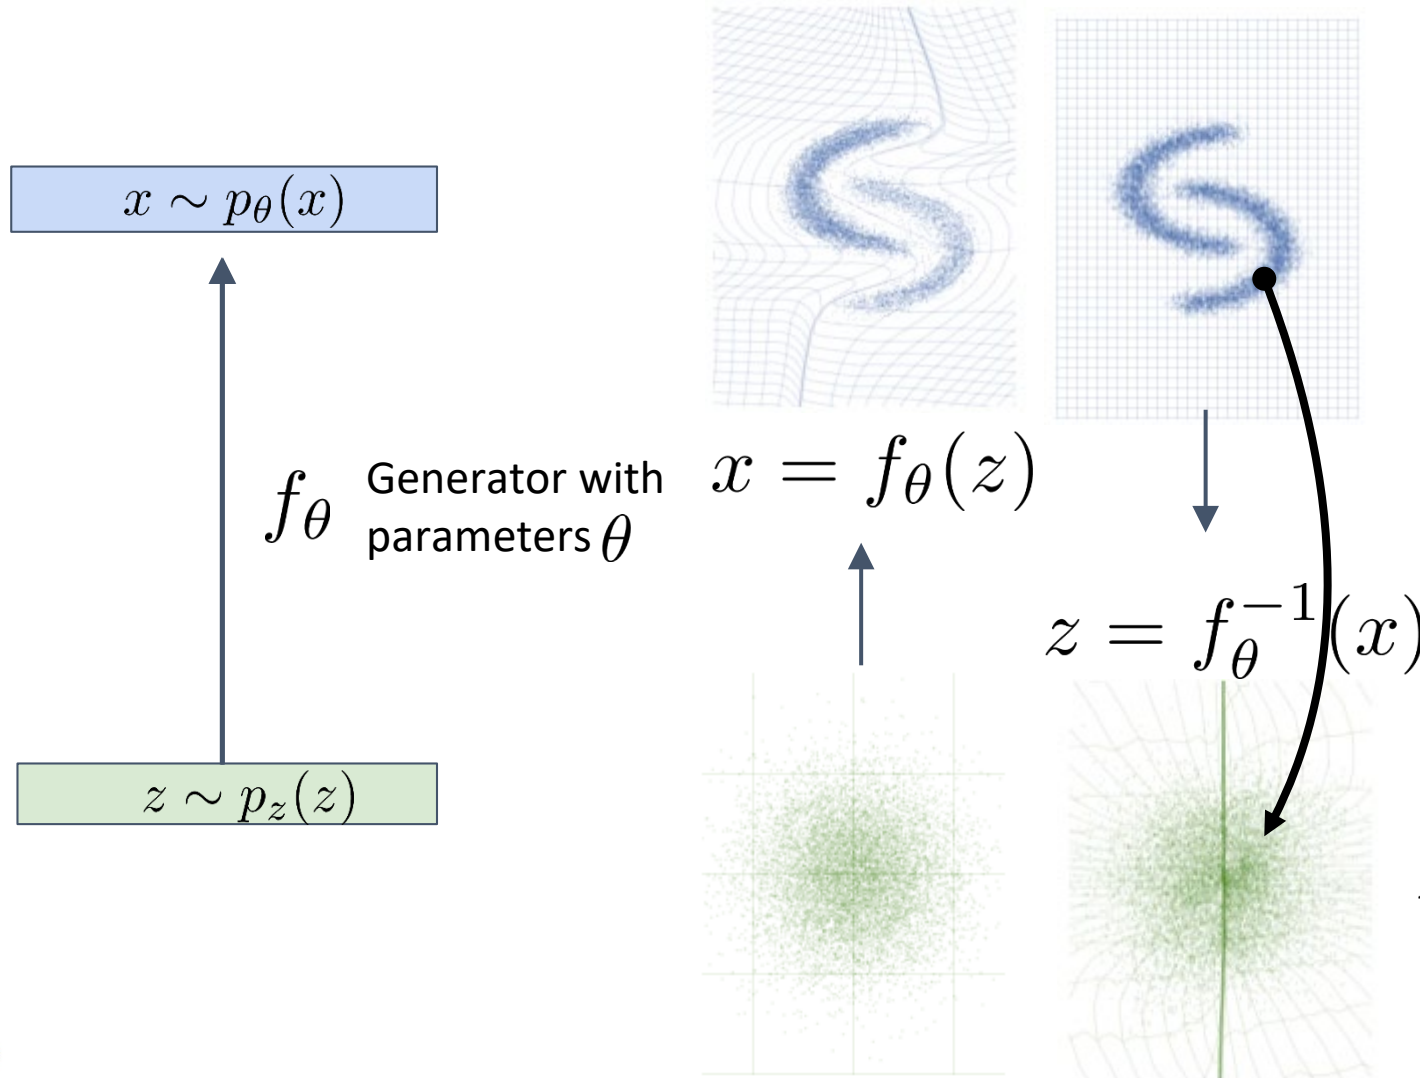
$$p_\theta(x) = \int p_\theta(x|z) p(z) dz$$

$p_z(z)$
(known)



Normalizing Flows

Sample?
Evaluate?



Normalizing Flows:

$$p_\theta(x) = p_z(f_\theta^{-1}(x))$$


... times the local density change caused by f_θ

$$\left| \det \left(\frac{\partial f_\theta^{-1}(x)}{\partial x} \right) \right|$$

$p_z(z)$
(known)

Normalizing Flows: Chaining

Sample?
Evaluate?



$$p_3(x_3) = p_z \left(f_1^{-1} \circ f_2^{-1} \circ f_3^{-1}(x_3) \right)$$

... times the local density change
caused by the chain of transformations

$$\prod_{i=1}^3 \left| \det \left(\frac{\partial f_i^{-1}}{\partial x_i} \right) \right|$$



Example: Glow

Invertible functions:

- Linear (1x1 conv) layer with weight matrix parameterized by its LU decomposition
- Affine coupling layer to propagate information between pixels



Training: 40 GPUs, 2 weeks




Example: PointFlow

Two flows: One to create the distribution of shape feature vectors, one for the distribution of points on a shape


Shape generation flow



Free shape generation



Shape interpolation



Summary: Normalizing Flows

Positives

- Creates a feature space
- Exact projection to feature space
- Exact likelihood evaluation


Negatives

- Only a limited set of functions (invertible and Jacobian easy to compute)
- Currently takes longer and needs more parameters than GANs for same quality




Autoregressive Models

- Create output step-by-step
- Each step depends on the output of all previous steps



Autoregressive Models

- Create output step-by-step
- Each step depends on the output of all previous steps



$$x = [x_1, x_2, \dots, x_n]$$

$$p_{\theta}(x) = p_{\theta}(x_1, x_2, \dots, x_n)$$



Chain rule of probability:

$$p_{\theta}(x) = \prod_{i=1}^n p_{\theta}(x_i | x_1, x_2, \dots, x_{i-1})$$

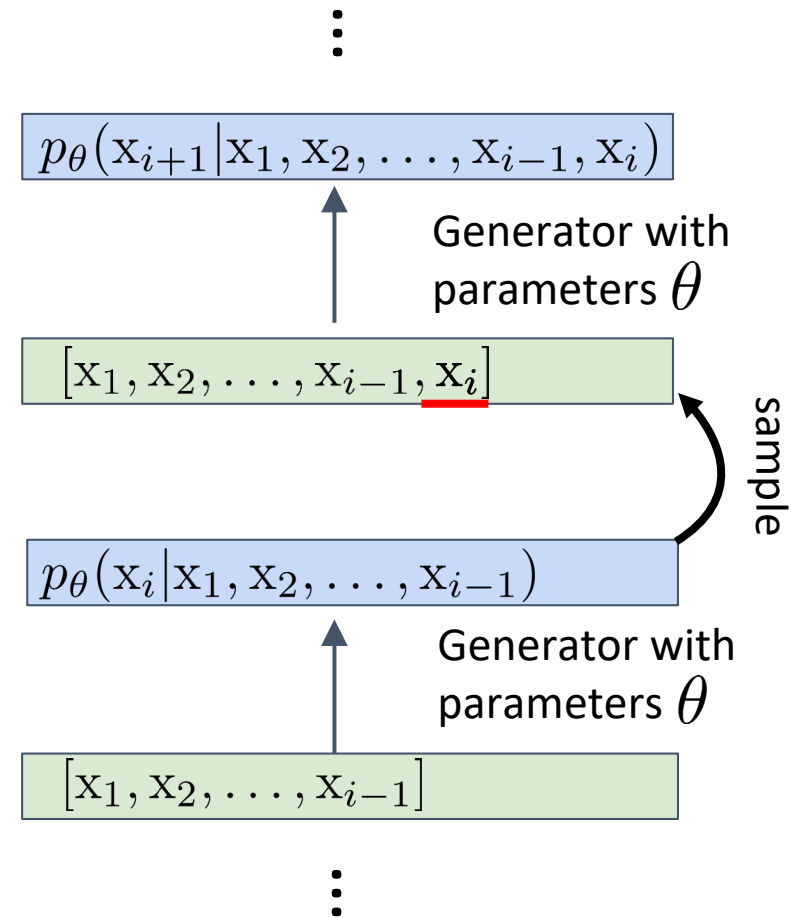


Autoregressive Models

- In each step, the model outputs a low-dimensional prob. distribution (e.g. over intensity values for one pixel)




Sample?
Evaluate?

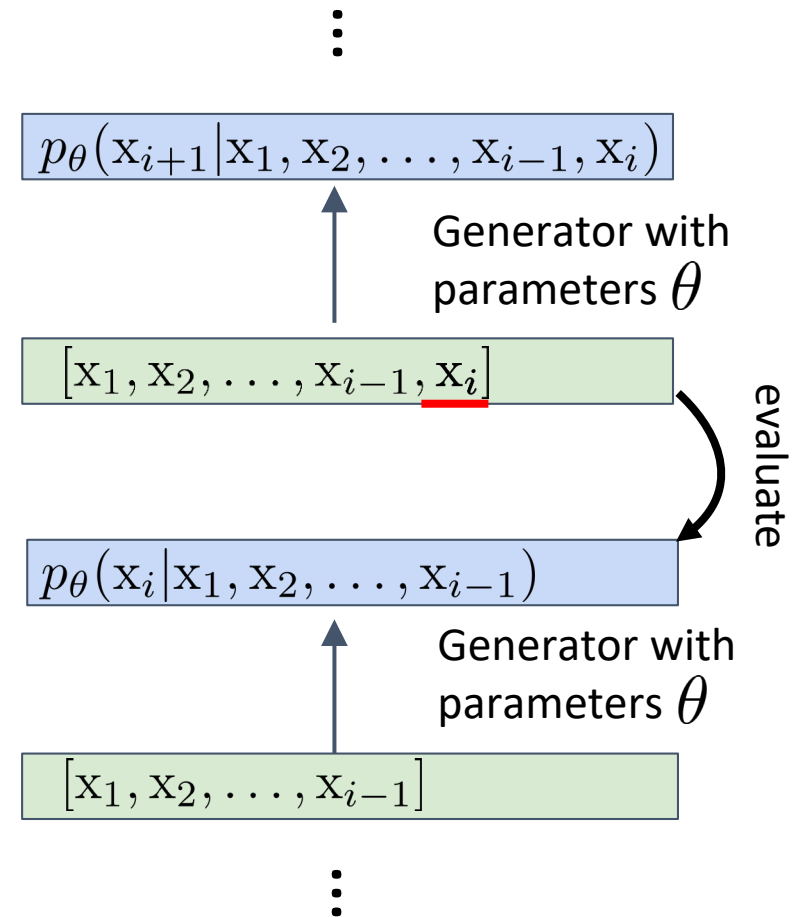


Autoregressive Models

- In each step, the model outputs a low-dimensional prob. distribution (e.g. over intensity values for one pixel)





Sample? 
Evaluate? 



Example: PixelRNN and PixelCNN


- Recursive network that has an **input** and a **state** (LSTM)
- Only recent steps are used as **input**, the **state** summarizes older steps



Sandbar



Lhasa Apso (dog)




Brown bear




Example: Graph Generation


GraphRNN




Learning Deep Generative Models of Graphs




Training Set




GraphRNN




Baseline



Training Set



New model



Baseline




Image Credit: Conditional Image Generation with PixelCNN Decoders, Oord et al.

GraphRNN: Generating Realistic Graphs with Deep Auto-regressive Models, You et al.

Learning Deep Generative Models of Graphs, Li et al.



Summary: Autoregressive Models

Positives

- Flexible output length
- Exact likelihood evaluation


Negatives

- No feature space
- Sequential generation (usually slow)




Generative Models

How do we measure the similarity of $p_{\theta}(x)$ and $p_{\text{data}}(x)$?




$p_{\text{data}}(x)$

1) Likelihood of data samples in $p_{\theta}(x)$



$$\approx KL(p_{\text{data}} \parallel p_{\theta})$$


2) Adversarial game



$$\approx JS(p_{\text{data}} \parallel p_{\theta})$$



Generative Adversarial Networks (GANs)




Generative Adversarial Networks (GANs)

Player 1: generator
Scores if discriminator
can't distinguish output
from real image

Player 2: discriminator
Scores if it can distinguish
between real and fake





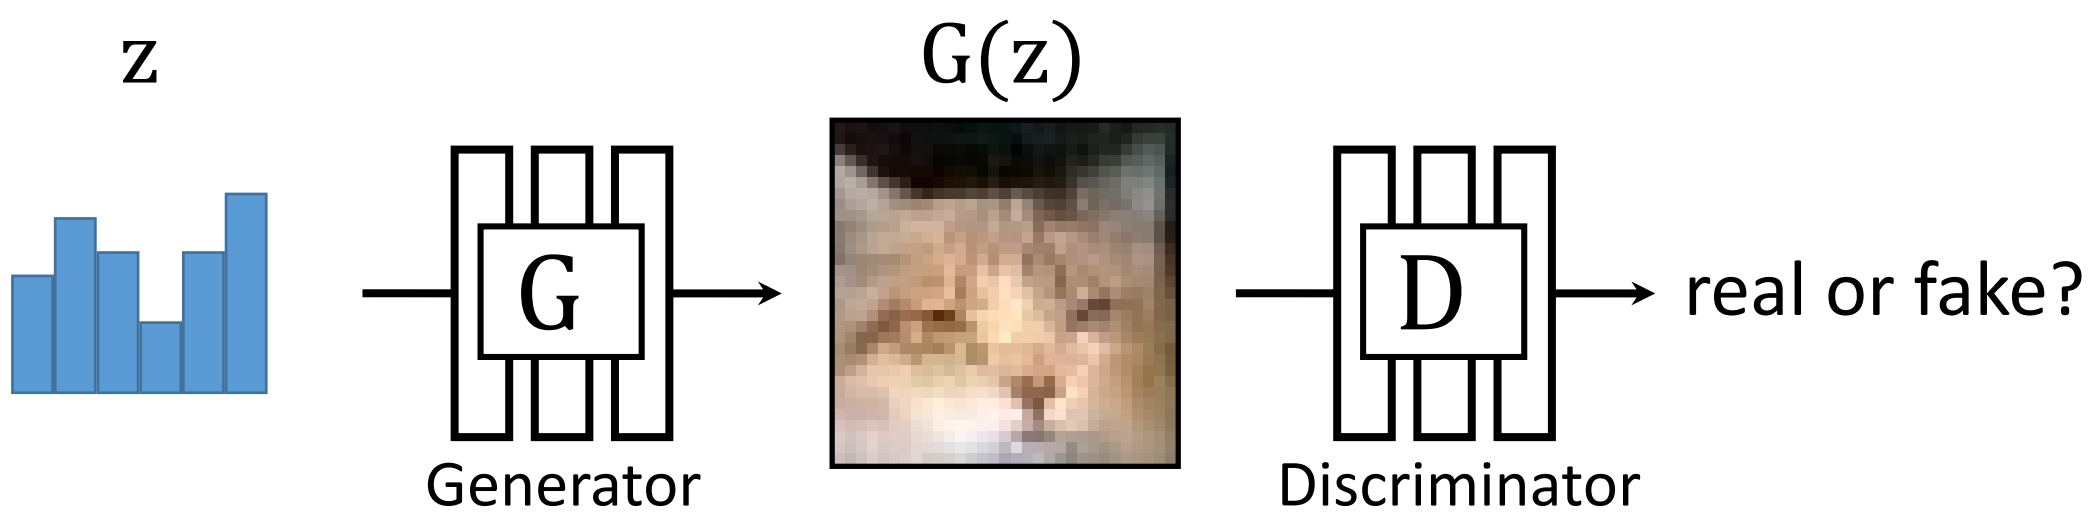
G : generate fake samples that can fool D



cat credit: aleju/cat-generator



slide credit: Phillip Isola & Jun-Yan Zhu



G : generate fake samples that can fool D

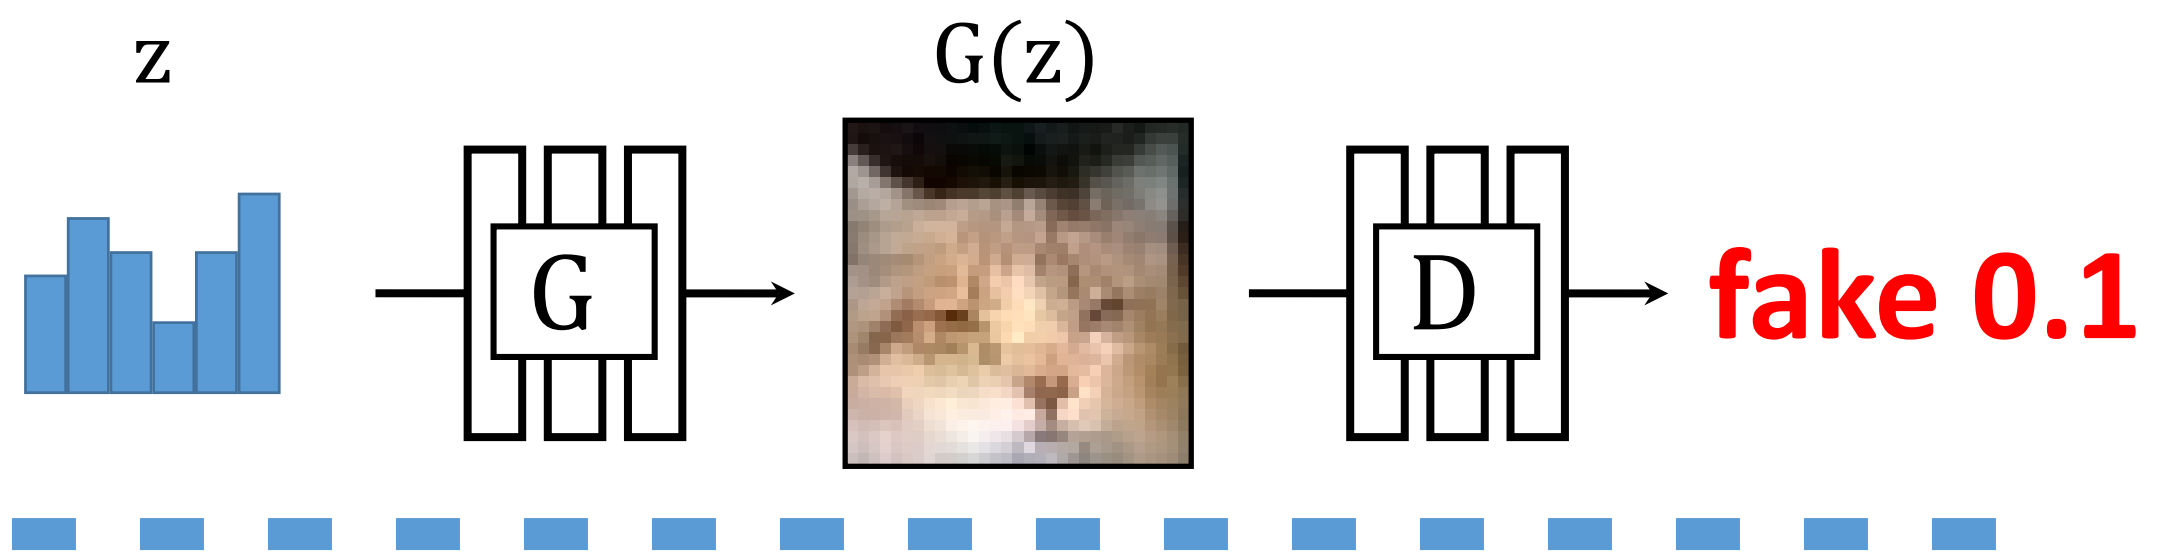
D : classify fake samples vs. real images





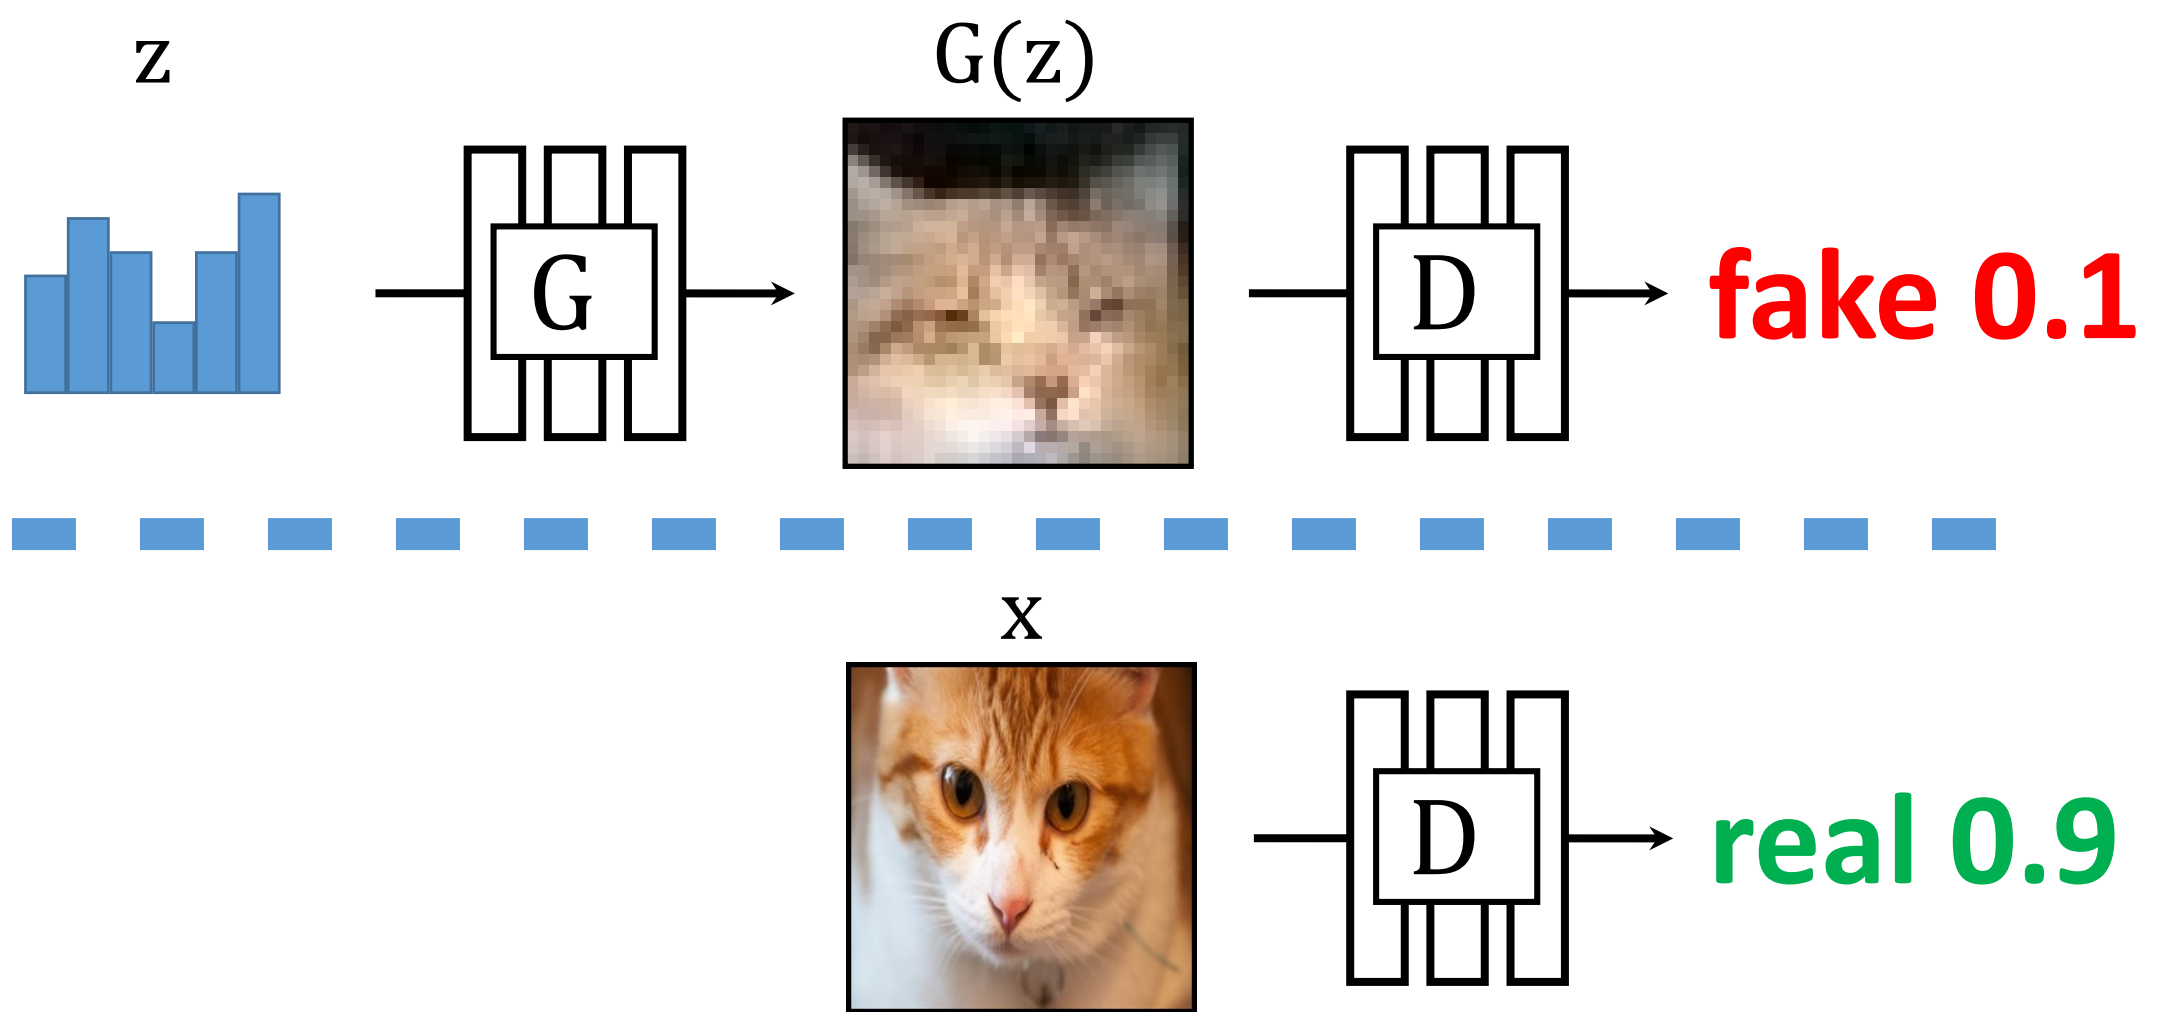
$$\min_G \max_D \mathbb{E}_{z,x} [\quad]$$





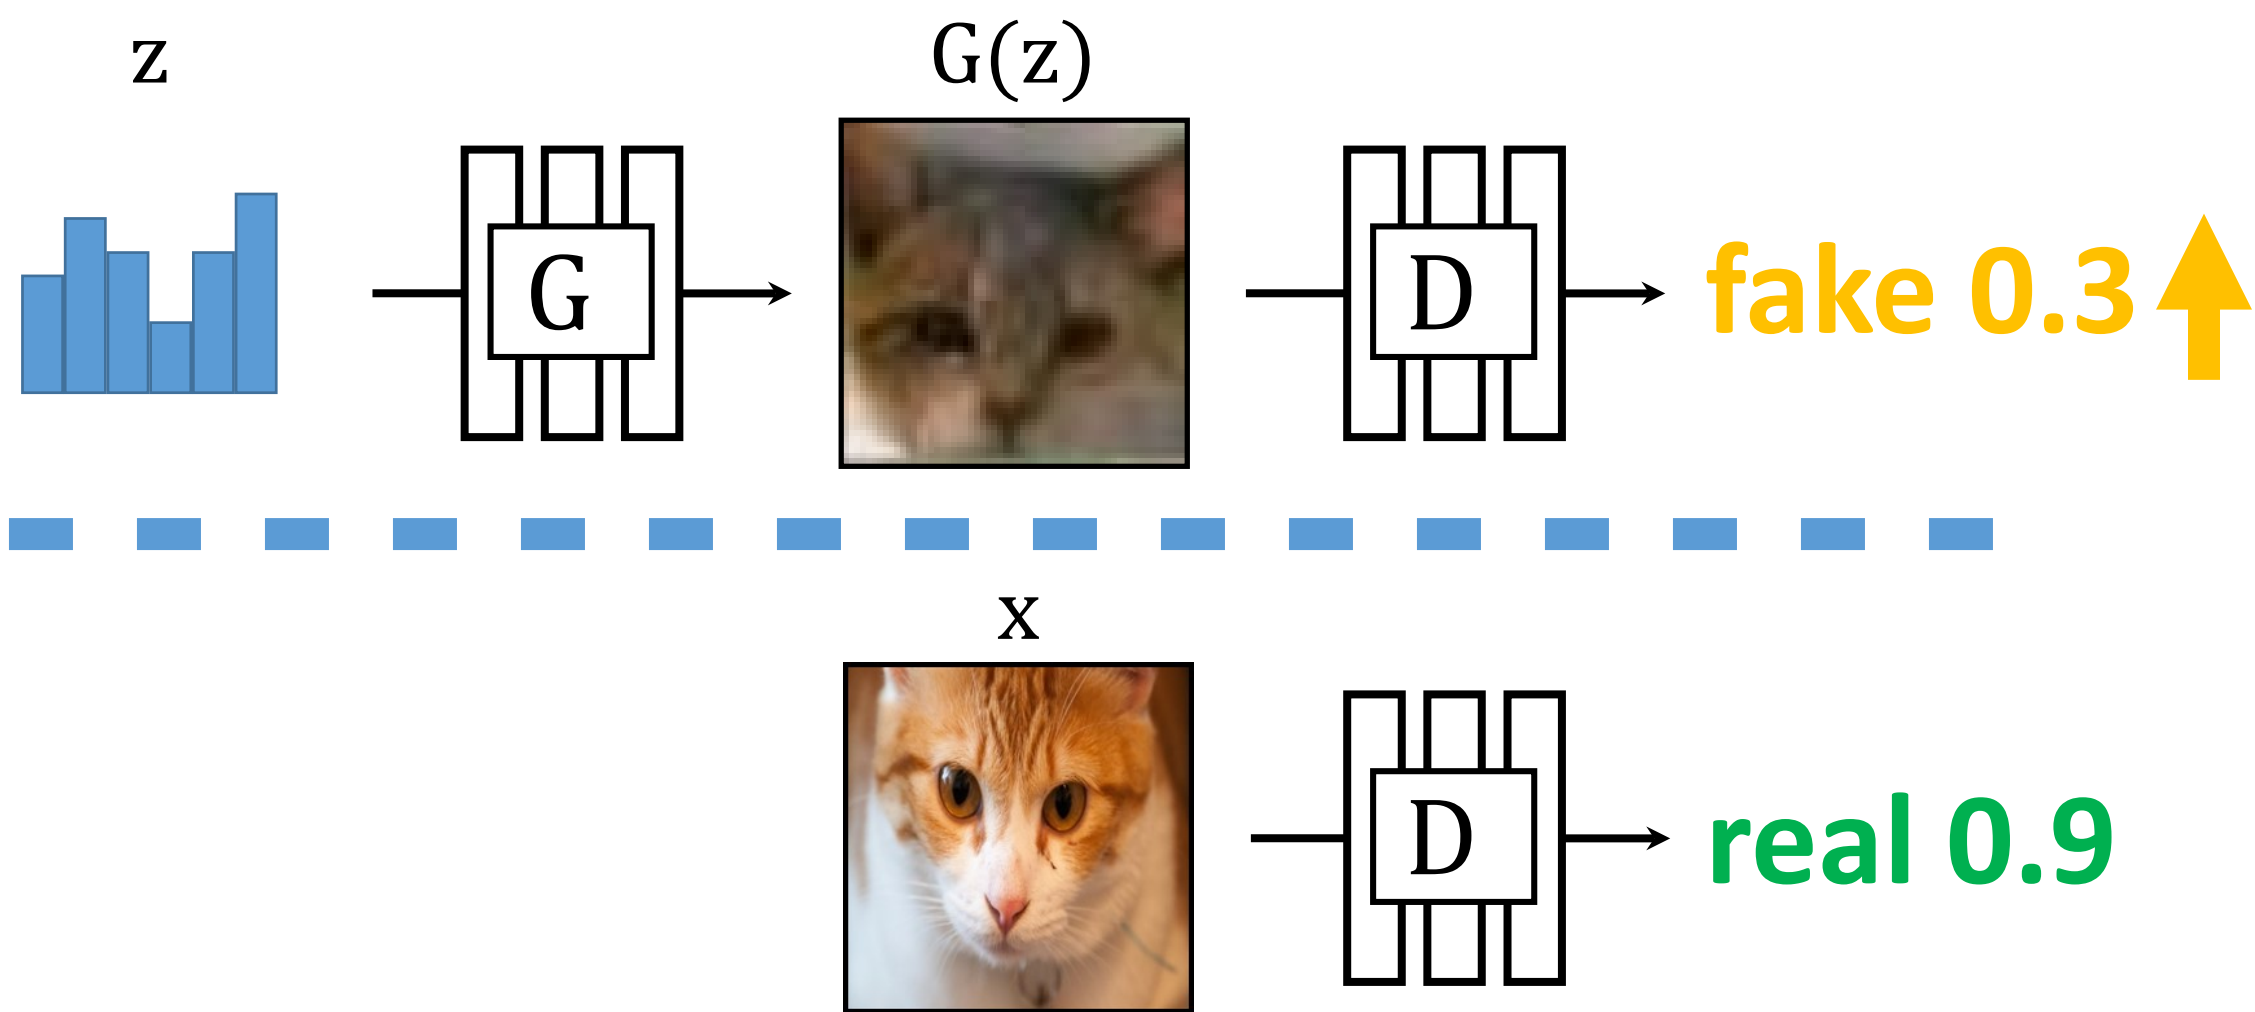
$$\min_G \max_D \mathbb{E}_{z,x} [\log \underbrace{D(G(z))}_{\text{fake}}]$$





$$\min_G \max_D \mathbb{E}_{z,x} [\log \underbrace{D(G(z))}_{\text{fake}} + \log(1 - \underbrace{D(x)}_{\text{real}})]$$





$$\min_G \max_D \mathbb{E}_{z,x} [\log \underbrace{D(G(z))}_{\text{Update } G} + \log(1 - D(x))]$$

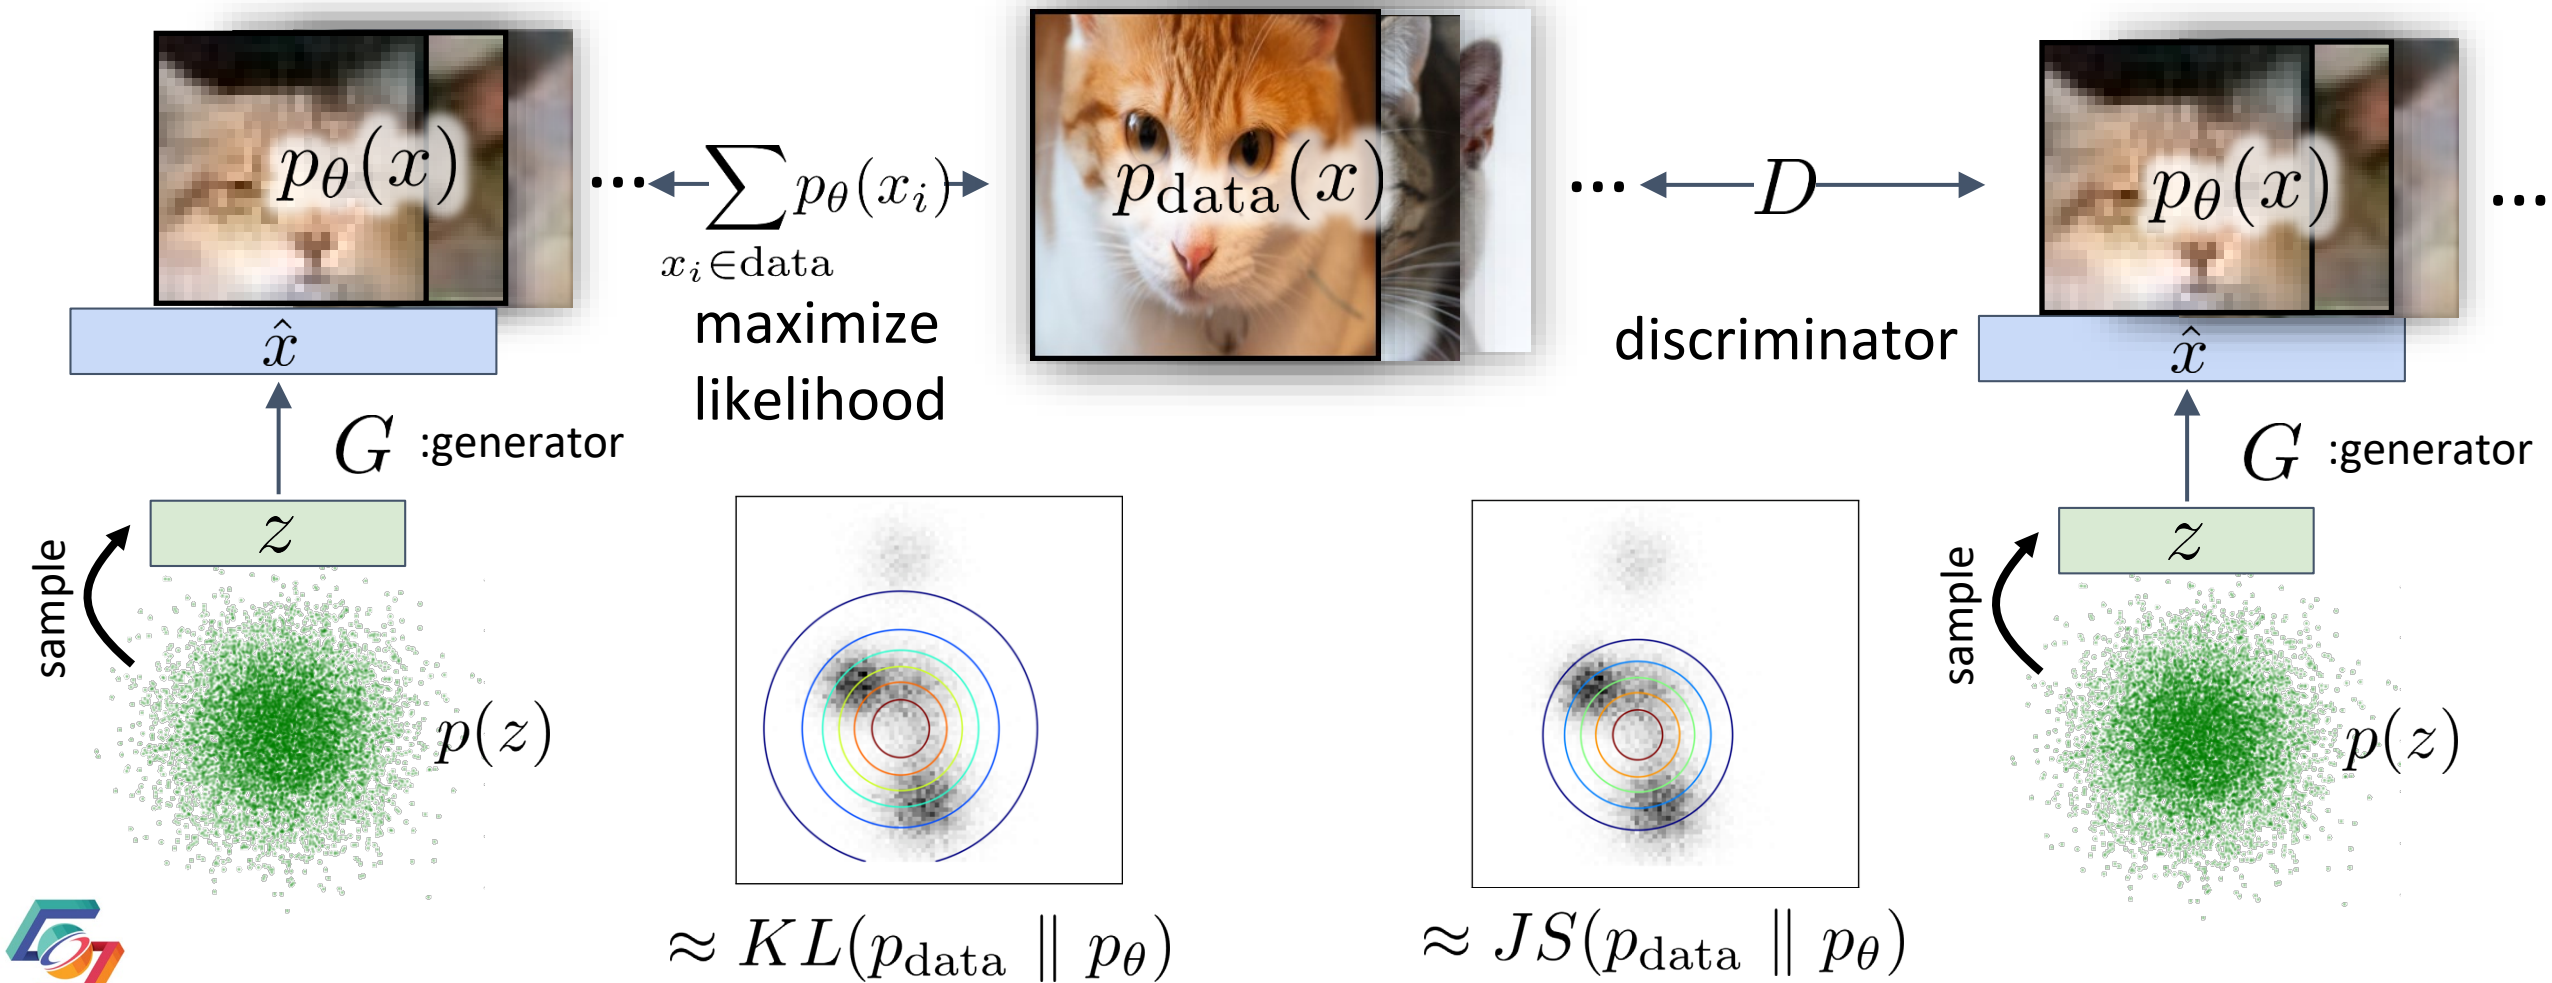


Generated Distributions of GANs vs ML

Sample? 
Evaluate? 

VAEs or
Norm. Flows


GANs



StyleGAN

Additional Tricks:

- Coarse-to-fine training
- Transformation of $p(z)$ to a more complex distr.
- ...



Summary: GANs

Positives

- Creates a feature space
- Currently highest-quality results

Negatives

- Can be unstable to train
- Not guaranteed to cover all of the data distribution
- Cannot evaluate likelihood



Open Problems

- More control
- Irregular data
- GAN training convergence
- Evaluating GANs




Conditional GAN: Pix2Pix

Labels to Street Scene



Labels to Facade




BW to Color



Aerial to Map



Day to Night



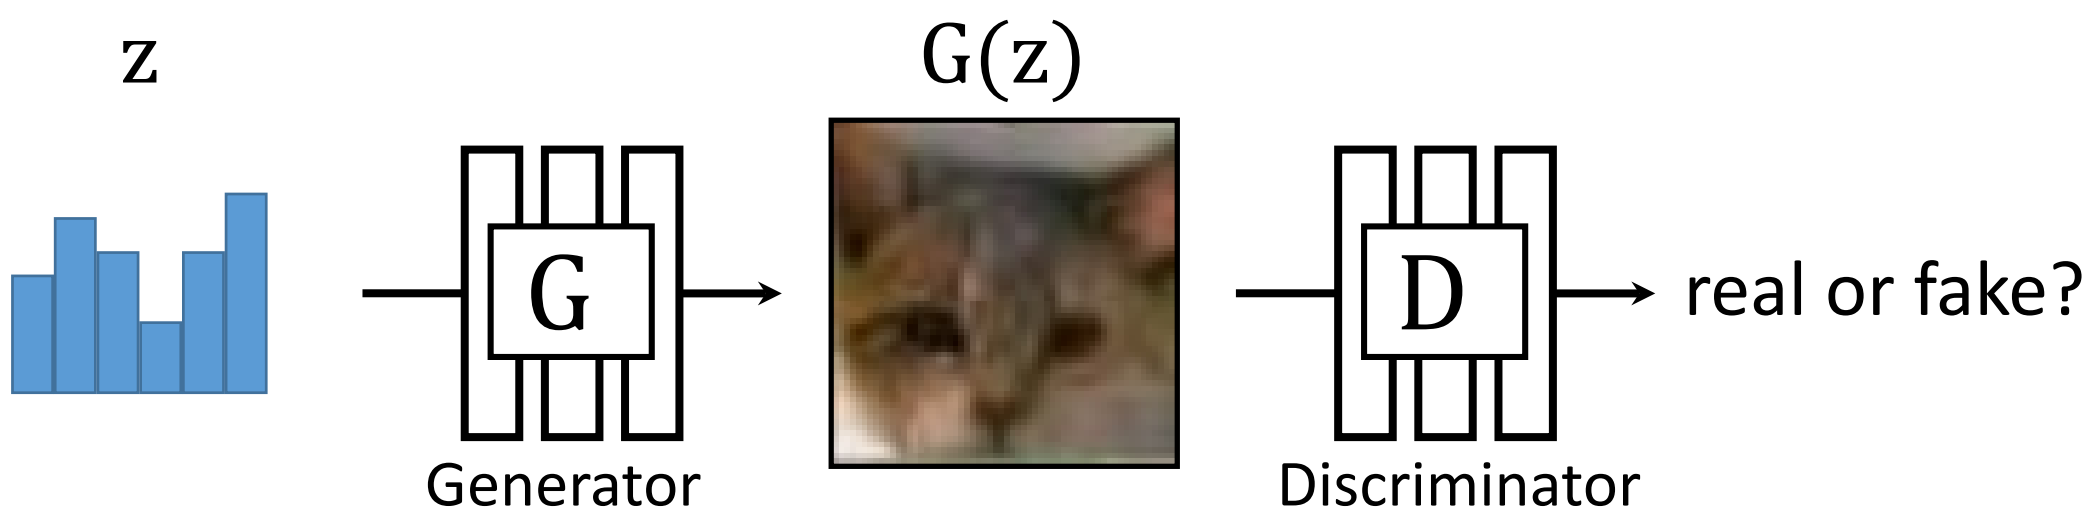
Edges to Photo



Image-to-image Translation with Conditional Adversarial Nets
Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, Alexei A. Efros. CVPR 2017

slide credit: Phillip Isola & Jun-Yan Zhu






$$\min_G \max_D \mathbb{E}_{z,x} [\log D(G(z)) + \log(1 - D(x))]$$






$$\min_G \max_D \mathbb{E}_{x,y} [\log D(G(x)) + \log(1 - D(y))]$$






$$\min_G \max_D \mathbb{E}_{x,y} [\log D(G(x)) + \log(1 - D(y))]$$





$$\min_G \max_D \mathbb{E}_{x,y} [\log D(G(x)) + \log(1 - D(y))]$$





$$\min_G \max_D \mathbb{E}_{x,y} [\log \underbrace{D(x, G(x))}_{\text{fake pair}} + \log(1 - \underbrace{D(x, y)}_{\text{real pair}})]$$


match joint distribution $p(G(x), y) \sim p(x, y)$



Edges → Images


Input

Output






Input

Output



Input

Output




Edges from [Xie & Tu, 2015]

slide credit: Phillip Isola & Jun-Yan Zhu



Sketches → Images


Input



Output




Input





Output



Input



Output



Trained on Edges → Images


Data from [Eitz, Hays, Alexa, 2012]

slide credit: Phillip Isola & Jun-Yan Zhu



Decomposition into Steps: FrankenGAN

input



Decomposition into Steps: FrankenGAN

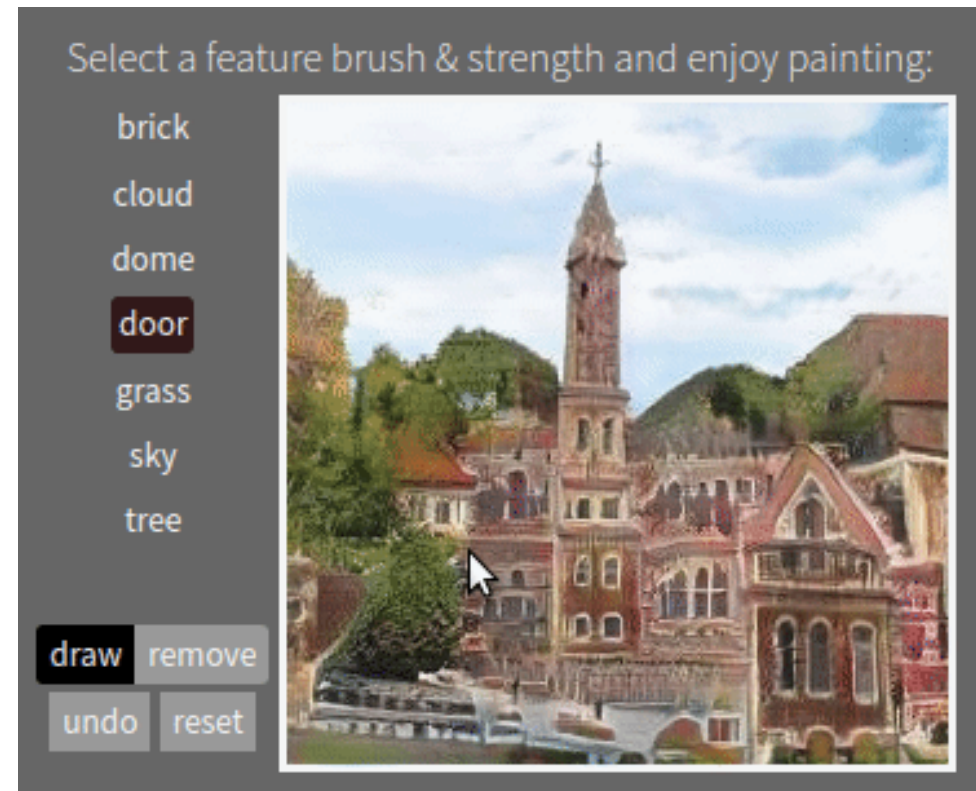
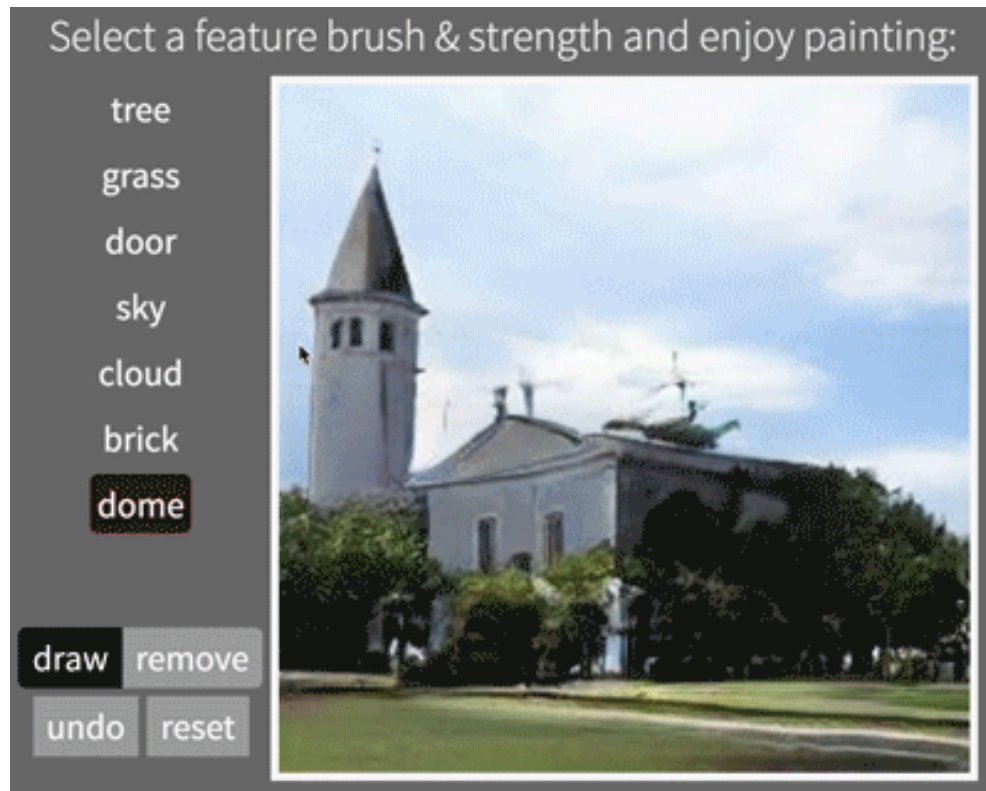




GAN Dissection

Question: How does a GAN create an image? What do individual neurons do?

Insight: Neurons are specialized to create objects of specific types




Open Problems

- More control
- Irregular data
- GAN training convergence
- Evaluating GANs




GRASS

Space of chairs?



GRASS


Part bounding boxes and their relationships represent a shape




Adjacency



Translational
symmetry



Rotational
symmetry




Reflectional
symmetry



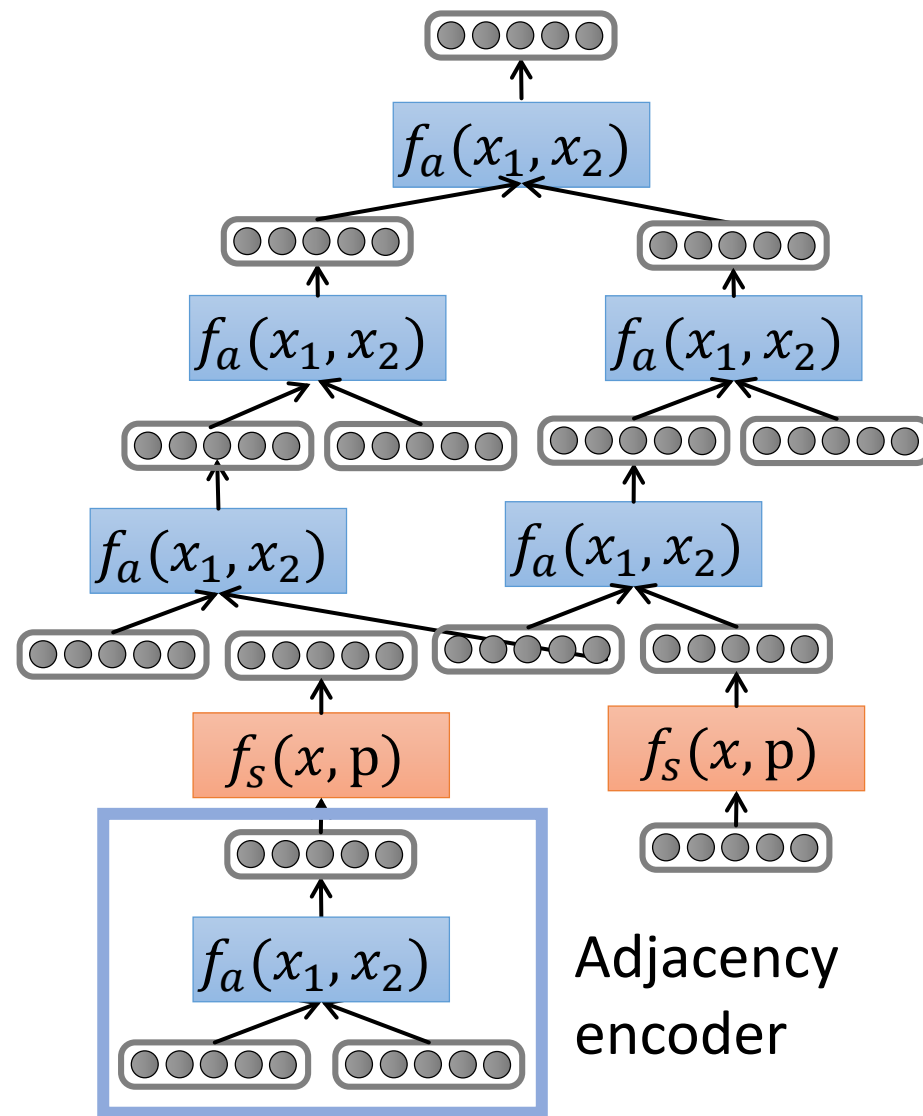

GRASS Training

VAE using a hierarchical encoder and decoder



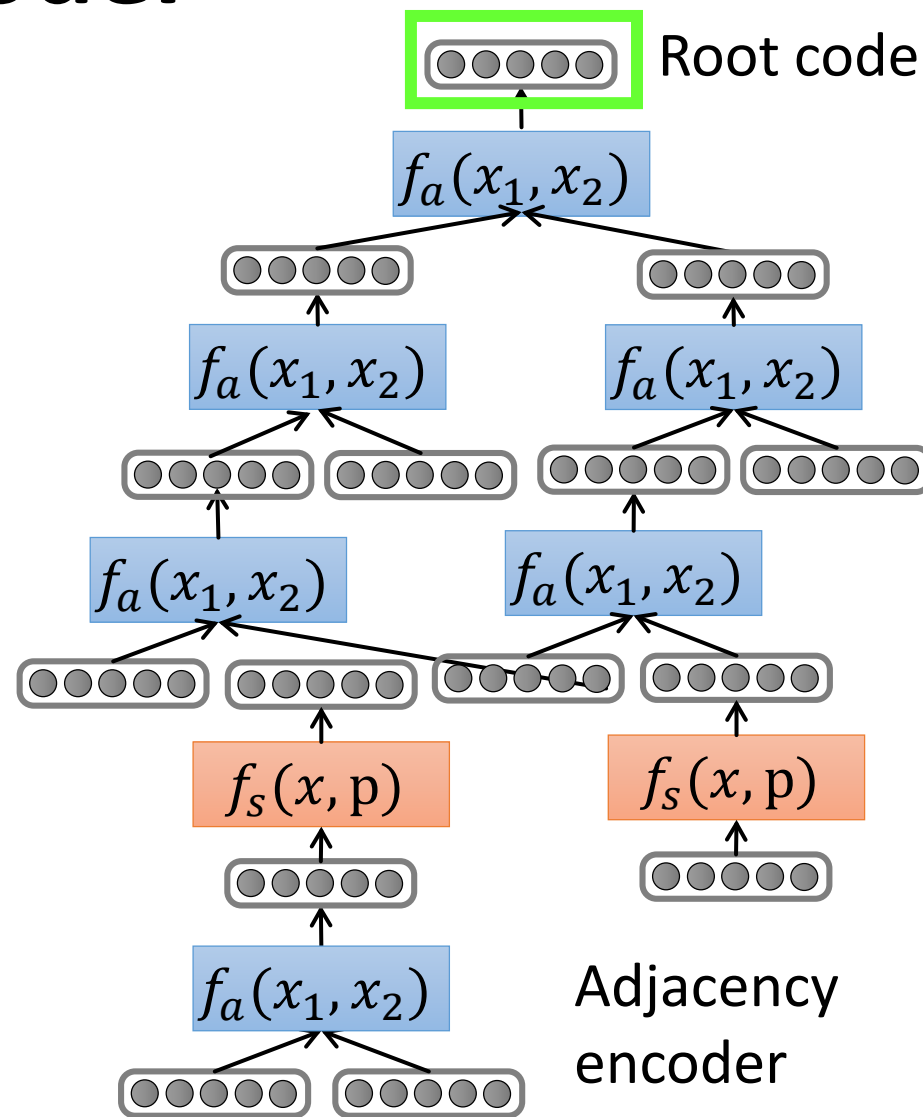

GRASS Hierarchical Encoder

Bottom-up merging




GRASS Hierarchical Encoder

Bottom-up merging



GRASS Results

interpolation




free generation




StructureNet

- Consistent structure across the dataset

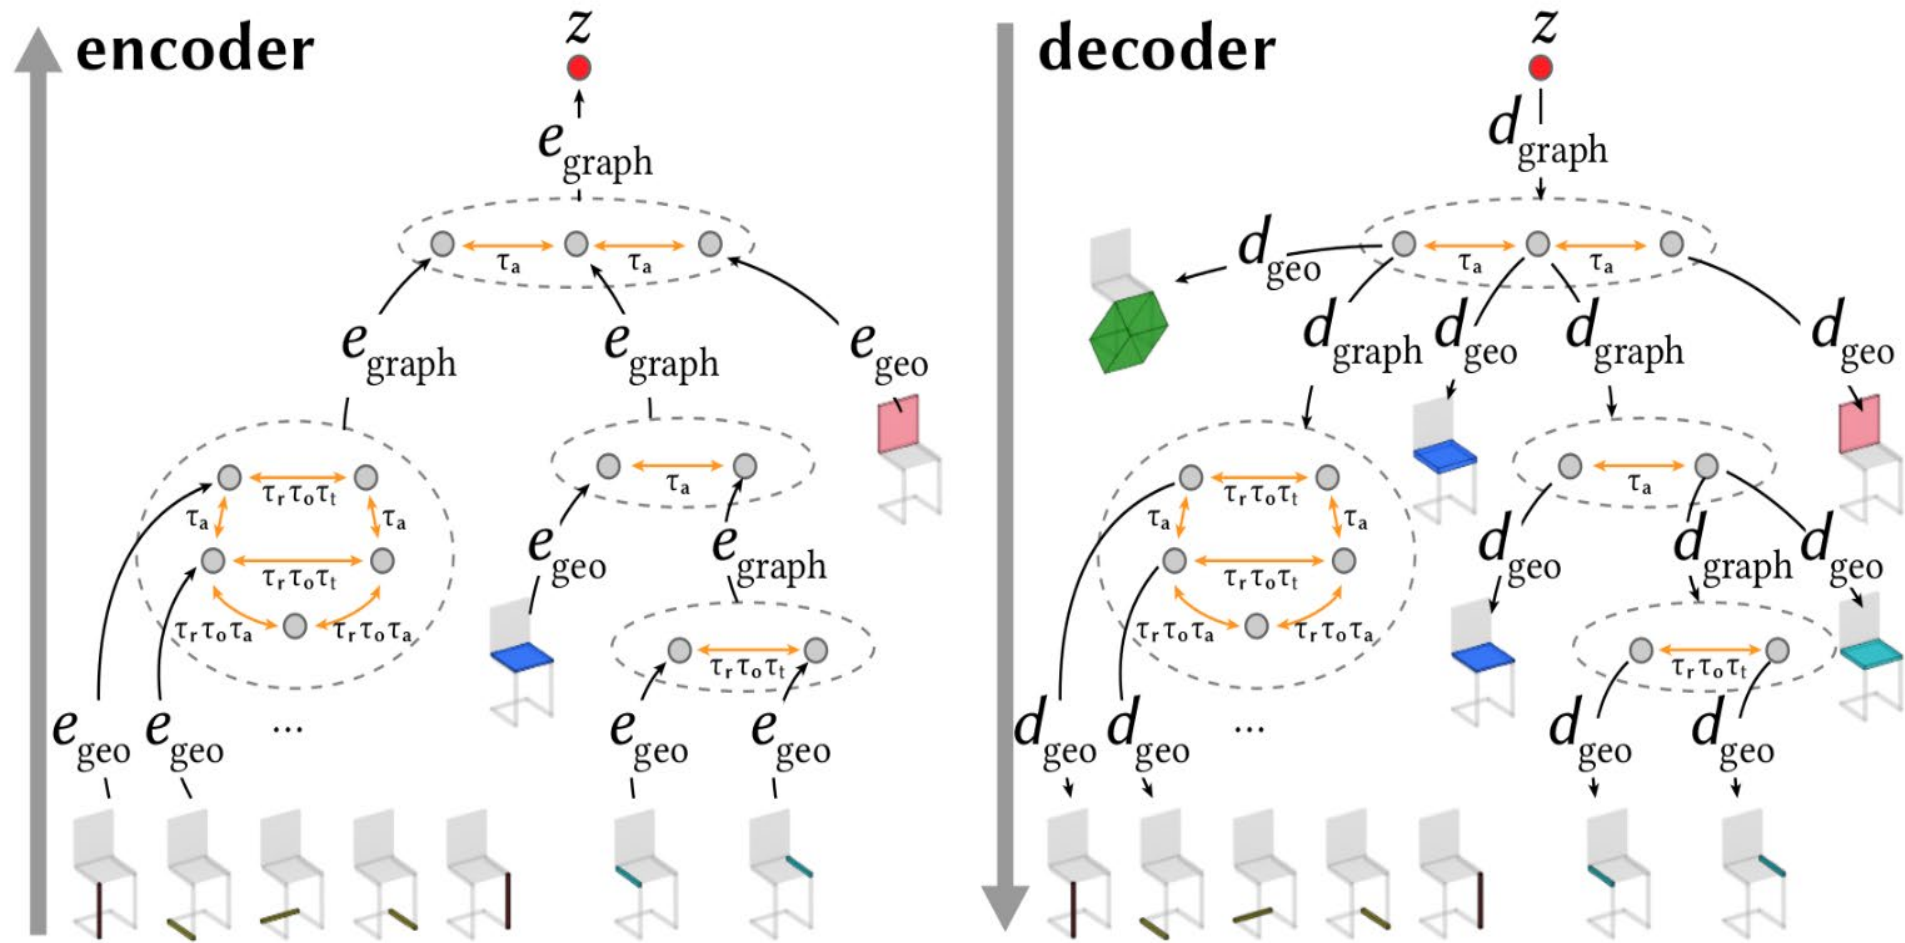


StructureNet

- Consistent structure across the dataset
- Richer structure




Hierarchical Graph Networks as Encoder and Decoder




StructureNet Results

interpolation



free generation



Open Problems

- More control
- Irregular data
- GAN training convergence
- Evaluating GANs




GAN Training Convergence

GAN training can be unstable


- Generator and discriminator do not always converge (Nash equilibrium)
- Vanishing discriminator gradients
- Mode Collapse



Vanishing Discriminator Gradients




Instance noise: adding noise to generated and real images



Roth et al. suggest an analytic convolution with a gaussian:



Stabilizing Training of Generative Adversarial Networks through Regularization, Roth et al. 2017

Wasserstein GANs: EMD as distance between p_{θ} and p_{data}



Mode Collapse

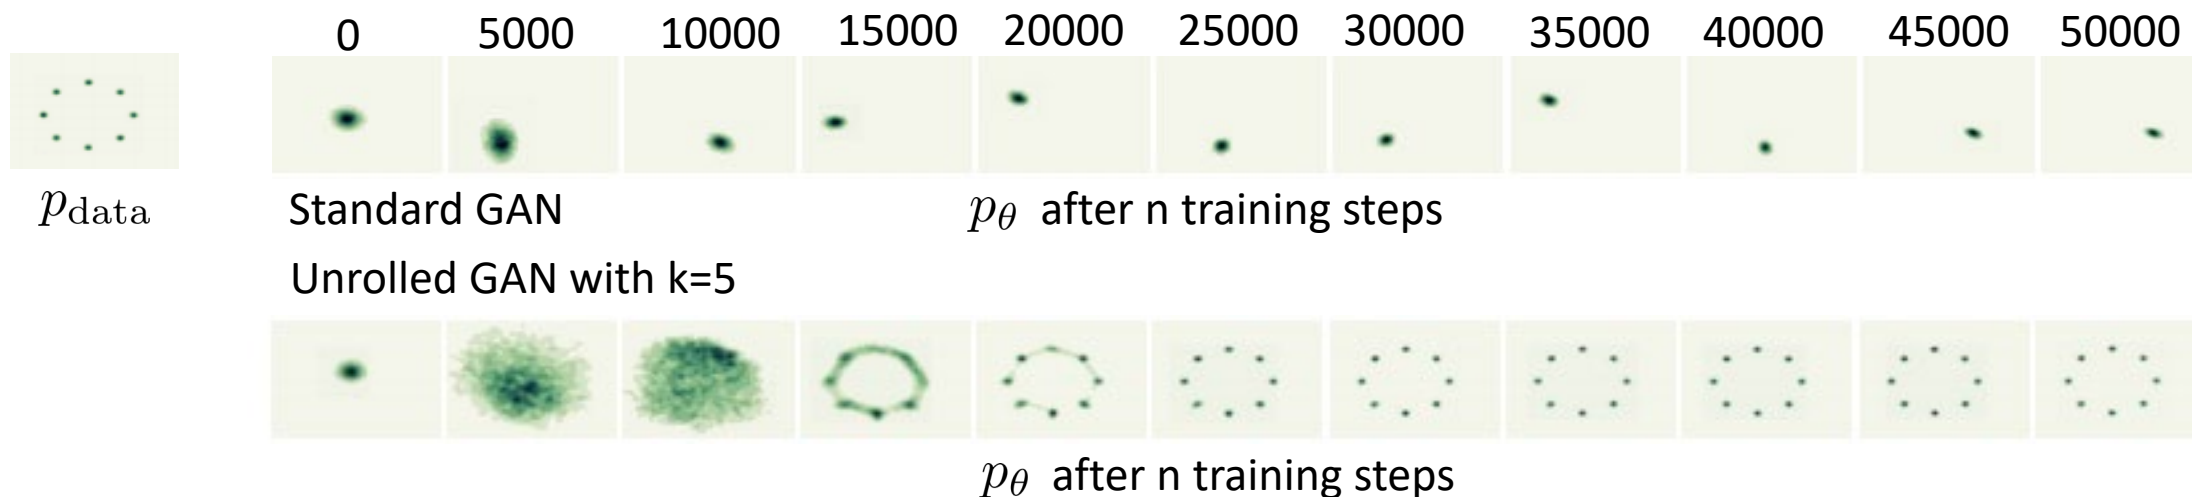
p_θ only covers one or a few modes of p_{data}



Mode Collapse

Solution attempts:

- Minibatch comparisons: Discriminator can compare instances in a minibatch (*Improved Techniques for Training GANs*, Salimans et al.)
- Unrolled GANs: Take k steps with the discriminator in each iteration, and backpropagate through all of them to update the generator




Summary

- Autoencoders
 - Can create a feature space, but bad generators
- VAEs
 - Lower quality generated samples (usually blurry)
 - Relatively stable to train
- Normalized Flows
 - Better quality generated samples
 - Invertible
 - Relatively stable, but expensive to train
- GANs
 - Can not find a latent representation for a given sample (no encoder)
 - Usually better generators than VAEs
 - Currently unstable training (active research)



Course Information (slides/code/comments)



<http://geometry.cs.ucl.ac.uk/creativeai/>

