

# Deep Learning for Computer Vision

Fall 2022

<https://cool.ntu.edu.tw/courses/189345> (NTU COOL)

<http://vllab.ee.ntu.edu.tw/dlcv.html> (Public website)

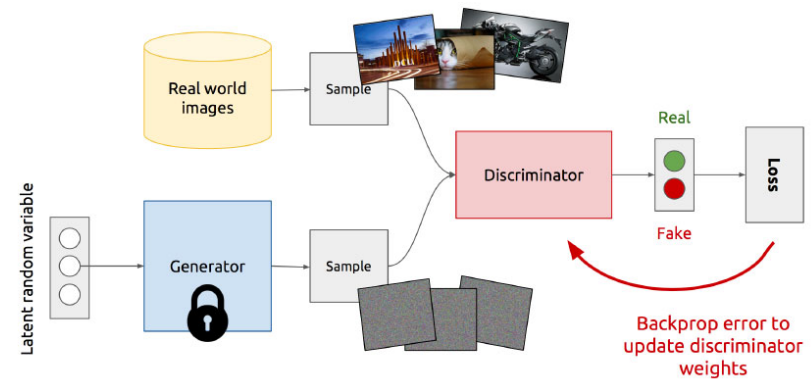


Yu-Chiang Frank Wang 王鈺強, Professor

Dept. Electrical Engineering, National Taiwan University

2022/10/4


# What's to Be Covered Today...

- Generative Models
  - Auto-Encoder vs. Variational Auto-Encoder
  - Generative Adversarial Network (GAN)
  - Diffusion Model
- Unfortunately, lots of equations today...  
I will try to make today's lecture as painless as possible!





# Discriminative vs. Generative Models

**Discriminative Model:**  
Learn a probability distribution  $p(y|x)$



**Generative Model:**  
Learn a probability distribution  $p(x)$




**Conditional Generative Model:** Learn  $p(x|y)$

Discriminative model: the possible labels for each input "compete" for probability mass. But no competition between **images**

# Discriminative vs. Generative Models (cont'd)

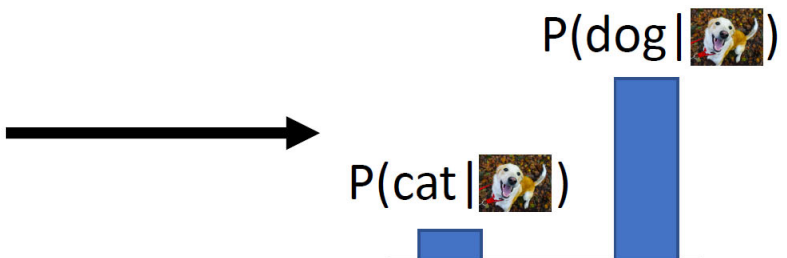

## Discriminative Model:

Learn a probability distribution  $p(y|x)$



## Generative Model:

Learn a probability distribution  $p(x)$



## Conditional Generative Model: Learn $p(x|y)$

Discriminative model: No way for the model to handle unreasonable inputs; it must give label distributions for all images



# Discriminative vs. Generative Models (cont'd)

## Discriminative Model:

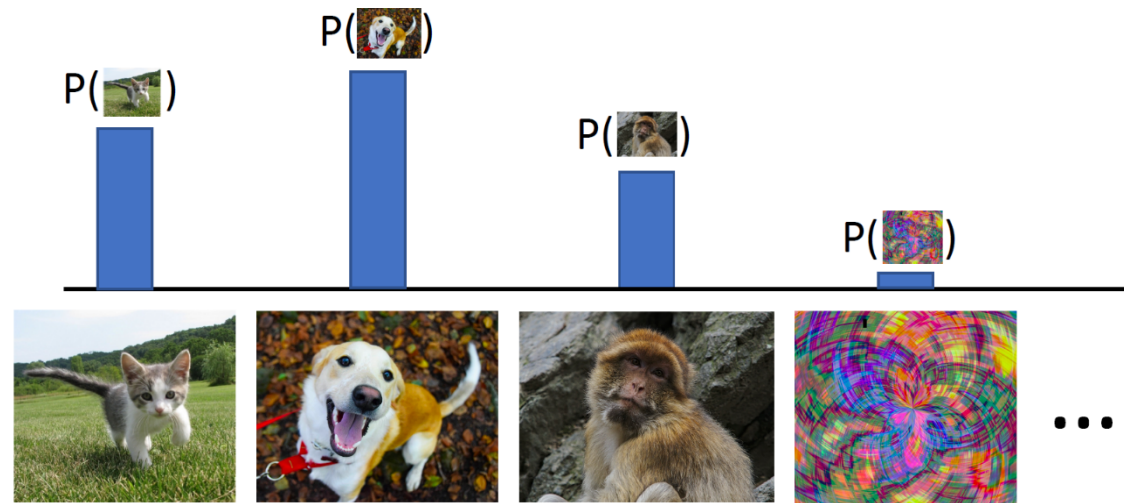
Learn a probability distribution  $p(y|x)$

## Generative Model:

Learn a probability distribution  $p(x)$

## Conditional Generative Model:

Learn  $p(x|y)$



Generative model: All possible images compete with each other for probability mass

Model can “reject” unreasonable inputs by assigning them small values

# Discriminative vs. Generative Models (cont'd)

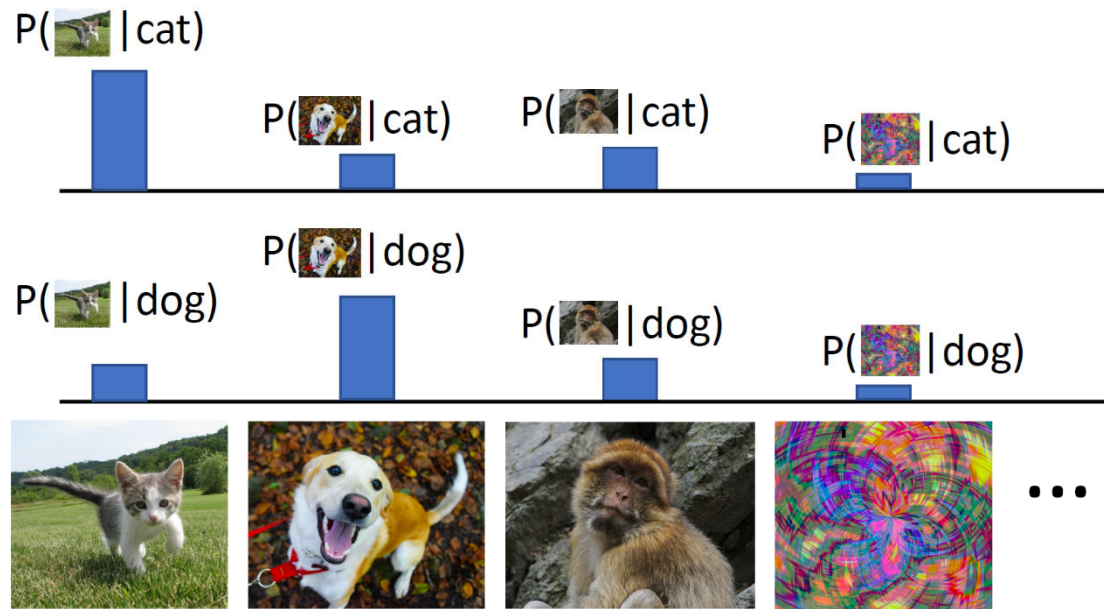
## Discriminative Model:

Learn a probability distribution  $p(y|x)$

## Generative Model:

Learn a probability distribution  $p(x)$

**Conditional Generative Model:** Learn  $p(x|y)$



Conditional Generative Model: Each possible label induces a competition among all images

# Discriminative vs. Generative Models (cont'd)

## Discriminative Model:

Learn a probability distribution  $p(y|x)$

## Generative Model:

Learn a probability distribution  $p(x)$

**Conditional Generative Model:** Learn  $p(x|y)$

Recall **Bayes' Rule**:

$$\underbrace{P(x | y)}_{\text{Conditional Generative Model}} = \frac{\underbrace{P(y | x)}_{\text{Discriminative Model}}}{\underbrace{P(y)}_{\text{Prior over labels}}} \underbrace{P(x)}_{\text{(Unconditional) Generative Model}}$$

We can build a conditional generative model from other components!

# Additional Remarks

- Discriminative Models
  - Learn a (posterior) probability distribution  $p(y|x)$
  - Assign labels to each instance  $x$
  - Supervised learning
- Generative Models
  - Learn a probability distribution  $p(x)$
  - Data representation, detect outliers, etc.
  - Unsupervised learning

# What Have Been Done Using Deep Generative Models?

- Progress on synthesizing images (ImageNet)




(Odena 2018)

Odena et al  
2016

Miyato et al  
2017

Zhang et al  
2018

Brock et al  
2018

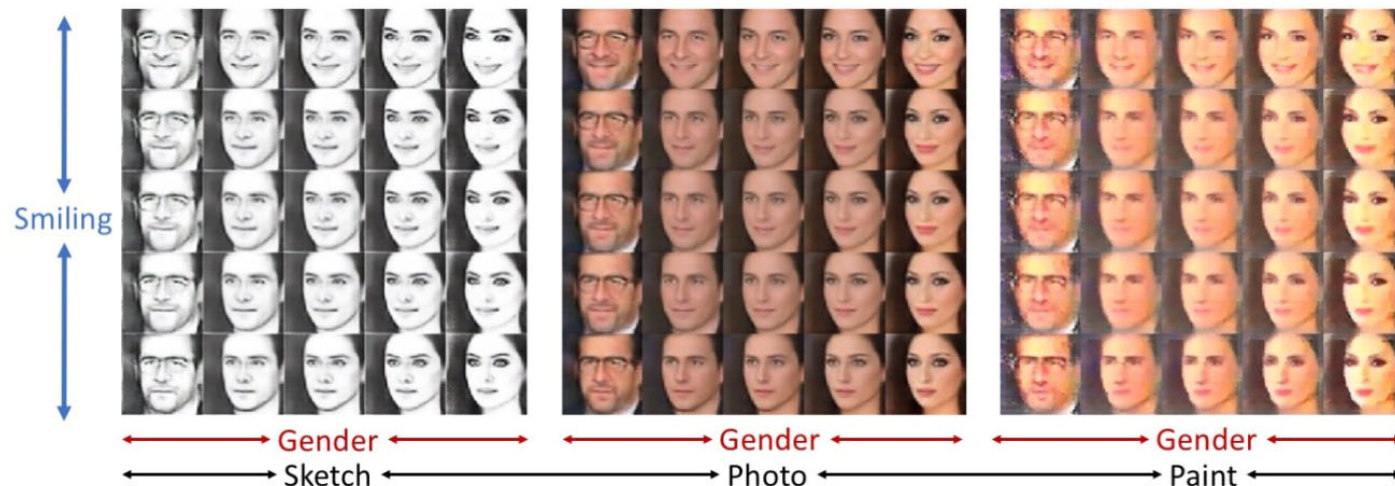


Super-Resolution via Repeated Refinements (SR3) by  
Class Diffusion Models (Google, 2021)



# Why We Need Generative Models?

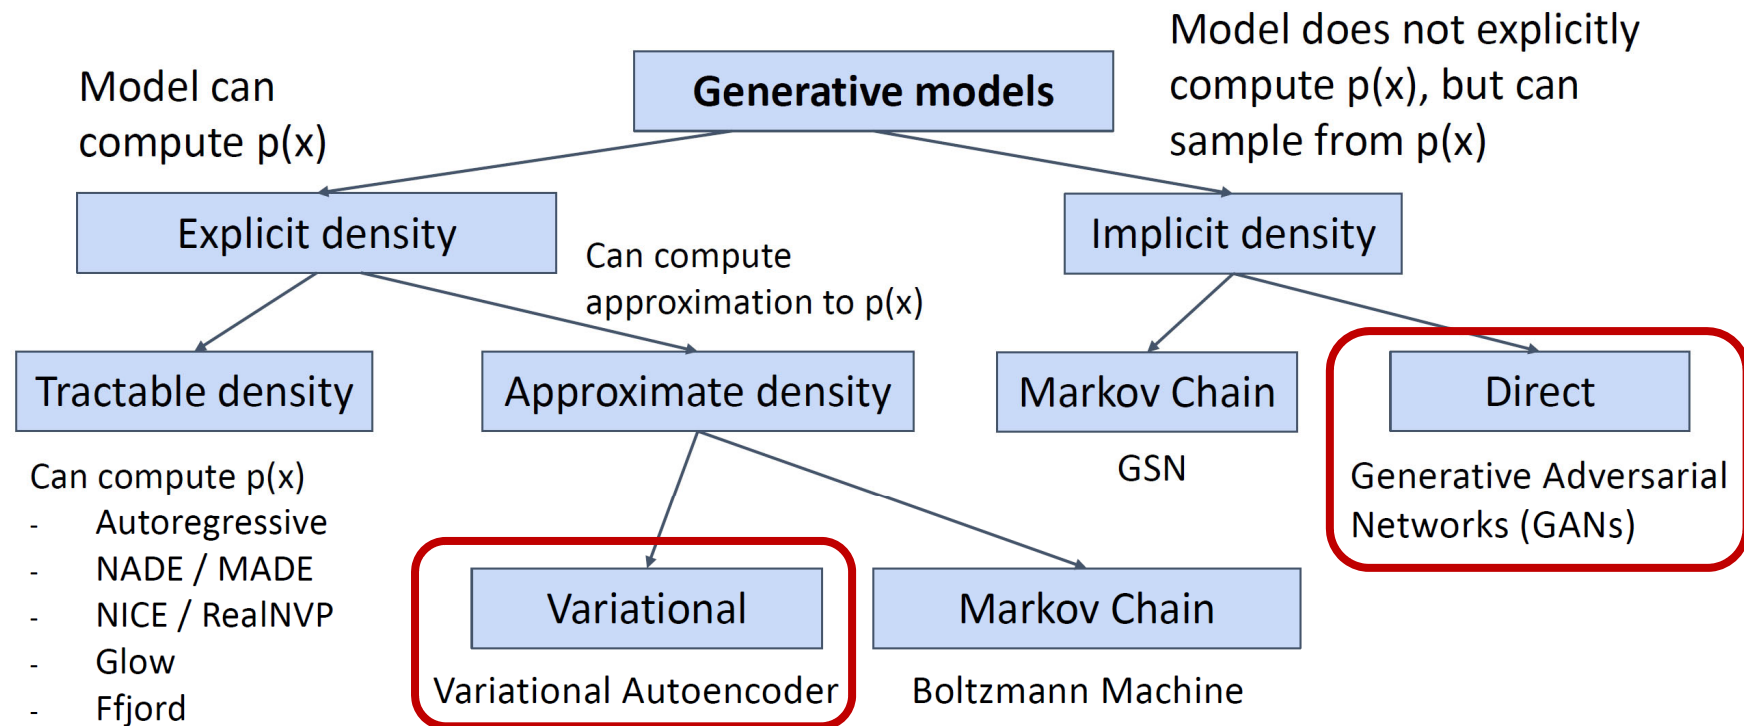
- Remarks
  - Able to process data information (e.g., priors like attribute, category, etc.) for synthesis, prediction, or recognition purposes
    - For example, with latent feature  $z$  derived from  $x$ , one may have  $P(z)$  may describe image variants.
    - Or,  $z$  in  $P(z)$  may annotate object categorical or attribute information.



- We will talk about a variety of visual applications based on generative models later.




# Taxonomy of Generative Models



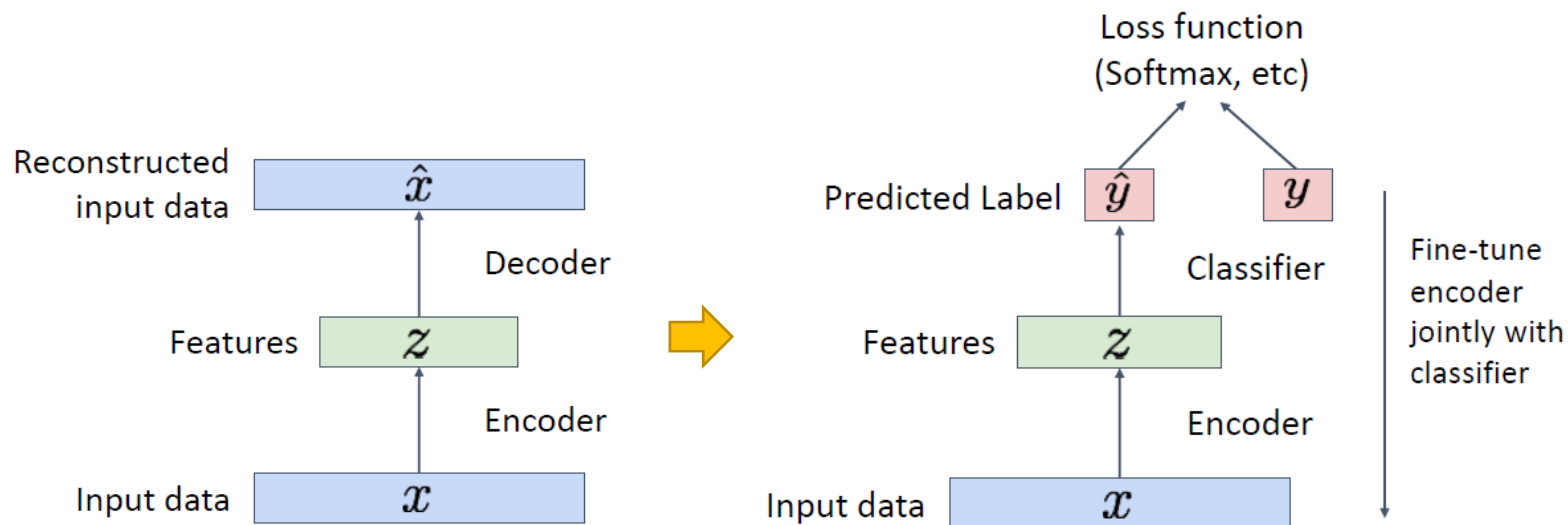
# Take a Deep Look to Discover Latent Variables/Representations

- Autoencoder
  - Autoencoding = encoding itself with recovery purposes
  - In other words, encode/decode data with reconstruction guarantees
  - Latent variables/features as deep representations
  - Example objective/loss function at output:
    - L2 norm between input and output, i.e.,  $\|\hat{x} - x\|_2$



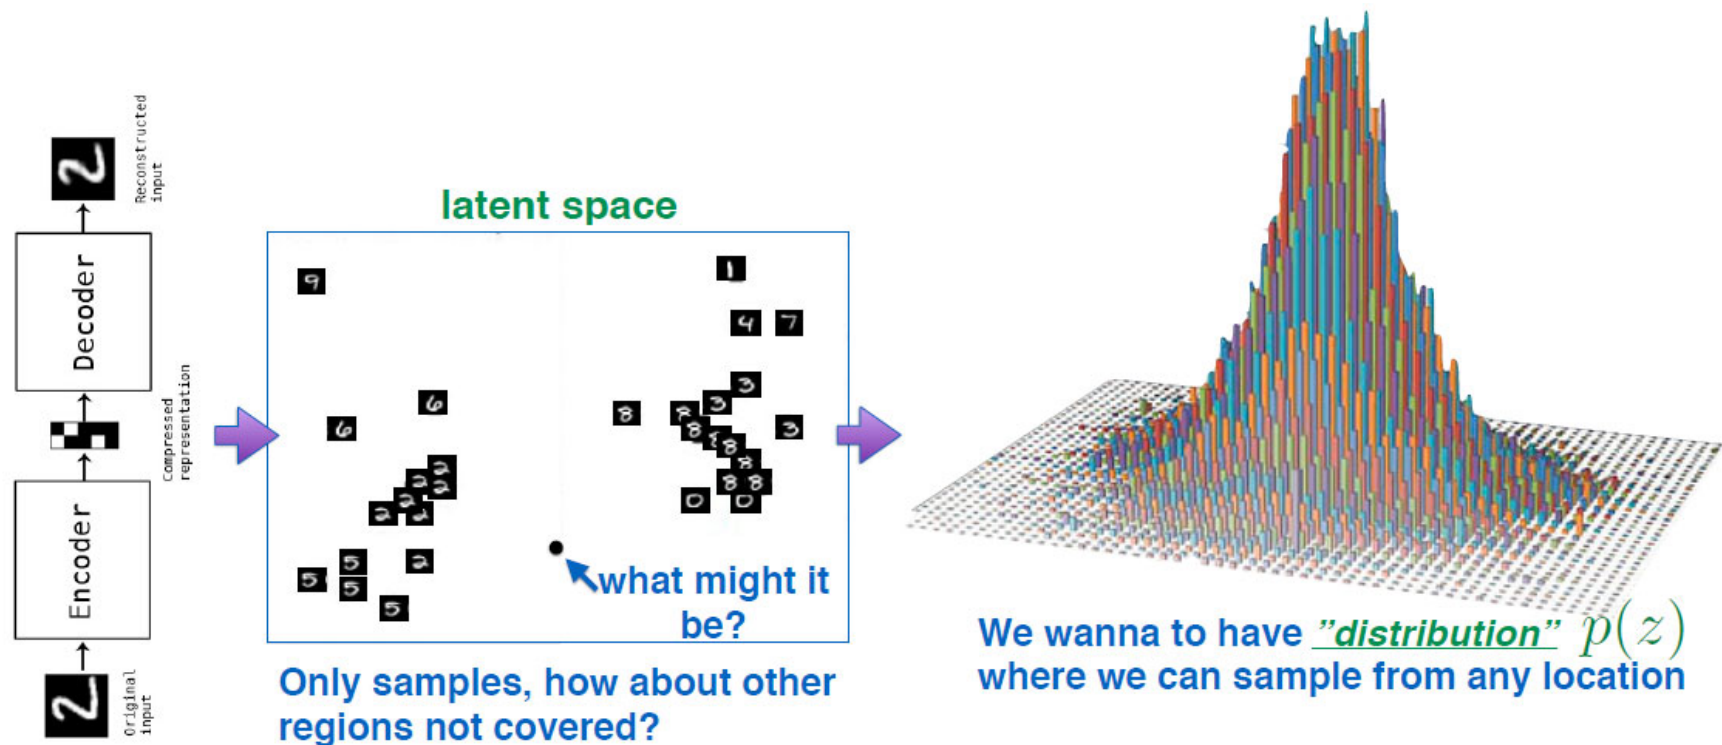
## Take a Deep Look to Discover Latent Variables/Representations (cont'd)

- Autoencoder (AE) for downstream tasks
  - Train AE with reconstruction guarantees
  - Keep encoder (and the derived features) for downstream tasks (e.g., classification)
  - Thus, a trained encoder can be applied to initialize a supervised model

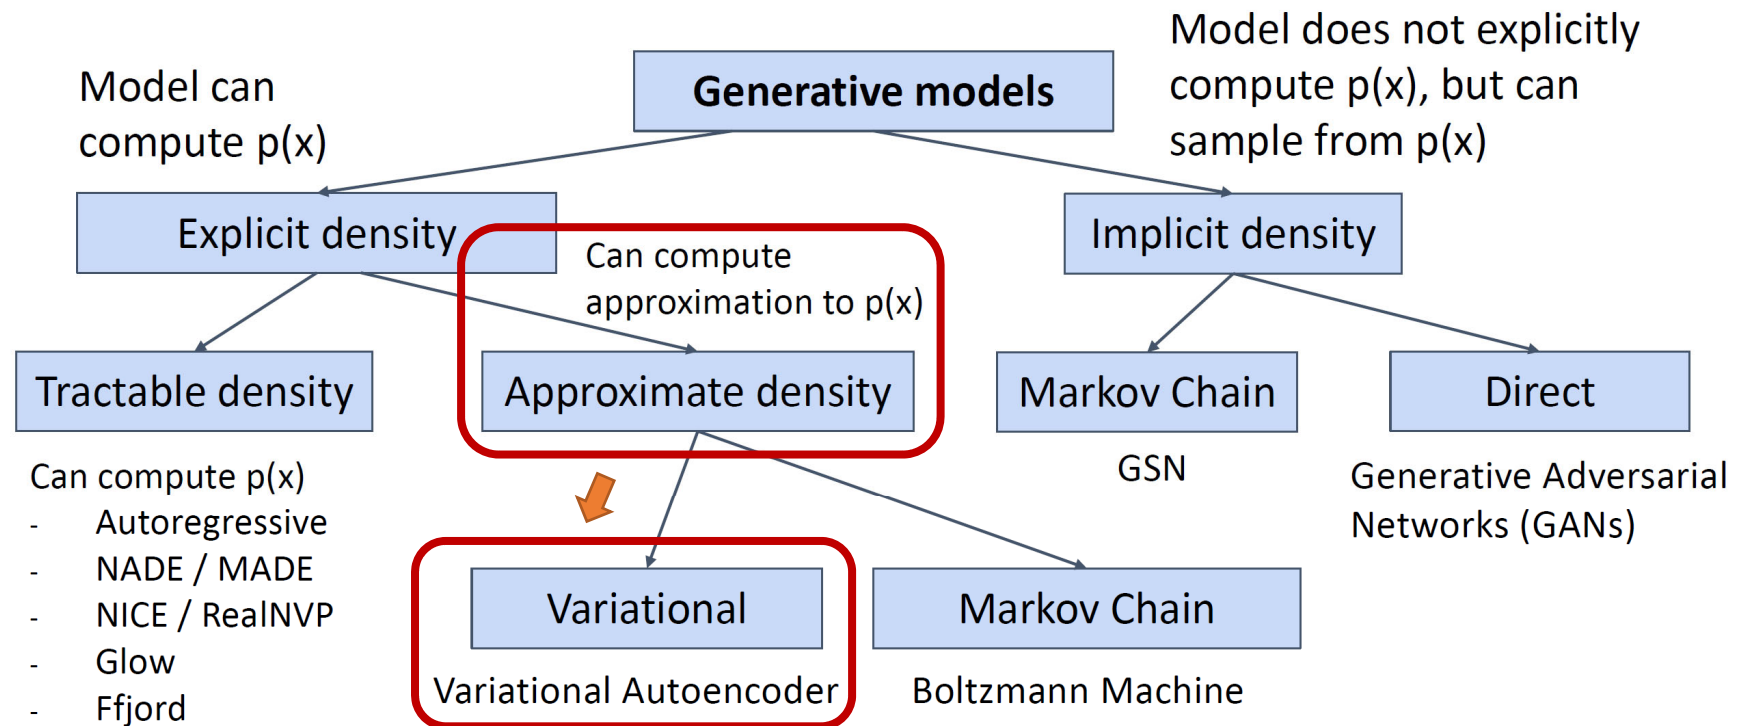


## Take a Deep Look to Discover Latent Variables/Representations (cont'd)

- What's the Limitation of Autoencoder?



# Taxonomy of Generative Models




# Variational Autoencoder

- Probabilistic Spin on AE

- Learn latent feature  $z$  from raw data  $x$
- Sample from the latent space (via model) to generate data

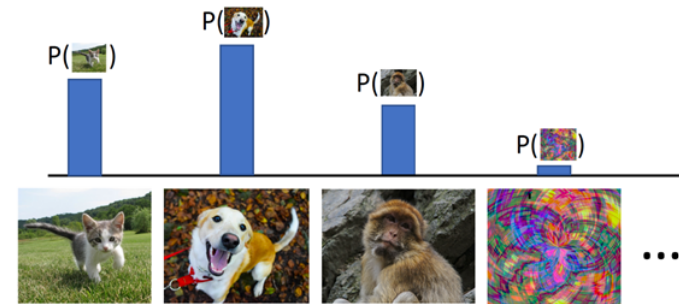
Sample from  
conditional  
 $p_{\theta^*}(x | z^{(i)})$

Sample  $z$   
from prior  
 $p_{\theta^*}(z)$





Assume simple prior  $p(z)$ , e.g. Gaussian

Represent  $p(x|z)$  with a neural network  
(Similar to **decoder** from autencoder)



- $p(x|z)$  is implemented via a (probabilistic) **decoder**




Decoder inputs  $z$ , outputs mean  $\mu_{x|z}$   
and (diagonal) covariance  $\Sigma_{x|z}$

↓  
Sample  $x$  from Gaussian with mean  
 $\mu_{x|z}$  and (diagonal) covariance  $\Sigma_{x|z}$



## Variational Autoencoder (cont'd)



- Remarks

- Train VAE via maximum likelihood of data  $p(x)$
- Note that we don't observe  $z$  & need to marginalize it:


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

- We can compute  $p_{\theta}(x|z)$  with the decoder module, and we assume Gaussian prior for  $p_{\theta}(z)$
- However, **can't integrate over all possible  $z$ !**
- What else can we do? Recall that we have Bayes' rule:


$$p_{\theta}(x) = \frac{p_{\theta}(x | z) p_{\theta}(z)}{p_{\theta}(z | x)}$$

We can't compute  $p_{\theta}(z | x)$ , but we can train the encoder module to learn

$$q_{\phi}(z | x) \approx p_{\theta}(z | x)$$



## Variational Autoencoder (cont'd)



- Again, we aim to maximize


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

we have...

**Decoder network** inputs  
latent code  $z$ , gives  
distribution over data  $x$


**Encoder network** inputs  
data  $x$ , gives distribution  
over latent codes  $z$

$$p_{\theta}(x | z) = N(\mu_{x|z}, \Sigma_{x|z}) \quad q_{\phi}(z | x) = N(\mu_{z|x}, \Sigma_{z|x})$$



- If we ensure  $q_{\phi}(z | x) \approx p_{\theta}(z | x)$   
then we have 
$$p_{\theta}(x) = \frac{p_{\theta}(x | z) p_{\theta}(z)}{p_{\theta}(z | x)} \approx \frac{p_{\theta}(x | z) p(z)}{q_{\phi}(z | x)}$$

# Training VAE



$$\log p_{\theta}(x) = \log \frac{p_{\theta}(x | z)p(z)}{p_{\theta}(z | x)} = \log \frac{p_{\theta}(x|z)p(z)q_{\phi}(z|x)}{p_{\theta}(z|x)q_{\phi}(z|x)}$$

$$= E_z[\log p_{\theta}(x|z)] - E_z \left[ \log \frac{q_{\phi}(z|x)}{p(z)} \right] + E_z \left[ \log \frac{q_{\phi}(z|x)}{p_{\theta}(z|x)} \right]$$

$$= E_{z \sim q_{\phi}(z|x)}[\log p_{\theta}(x|z)] - D_{KL}(q_{\phi}(z|x), p(z)) + D_{KL}(q_{\phi}(z|x), p_{\theta}(z|x))$$

**Data reconstruction**

**KL divergence**  
between sample  
distribution from the  
encoder and the prior

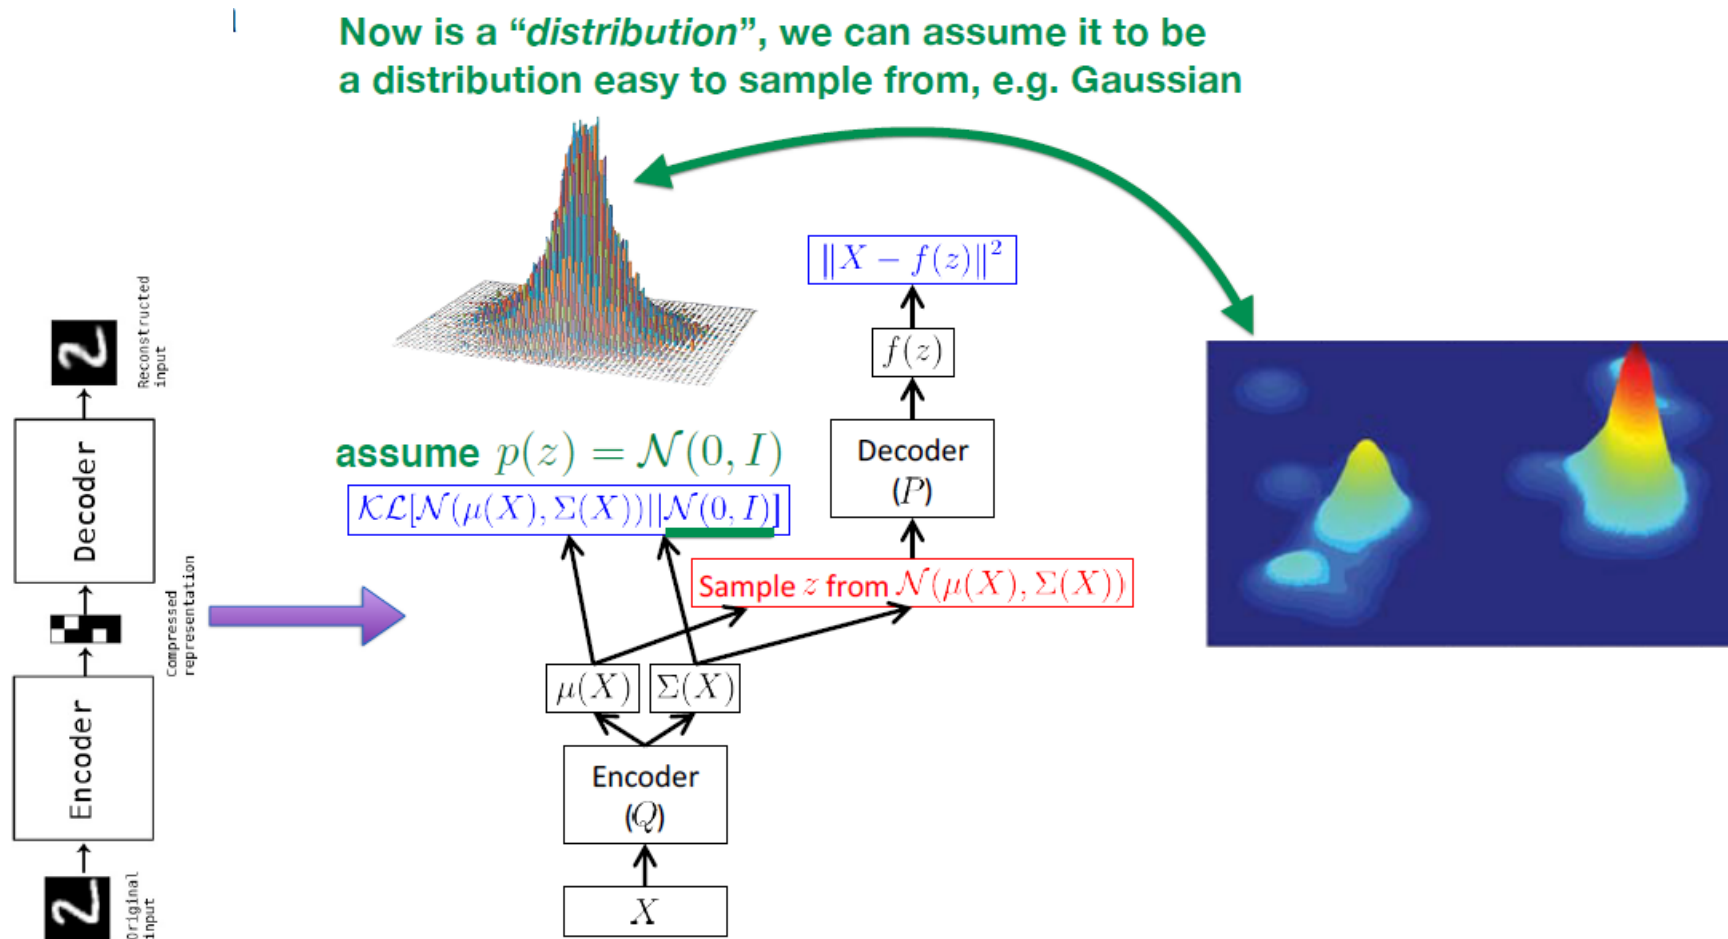
**KL divergence** between  
sample distribution  
from the encoder and  
the posterior of data

$$\Rightarrow \log p_{\theta}(x) \geq E_{z \sim q_{\phi}(z|x)}[\log p_{\theta}(x|z)] - D_{KL}(q_{\phi}(z|x), p(z))$$

i.e., **variational lower bound** on the **data likelihood**  $p_{\theta}(x)$

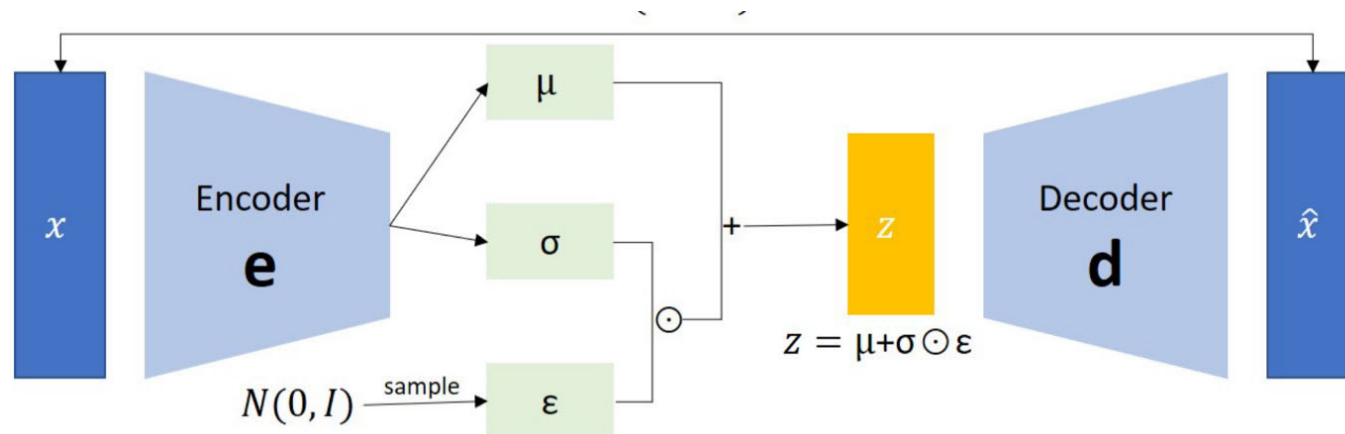
# Summary:

## From Autoencoder to Variational Autoencoder

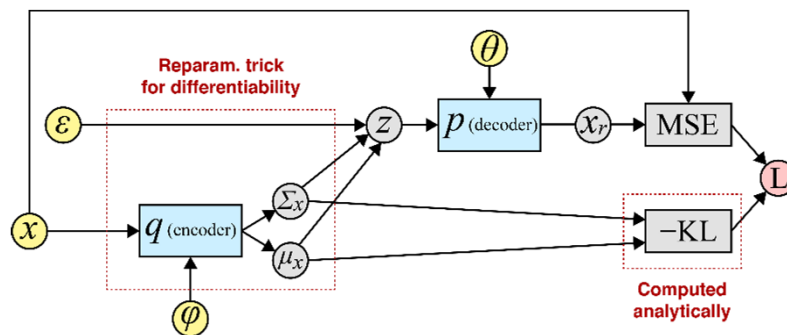
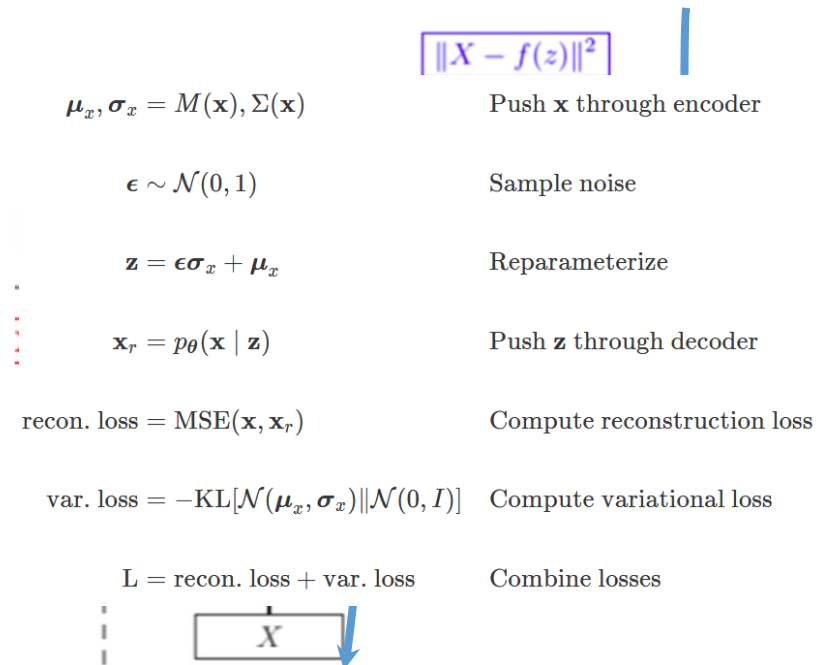


# Reparameterization Trick in VAE

- Remarks
  - Given  $x$ , sample  $z$  from latent distribution (described by output parameters of encoder)
  - However, this creates a bottleneck since **backpropagation cannot flow through**
  - Alternatively, we apply  $z = \mu + \sigma \odot \varepsilon$  ( $\varepsilon$  simply generated by **Normal distribution**).
  - This enables BP gradients in encoder through  $\mu$  and  $\sigma$ , while maintaining stochasticity via  $\varepsilon$  (for generative model purposes).



# Implementation of VAE



**Initialize** parameters of encoder and decoder

**Repeat:**

Get mini-batch of  $\mathbf{X}$

$\mathbf{mu\_X}, \mathbf{var\_X} = \text{encoder}(\mathbf{X})$

$\epsilon = \text{sampling from Normal}(\mathbf{0}, \mathbf{I})$

$\mathbf{z} = \mathbf{mu\_X} + \epsilon * \mathbf{var\_X}$

$\mathbf{X}' = \text{decoder}(\mathbf{z})$

$\text{recon\_loss} = \text{MSE}(\mathbf{X}, \mathbf{X}')$

$\text{latent\_loss} =$

$\text{KLD}(\text{Normal}(\mathbf{mu\_X}, \mathbf{var\_X}) || \text{Normal}(\mathbf{0}, \mathbf{I}))$

$\text{all\_loss} = \text{recon\_loss} + \text{latent\_loss}$

$\text{all\_loss.backward}()$


**Until:** parameters of encoder & decoder converge

**Return** parameters of encoder and decoder

First sample noise  $\epsilon$  from  $\text{Normal}(\mathbf{0}, \mathbf{I})$ , then reparameterize  $\mathbf{z}$  by  $\mathbf{mu\_X} + \epsilon * \mathbf{var\_X}$ , (equivalently sampled  $\text{Normal}(\mathbf{mu\_X}, \mathbf{var\_X})$ ). The model is now differentiable!



Before We Move On...



$$\log p_{\theta}(x) = \log \frac{p_{\theta}(x|z)p(z)}{p_{\theta}(z|x)} = \log \frac{\boxed{p_{\theta}(x|z)} \boxed{p(z)} \boxed{q_{\phi}(z|x)}}{\boxed{p_{\theta}(z|x)} \boxed{q_{\phi}(z|x)}}$$

$$= E_z [\log \boxed{p_{\theta}(x|z)}] - E_z \left[ \log \frac{\boxed{q_{\phi}(z|x)}}{\boxed{p(z)}} \right] + E_z \left[ \log \frac{\boxed{q_{\phi}(z|x)}}{\boxed{p_{\theta}(z|x)}} \right]$$

$$= E_{z \sim q_{\phi}(z|x)} [\log p_{\theta}(x|z)] - D_{KL} \left( q_{\phi}(z|x), p(z) \right) + D_{KL} (q_{\phi}(z|x), p_{\theta}(z|x))$$

Data reconstruction

KL divergence  
between sample  
distribution from the  
encoder and the prior


KL divergence between  
sample distribution  
from the encoder and  
the posterior of data

$$\Rightarrow \log p_{\theta}(x) \geq E_{z \sim q_{\phi}(z|x)} [\log p_{\theta}(x|z)] - D_{KL} \left( q_{\phi}(z|x), p(z) \right)$$

i.e., **variational lower bound** on the **data likelihood**  $p_{\theta}(x)$


# From Autoencoder to Variational Autoencoder (cont'd)

- Example Results



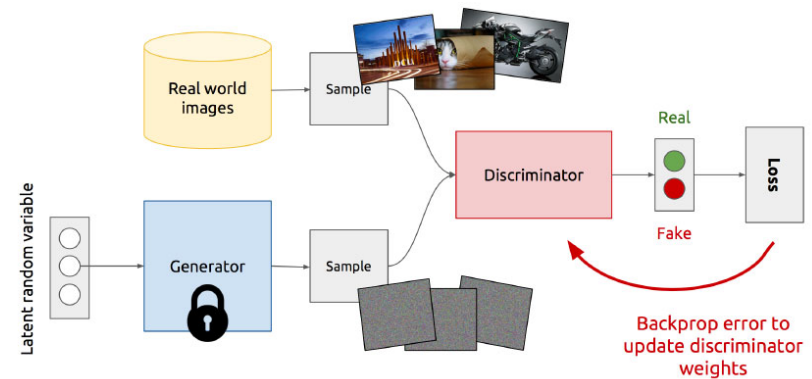


# From Autoencoder to Variational Autoencoder (cont'd)

- Example Results
  - $A' - A + B = B'$

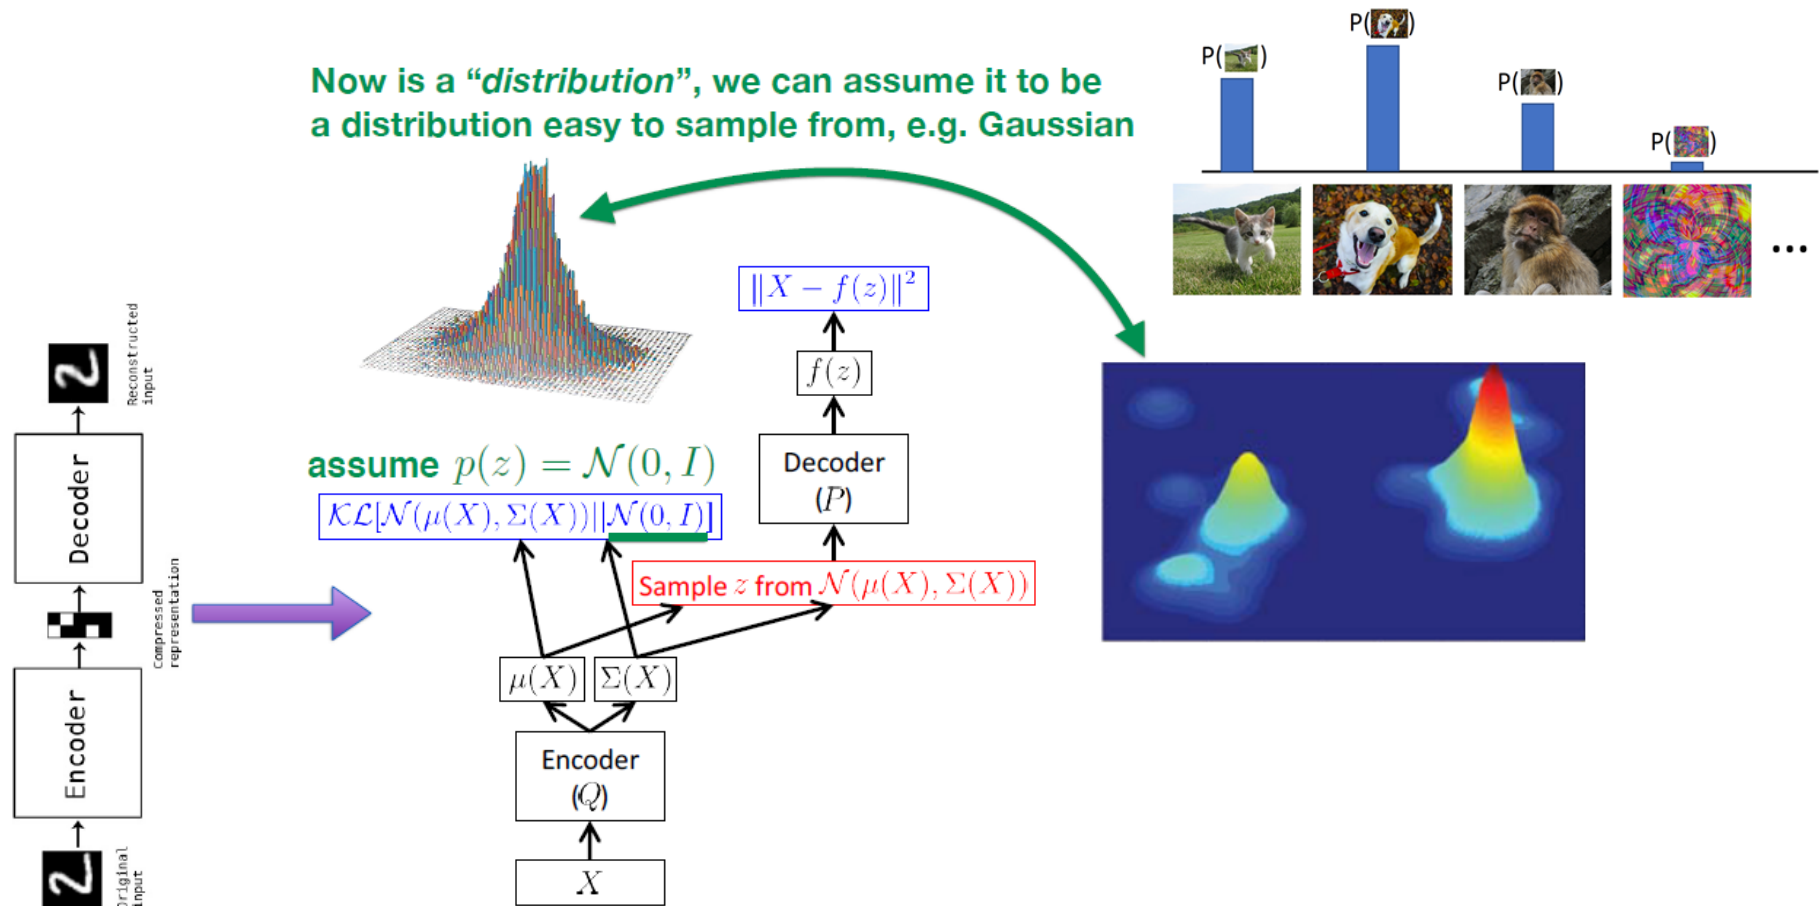


# What's to Be Covered Today...

- Generative Models
  - Auto-Encoder vs. Variational Auto-Encoder
  - Generative Adversarial Network (GAN)
  - Diffusion Model
- HW #1 is due Oct. 10<sup>th</sup> Mon 23:59
- HW #2 will be out next week...

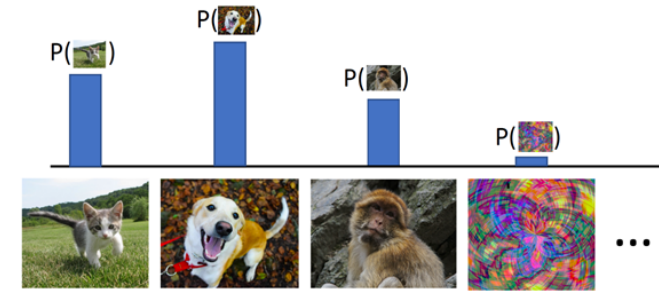


# From VAE to *Generative Adversarial Networks (GAN)*



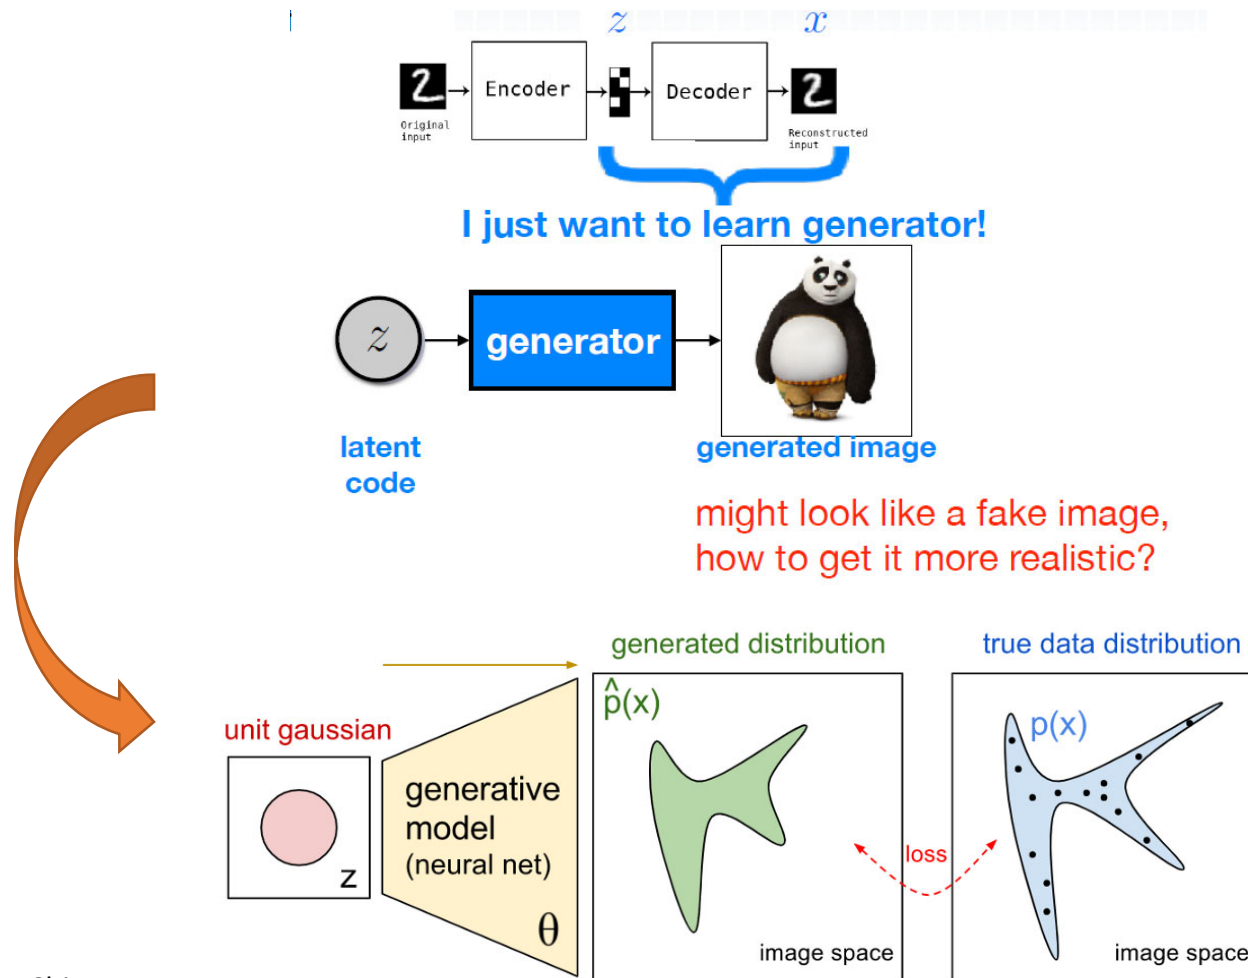


# From VAE to GAN (cont'd)



- Remarks

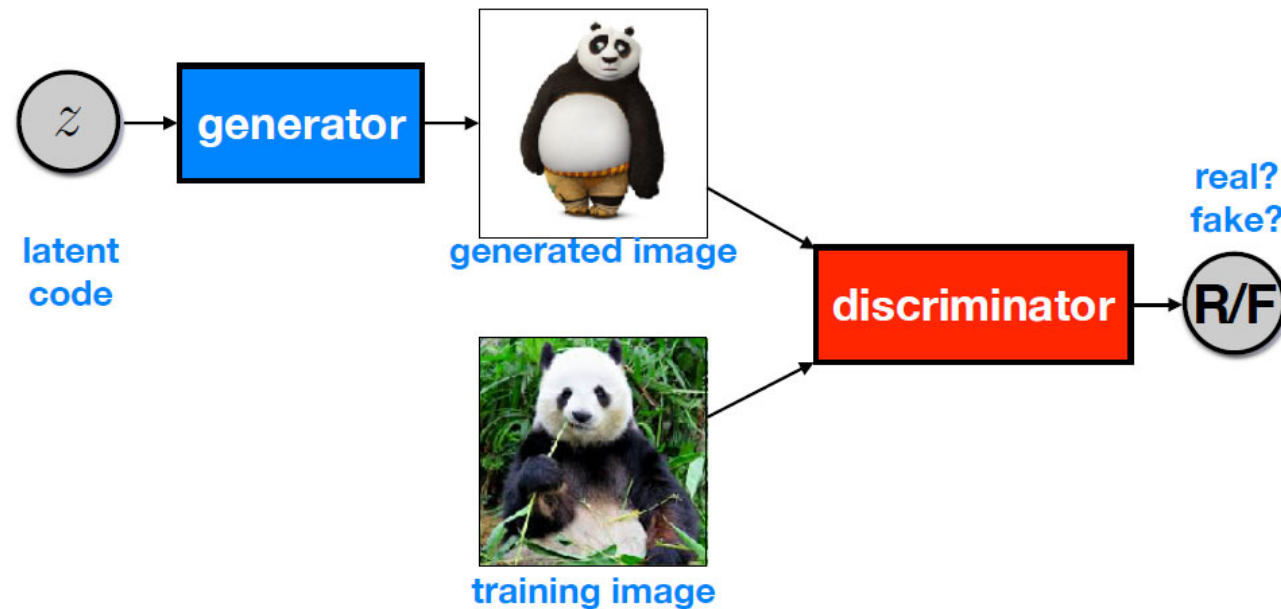
- What if we only need the decoder/generator in practice?
- How do we know if the output images are sufficiently good?





# Generative Adversarial Network

- Idea
  - **Generator** to convert a vector  $z$  (sampled from  $P_z$ ) into fake data  $x$  (from  $P_G$ ), while we need  $P_G = P_{\text{data}}$
  - **Discriminator** classifies data as real or fake (1/0)
  - How? Impose an **adversarial loss** on the observed data distribution!

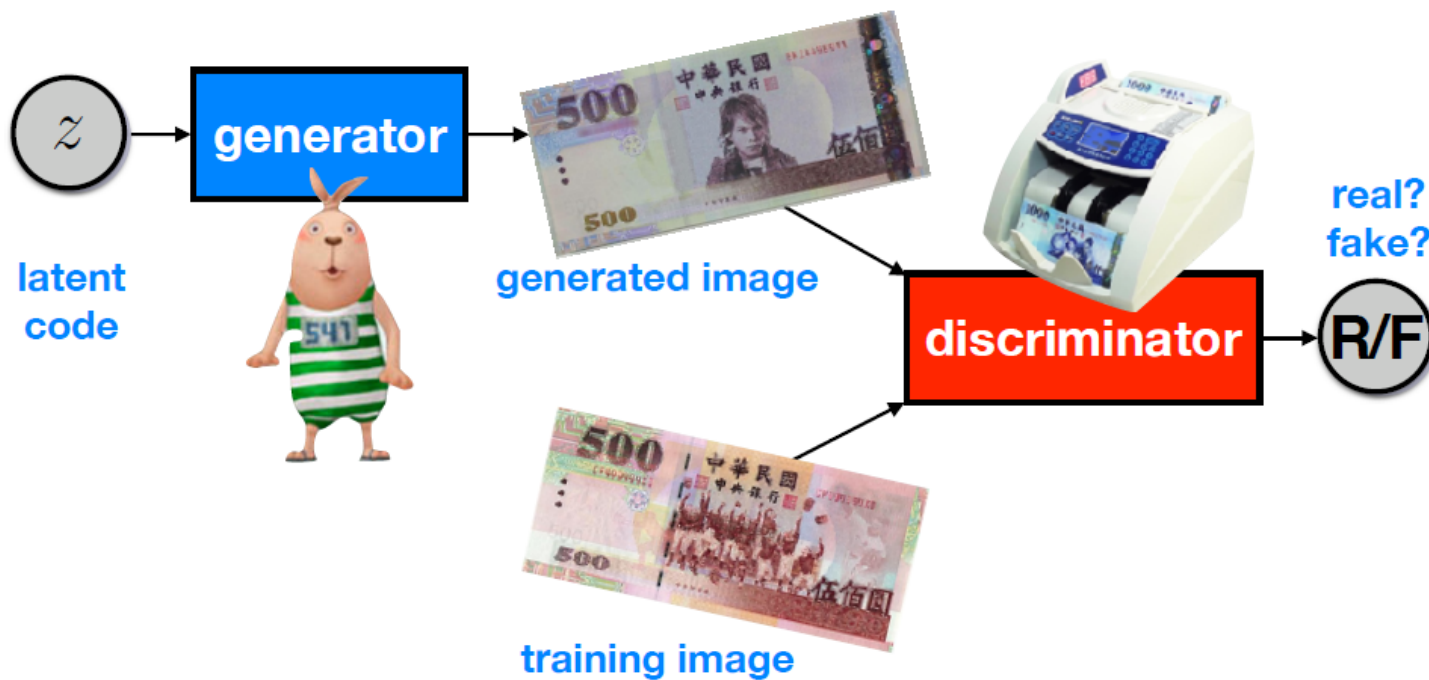


# Generative Adversarial Network (cont'd)

- Idea
  - Impose adversarial loss on data distribution
  - Let's see a practical example...

generator: try to generate more realistic images to cheat discriminator

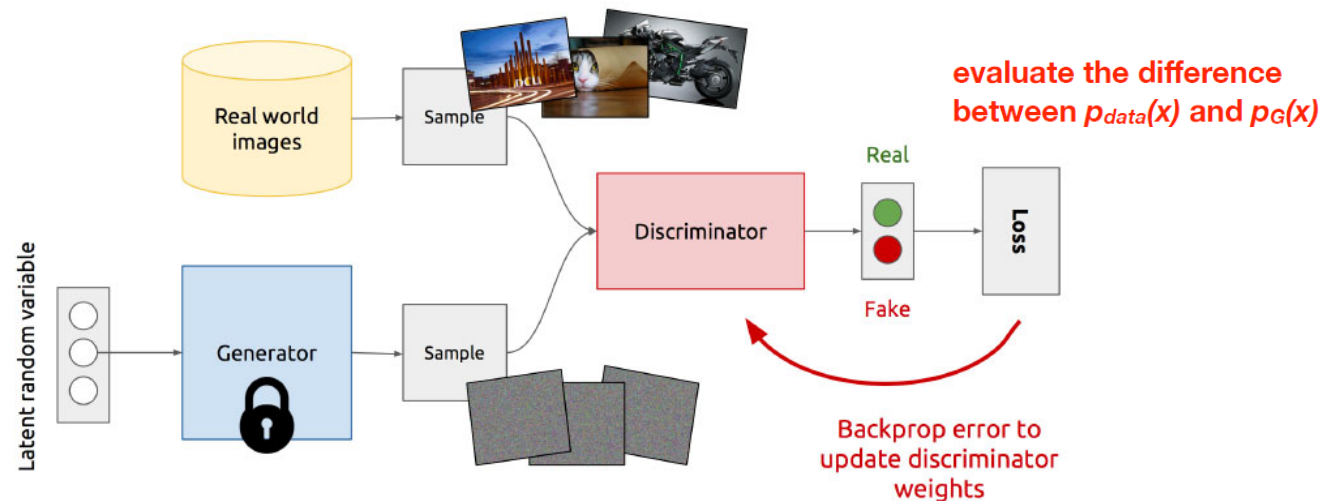
discriminator: try to distinguish whether the image is generated or real



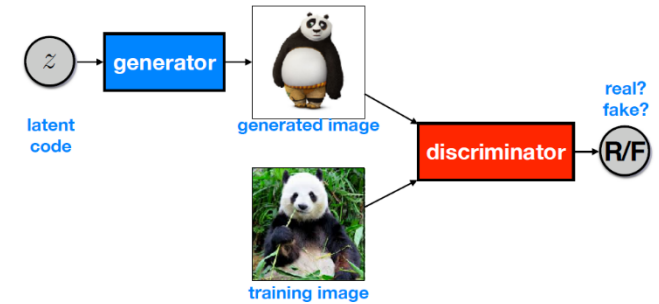
# GAN (cont'd)

- Remarks
  - A function maps **normal distribution**  $N(\mathbf{0}, I)$  to  $P_{data}$
  - How good we are in mapping  $P_g$  to  $P_{data}$ ?
    - Train & ask the discriminator!
  - Conduct a two-player min-max game

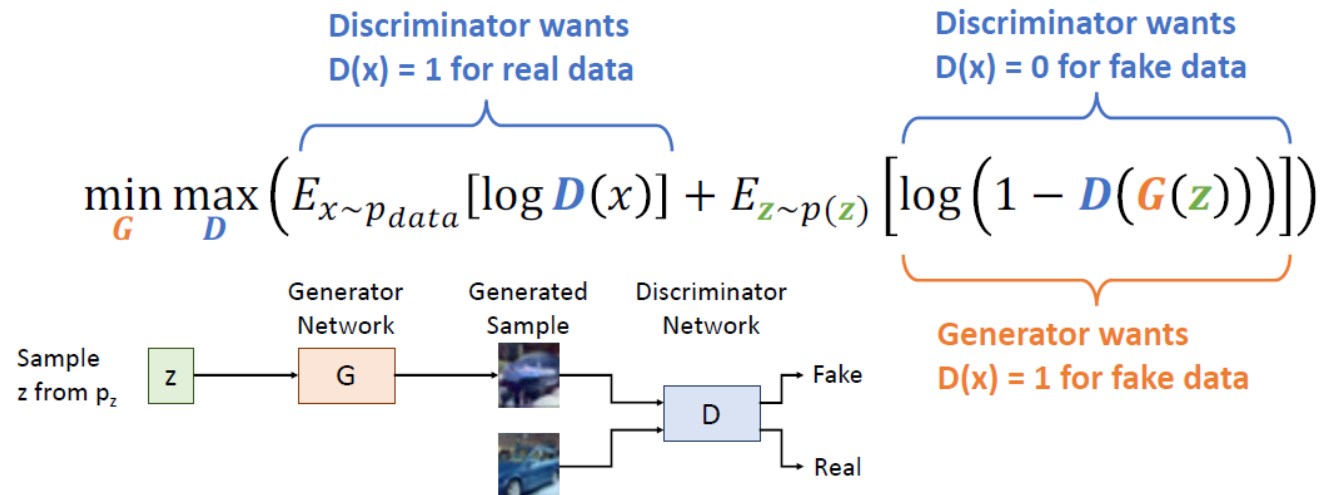
$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$



# Training Objective of GAN



- Jointly train generator G and discriminator D with a min-max game



- Train G & D with alternating gradient updates

$$\min_G \max_D V(G, D)$$

For t in 1, ... T:


- (Update D)  $D = D + \alpha_D \frac{\partial V}{\partial D}$
- (Update G)  $G = G - \alpha_G \frac{\partial V}{\partial G}$

## Training Objective of GAN (*optional trick*)

- Potential Problem

$$\min_G \max_D \left( E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))] \right)$$

- At start of training, G is not OK yet (obviously);  
D easily tells apart real/fake data (i.e.,  $D(G(z))$  close to 0).
- Solution:**
  - Instead of training G to minimize  $\log(1-D(z))$  in the beginning, we train G to minimize  $-\log(D(G(z)))$ .
  - With strong gradients from G, we start the training of the above min-max game.



# Optimality of GAN

- Why the min-max game as objective a good idea?

$$\begin{aligned} & \min_G \max_D \left( E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))] \right) \\ &= \min_G \max_D \left( E_{x \sim p_{data}} [\log D(x)] + E_{x \sim p_G} [\log (1 - D(x))] \right) \\ &= \min_G \int_X \max_D (p_{data}(x) \log D(x) + p_G(x) \log (1 - D(x))) dx \\ & \left. \begin{aligned} f(y) &= a \log y + b \log(1 - y) \\ f'(y) &= \frac{a}{y} - \frac{b}{1 - y} \end{aligned} \right\} f'(y) = 0 \Leftrightarrow y = \frac{a}{a+b} \text{ (local max)} \\ & \hspace{10em} \downarrow \\ & \text{Optimal Discriminator: } D_G^*(x) = \frac{p_{data}(x)}{p_{data}(x) + p_G(x)} \end{aligned}$$



## Optimality of GAN

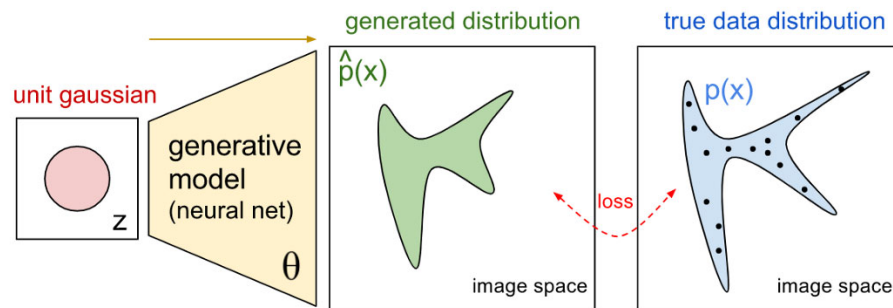
- Why the min-max game as objective a good idea? (cont'd)

$$\begin{aligned} & \min_G \max_D \left( E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} \left[ \log \left( 1 - D(G(z)) \right) \right] \right) \\ \Rightarrow & \min_G \int_X \left( p_{data}(x) \log \frac{p_{data}(x)}{p_{data}(x) + p_G(x)} + p_G(x) \log \frac{p_G(x)}{p_{data}(x) + p_G(x)} \right) dx \\ & = \min_G \left( E_{x \sim p_{data}} \left[ \log \frac{2}{2} \frac{p_{data}(x)}{p_{data}(x) + p_G(x)} \right] + E_{x \sim p_G} \left[ \log \frac{2}{2} \frac{p_G(x)}{p_{data}(x) + p_G(x)} \right] \right) \\ & = \min_G \left( E_{x \sim p_{data}} \left[ \log \frac{2 * p_{data}(x)}{p_{data}(x) + p_G(x)} \right] + E_{x \sim p_G} \left[ \log \frac{2 * p_G(x)}{p_{data}(x) + p_G(x)} \right] - \log 4 \right) \end{aligned}$$

# Optimality of GAN

- Why the min-max game as objective a good idea? (cont'd)

$$\begin{aligned}
 & \min_G \max_D \left( E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))] \right) \\
 &= \min_G \left( E_{x \sim p_{data}} \left[ \log \frac{2 * p_{data}(x)}{p_{data}(x) + p_G(x)} \right] + E_{x \sim p_G} \left[ \log \frac{2 * p_G(x)}{p_{data}(x) + p_G(x)} \right] - \log 4 \right) \\
 &= \min_G \left( KL \left( p_{data}, \frac{p_{data} + p_G}{2} \right) + KL \left( p_G, \frac{p_{data} + p_G}{2} \right) - \log 4 \right)
 \end{aligned}$$



**Kullback-Leibler Divergence:**

$$KL(p, q) = E_{x \sim p} \left[ \log \frac{p(x)}{q(x)} \right]$$

# Optimality of GAN

- Why the min-max game as objective a good idea? (cont'd)

$$\begin{aligned}
 & \min_G \max_D \left( E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))] \right) \\
 &= \min_G \left( E_{x \sim p_{data}} \left[ \log \frac{2 * p_{data}(x)}{p_{data}(x) + p_G(x)} \right] + E_{x \sim p_G} \left[ \log \frac{2 * p_G(x)}{p_{data}(x) + p_G(x)} \right] - \log 4 \right) \\
 &= \min_G \left( KL \left( p_{data}, \frac{p_{data} + p_G}{2} \right) + KL \left( p_G, \frac{p_{data} + p_G}{2} \right) - \log 4 \right) \\
 &= \min_G (2 * JSD(p_{data}, p_G) - \log 4)
 \end{aligned}$$

JSD is always nonnegative, and zero only when the two distributions are equal!

Thus  $p_{data} = p_G$  is the global min, QED

**Jensen-Shannon Divergence:**

$$JSD(p, q) = \frac{1}{2} KL \left( p, \frac{p + q}{2} \right) + \frac{1}{2} KL \left( q, \frac{p + q}{2} \right)$$

## Remarks on Optimality of GAN

$$\min_{\underset{G}{G}} \max_{\underset{D}{D}} \left( E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} \left[ \log \left( 1 - D(G(z)) \right) \right] \right) \\ = \min_{\underset{G}{G}} (2 * JSD(p_{data}, p_G) - \log 4)$$

- Summary

- The global min of the minmax game happens when

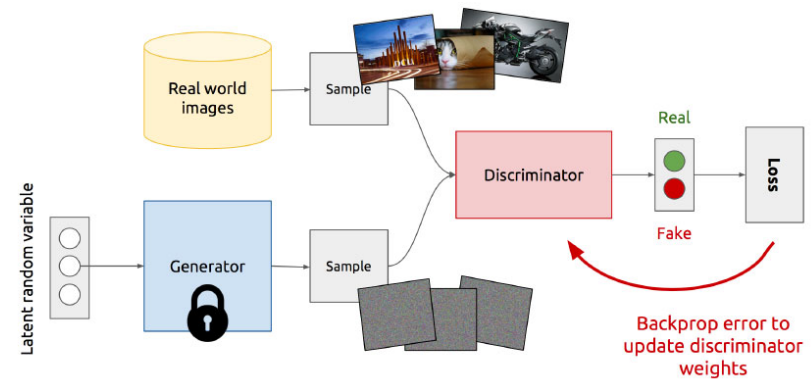


1.  $D_G^*(x) = \frac{p_{data}(x)}{p_{data}(x) + p_G(x)}$  (Optimal discriminator for any G) ➡
2.  $p_G(x) = p_{data}(x)$  (Optimal generator for optimal D) ➡

- **Caution!**

- G and D are learned models (i.e., DNNs) with fixed architectures.  
We don't know whether we can actually represent the **optimal D & G**.
- Optimality of GAN does not tell anything about **convergence** to the optimal D/G.

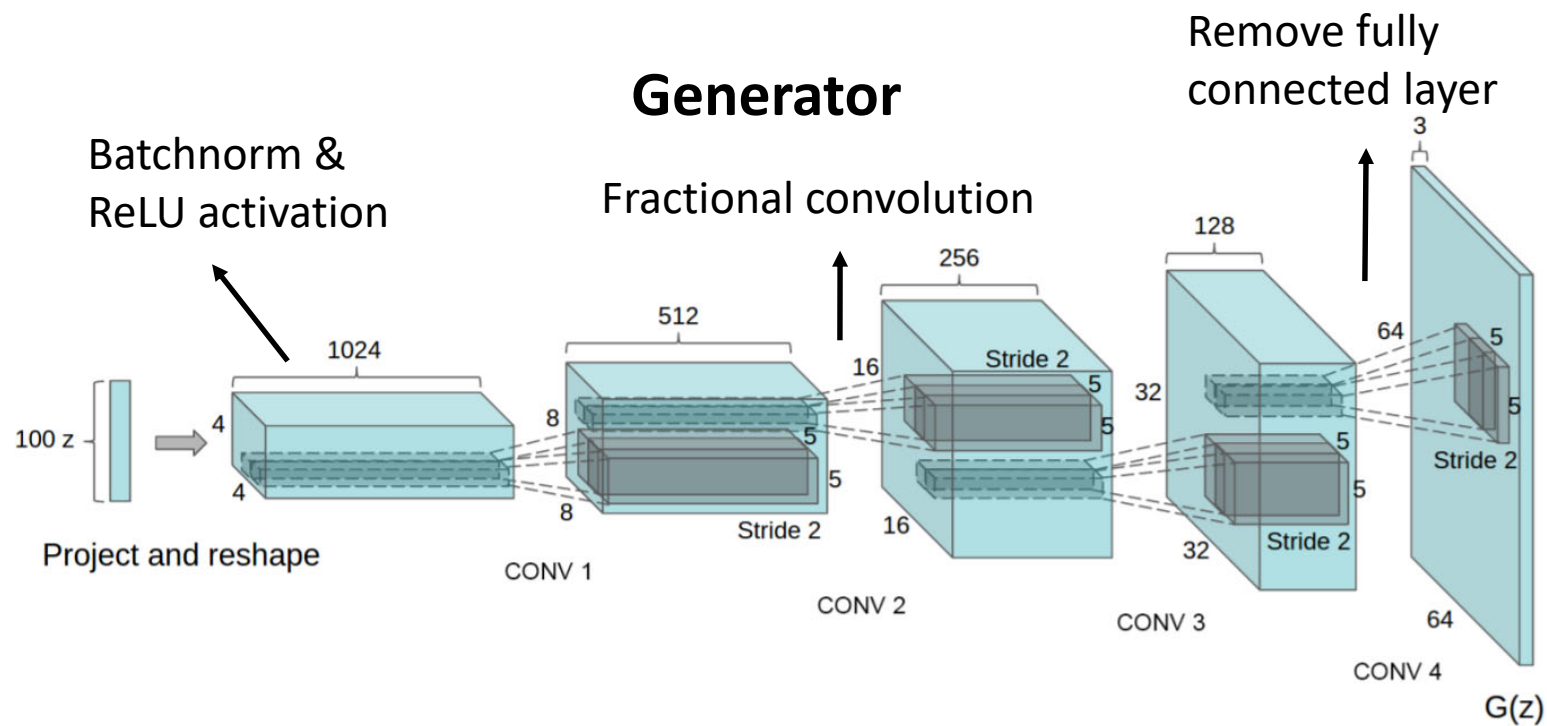
# What's to Be Covered Today...

- Generative Models
  - Auto-Encoder vs. Variational Auto-Encoder
  - Generative Adversarial Network (GAN)
    - Challenges & Variants of GAN
  - Diffusion Model
- HW #1 is due Oct. 10<sup>th</sup> Mon 23:59
- HW #2 will be out next week...



# Deep Convolutional GAN (DC-GAN)

- Remarks
  - ICLR 2016
  - A CNN+GAN architecture
  - Empirically make training of GAN more stable






# Deep Convolutional GAN (DC-GAN)

- Example Results




Collected face dataset



LSUN dataset

# Conditional GANs

- Remarks
  - ICLR 2016
  - Conditional generative model  $p(x|y)$  instead of  $p(x)$
  - Both G and D take the label  $y$  as an additional input...Why?  
Why not just use D as designed in the standard GAN?



# Conditional GANs

- Example Results


Welsh springer spaniel



Fire truck




Daisy




## Problems in Training GANs: Vanishing Gradients

- What Might Go Wrong?
  - GAN training is often unstable.
  - In other words, training might not converge properly.
  - The discriminator which we prefer is...



0



0.2



0.6




1




## Problems in Training GANs: Vanishing Gradients (cont'd)

- What Might Go Wrong?
  - GAN training is often unstable.
  - In other words, training might not converge properly.
  - The discriminator we trained might be as follows.  
In other words, no gradient to guide the generator to output proper images.



0



0



0




1

- This is known as the problem of *vanishing gradients*.

# Problems in Training GANs: Mode Collapse

- Remarks
  - The generator only outputs a limited number of image variants regardless of the inputs.





## Problems in Training GANs: Mode Collapse (cont'd)

- Remarks
  - The generator only outputs a limited number of image variants regardless of the inputs.





Photo credit:  
<https://openreview.net/pdf?id=rkmu5b0a->



## Problems in Training GANs: Mode Collapse (cont'd)

- Why Mode Collapse Happens?
  - The objective of GANs assesses the **image authenticity**, not **diversity**.
  - Imbalance training between generator/discriminator (exploding/vanishing gradients)



# Energy-Based GAN

- Energy Function
  - Converting input data into scalar outputs, viewed as energy values
  - Desired configuration is expected to output low energy values & vice versa.
- Energy Function as Discriminator
  - Use of autoencoder; can be pre-trained!
  - Reconstruction loss outputs a range of values instead of binary logistic loss.
  - Empirically better convergence




# EB-GAN

- Example Results



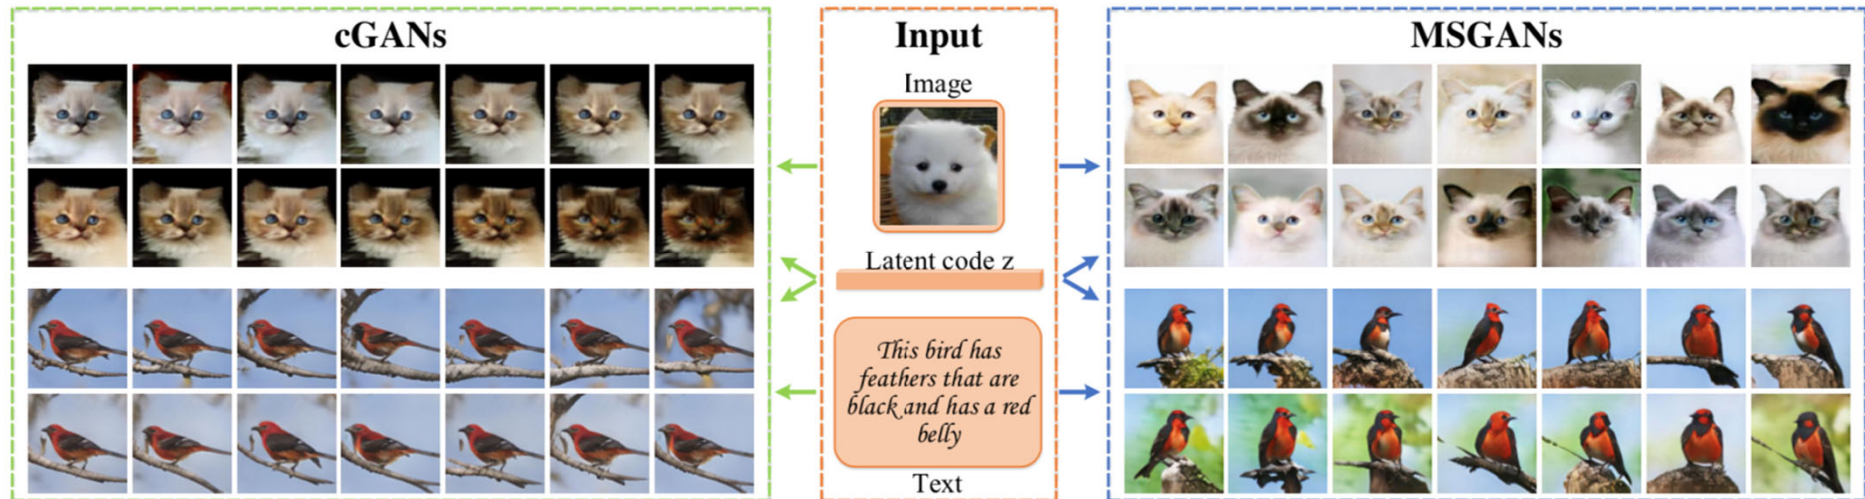
DCGAN



EBGAN

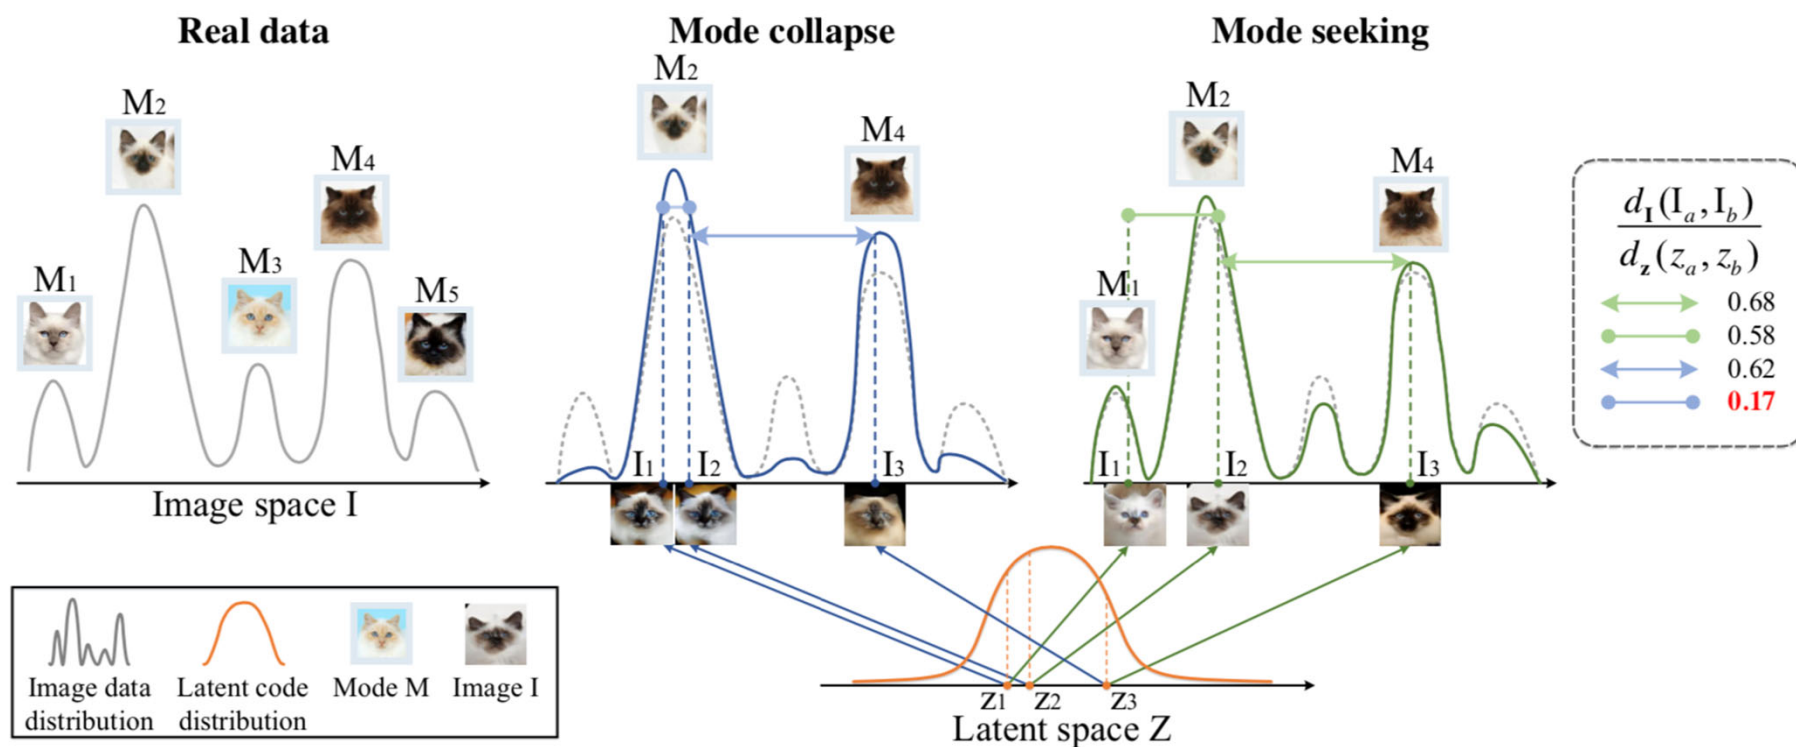
# MSGAN

- Mode Seeking Generative Adversarial Networks for Diverse Image Synthesis
- With the goal of producing **diverse** image outputs.
- To address the **mode collapse** issue by **conditional GANs**



# MSGAN (cont'd)

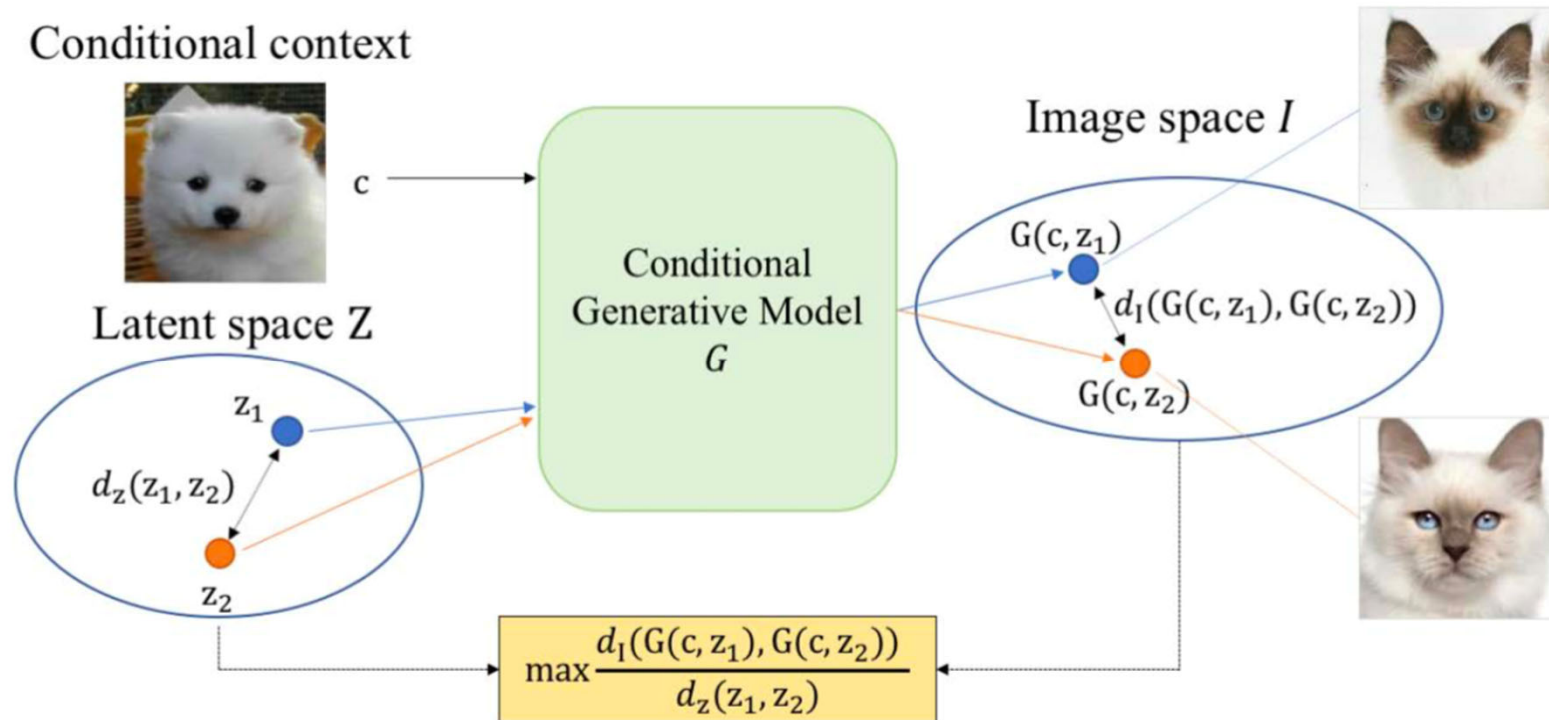
- Motivation (for unconditional GAN)





# MSGAN (cont'd)

- Proposed Regularization (for conditional GAN)



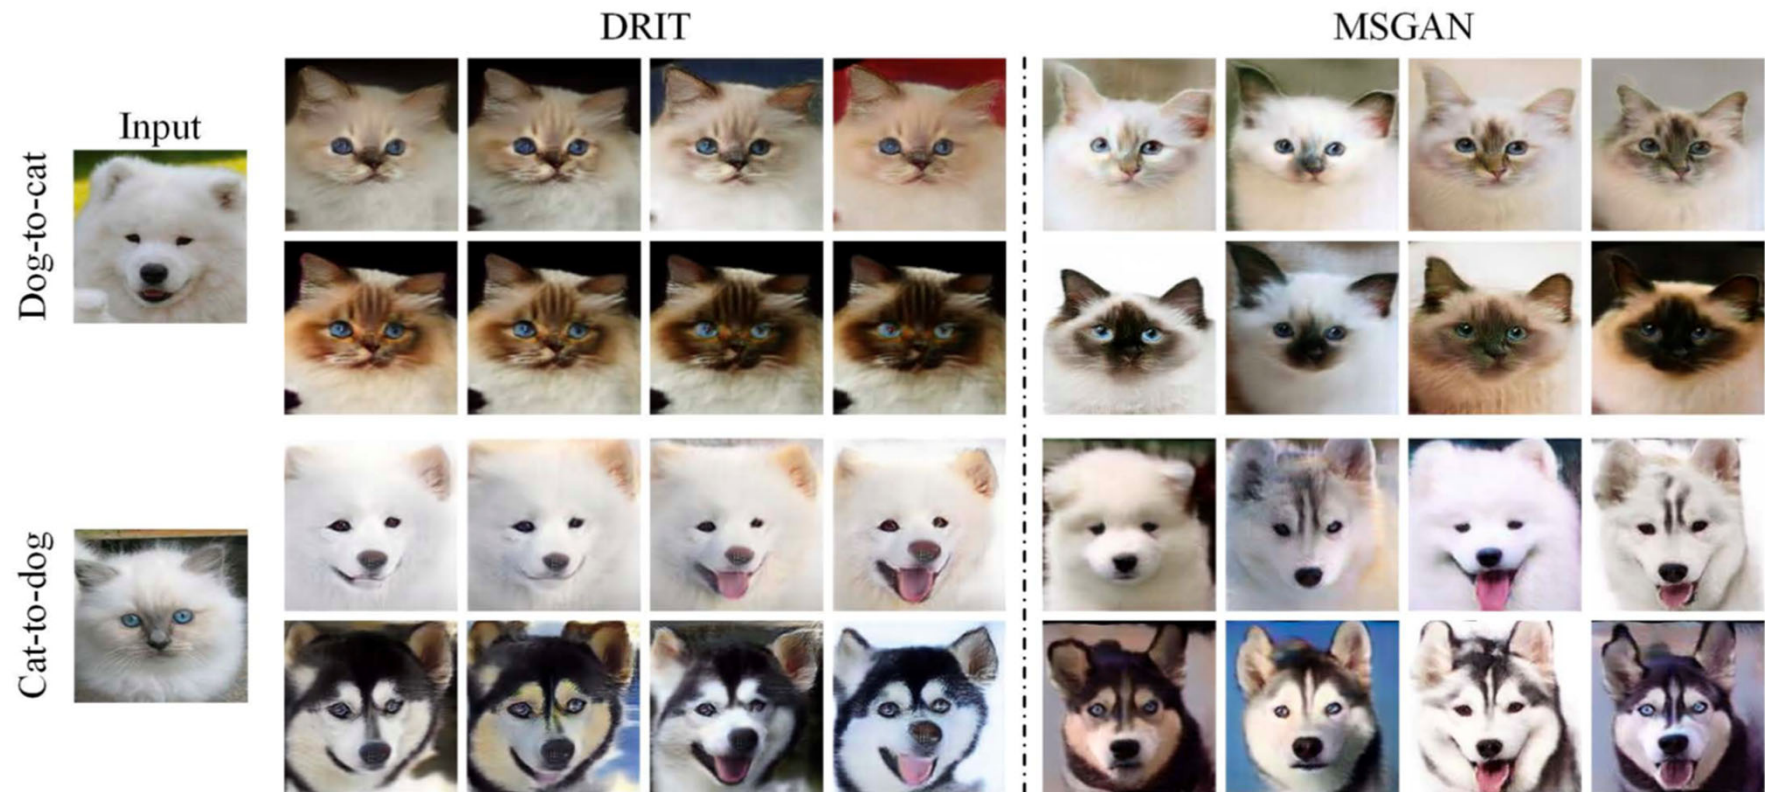
# MSGAN (cont'd)

- Qualitative results
  - Conditioned on [paired images](#)



# MSGAN (cont'd)

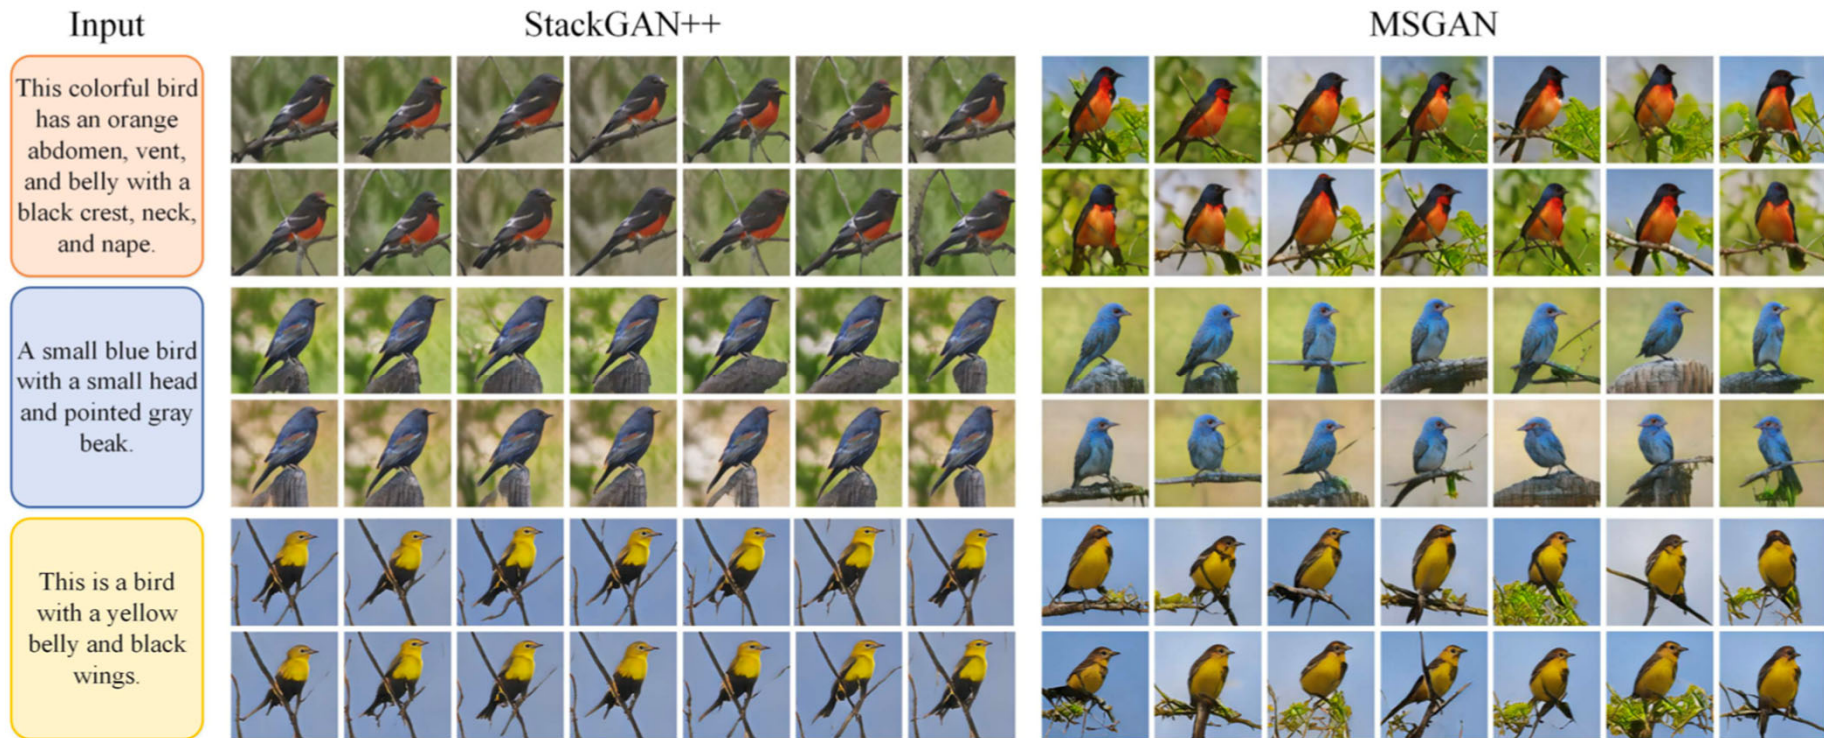
- Qualitative results
  - Conditioned on **unpaired images**





# MSGAN (cont'd)

- Qualitative results
  - Conditioned on **text** (will talk about Vision & Language later this semester)



# Style-based GAN (if time permits)

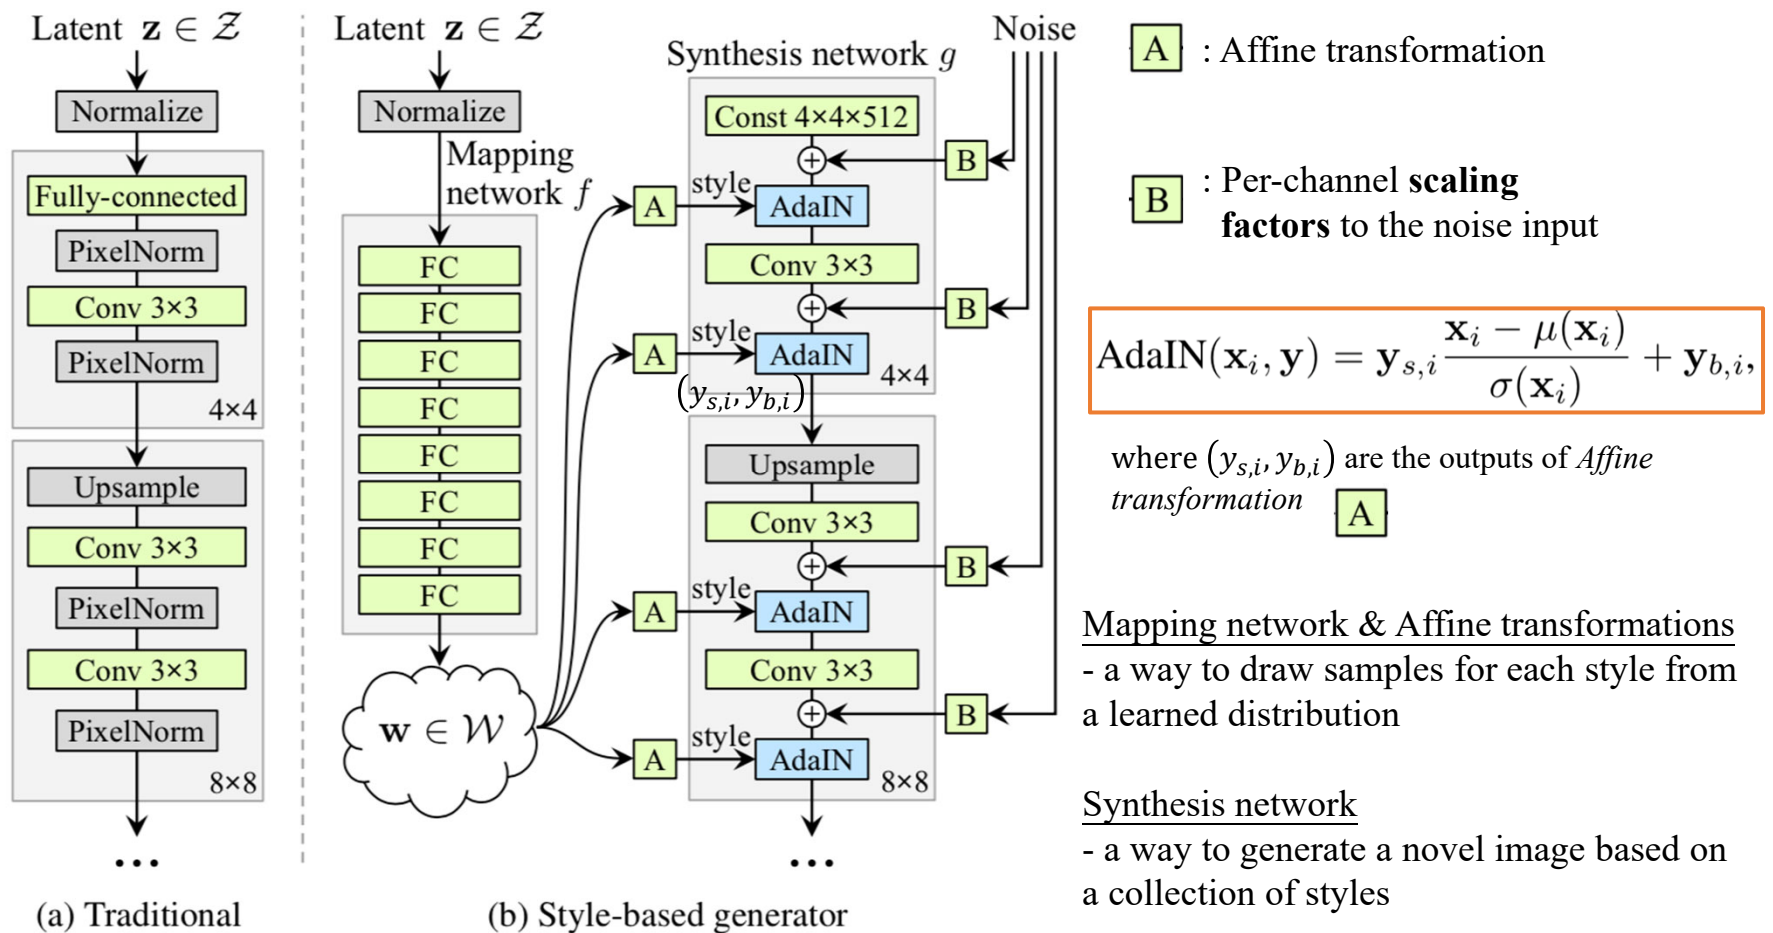
- A Style-Based Generator Architecture for Generative Adversarial Networks (CVPR'19)
- Design **style-based generator** to achieve **high-resolution** image synthesis
- No particular designs on loss functions, regularization, and hyper-parameters





# Style-based GAN (cont'd)

- Style-based generator



# Style-based GAN (cont'd)


- Qualitative results







# SinGAN: Learning a Generative Model from a Single Natural Image

- ICCV 2019 Best Paper Award
- Remarks:
  - Learning from a **single image**
  - Handle **multiple image manipulation tasks**
    - Super-resolution, style conversion, harmonization, image editing, et.




# SinGAN: Learning a Generative Model from a Single Natural Image

- Related Works
  - While single-image based learning models exist, most existing methods are designed to handle textural images but not natural ones.



# SinGAN: Learning a Generative Model from a Single Natural Image

- Goal
  - Output images with **arbitrary sizes** and **aspect ratios** (via fully conv models) by changing dimensions of noise and the input size

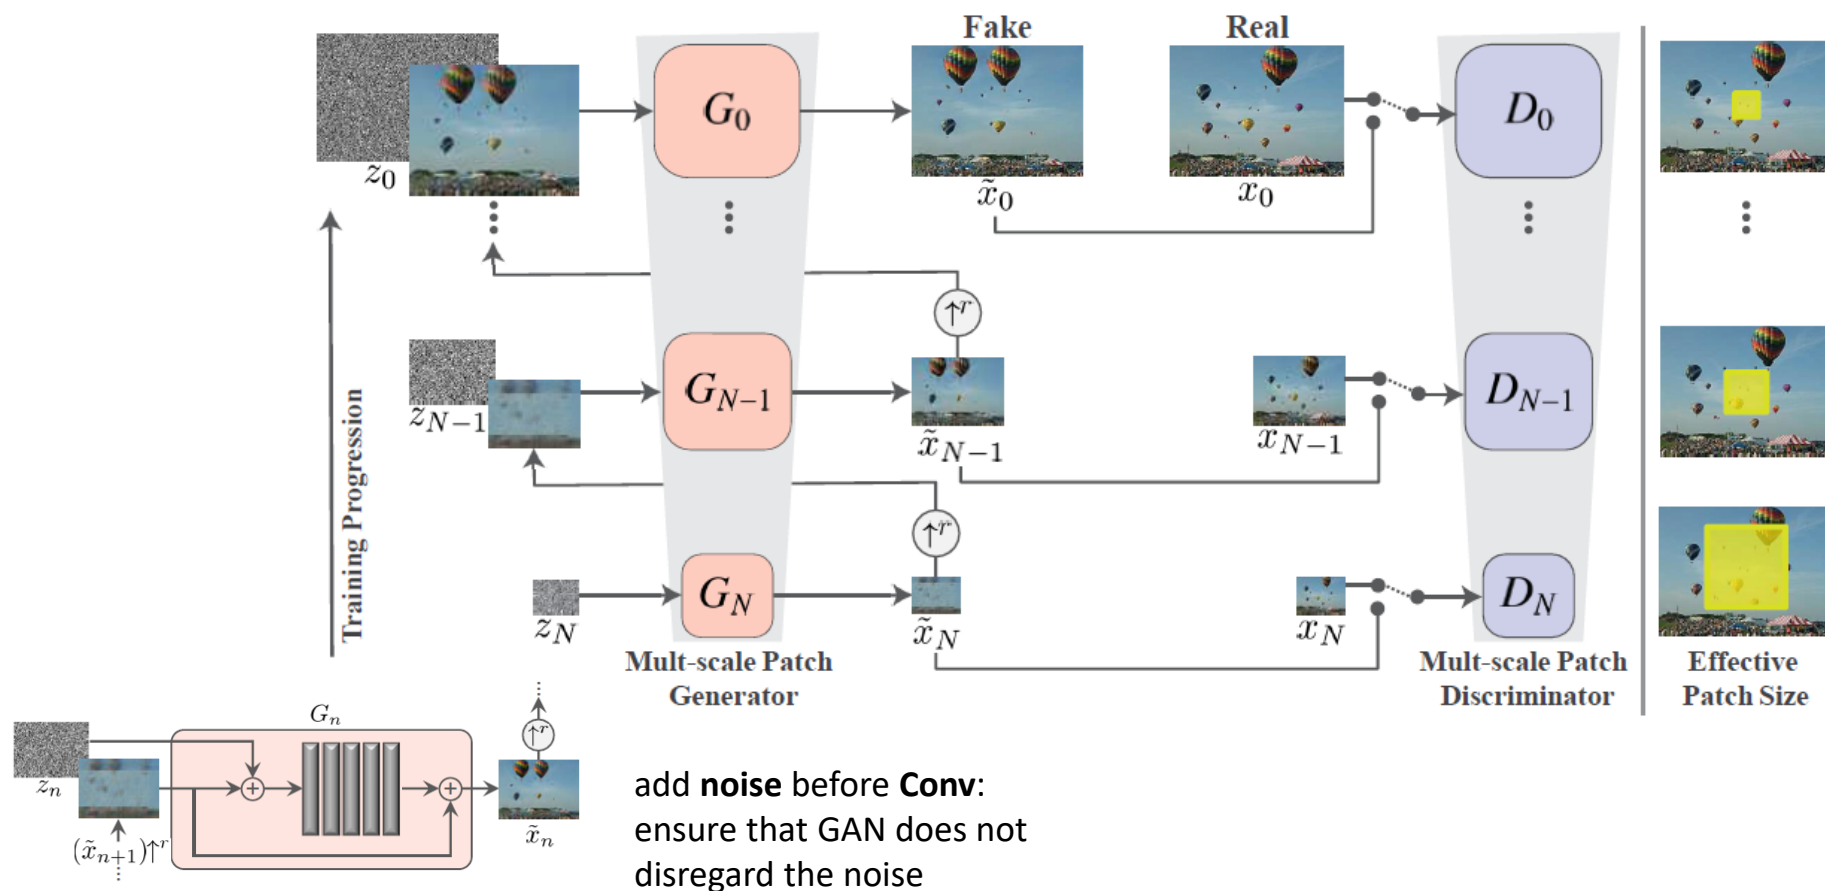


# SinGAN: Learning a Generative Model from a Single Natural Image

- Framework

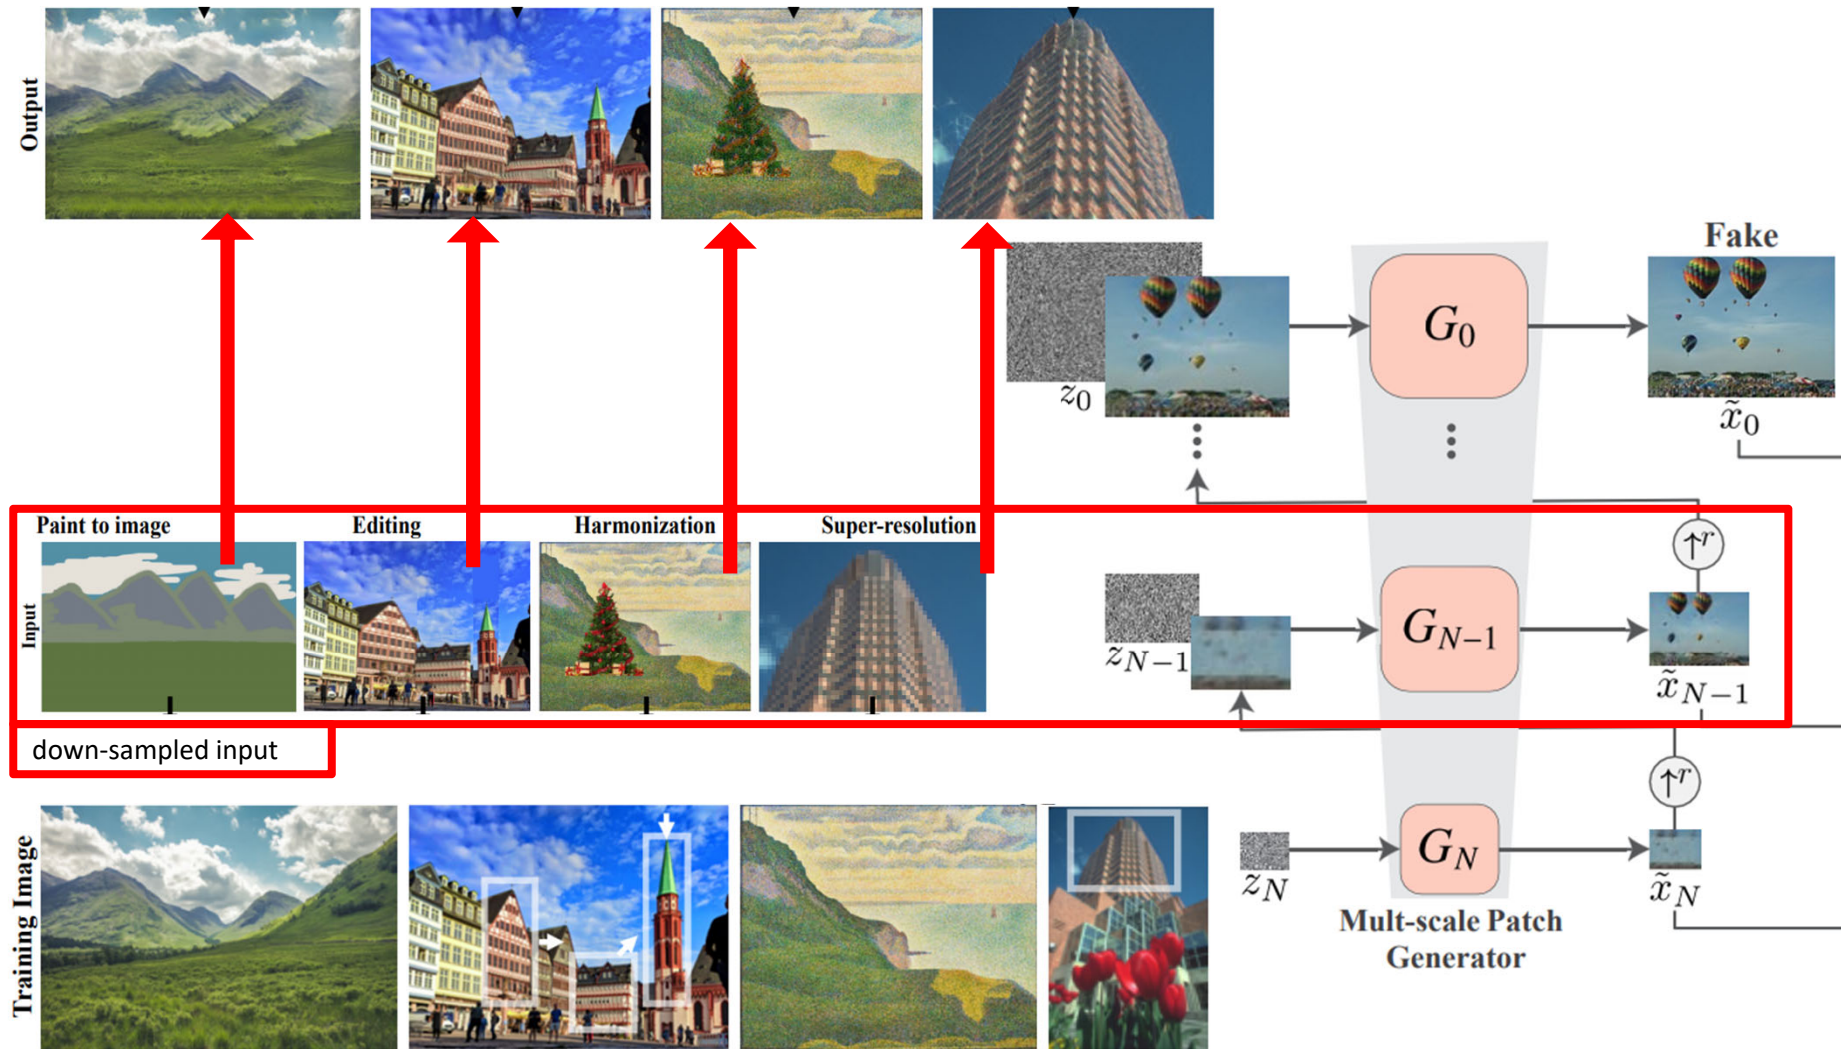
$$\min_{G_n} \max_{D_n} \mathcal{L}_{\text{adv}}(G_n, D_n) + \alpha \mathcal{L}_{\text{rec}}(G_n)$$

fix kernel (receptive field) size at each scale:  
capture structures of decreasing size as we go up



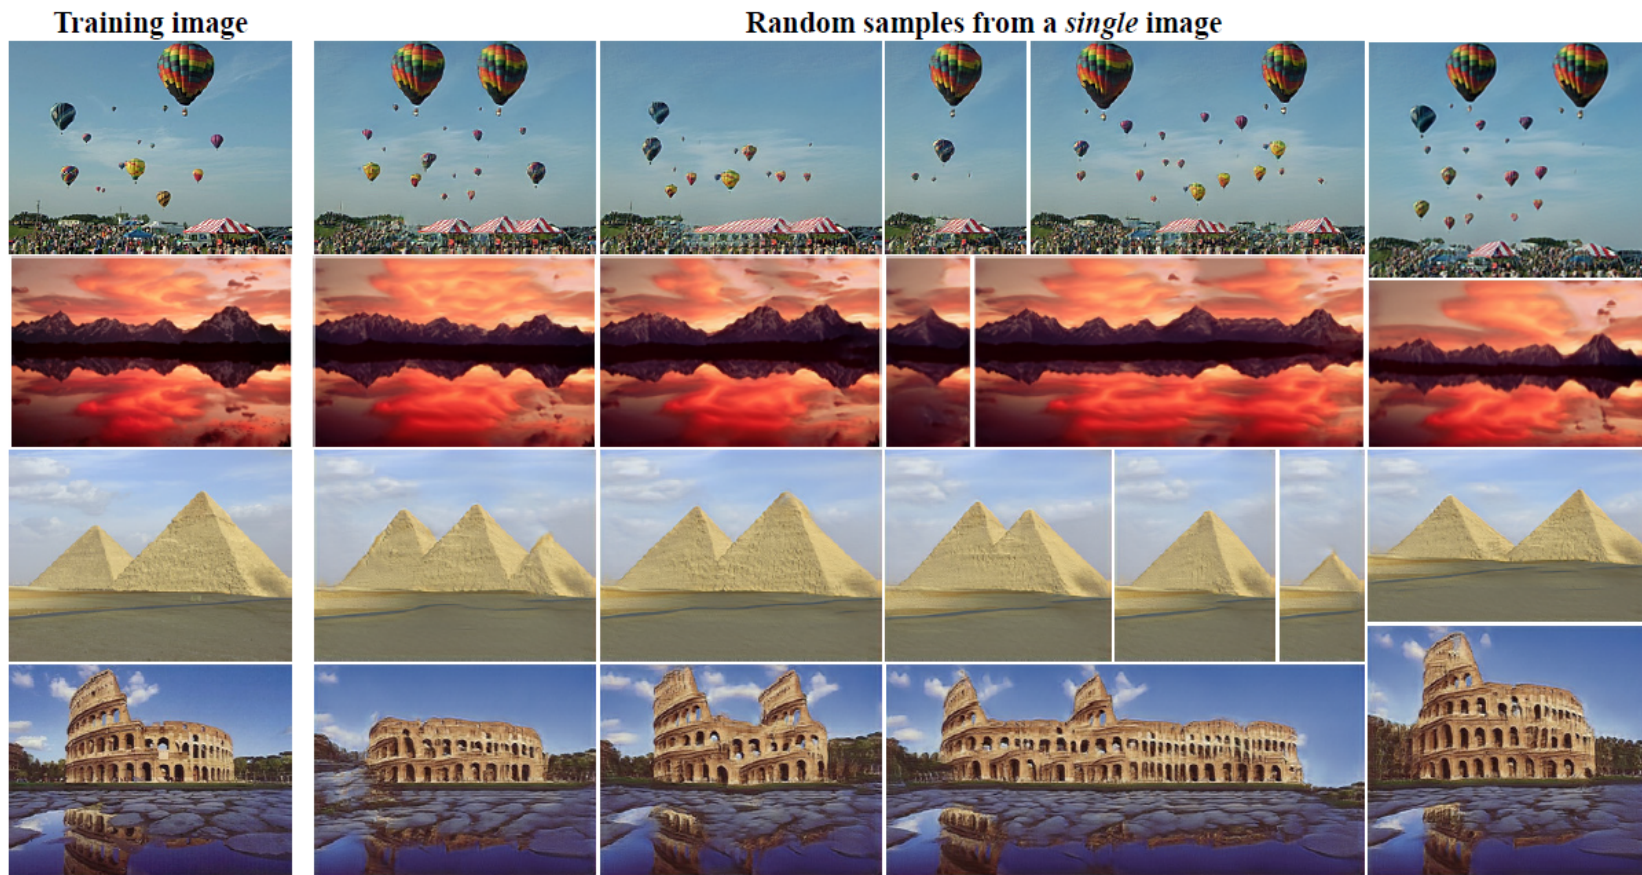


# Inference Stage for SinGAN



# SinGAN: Learning a Generative Model from a Single Natural Image

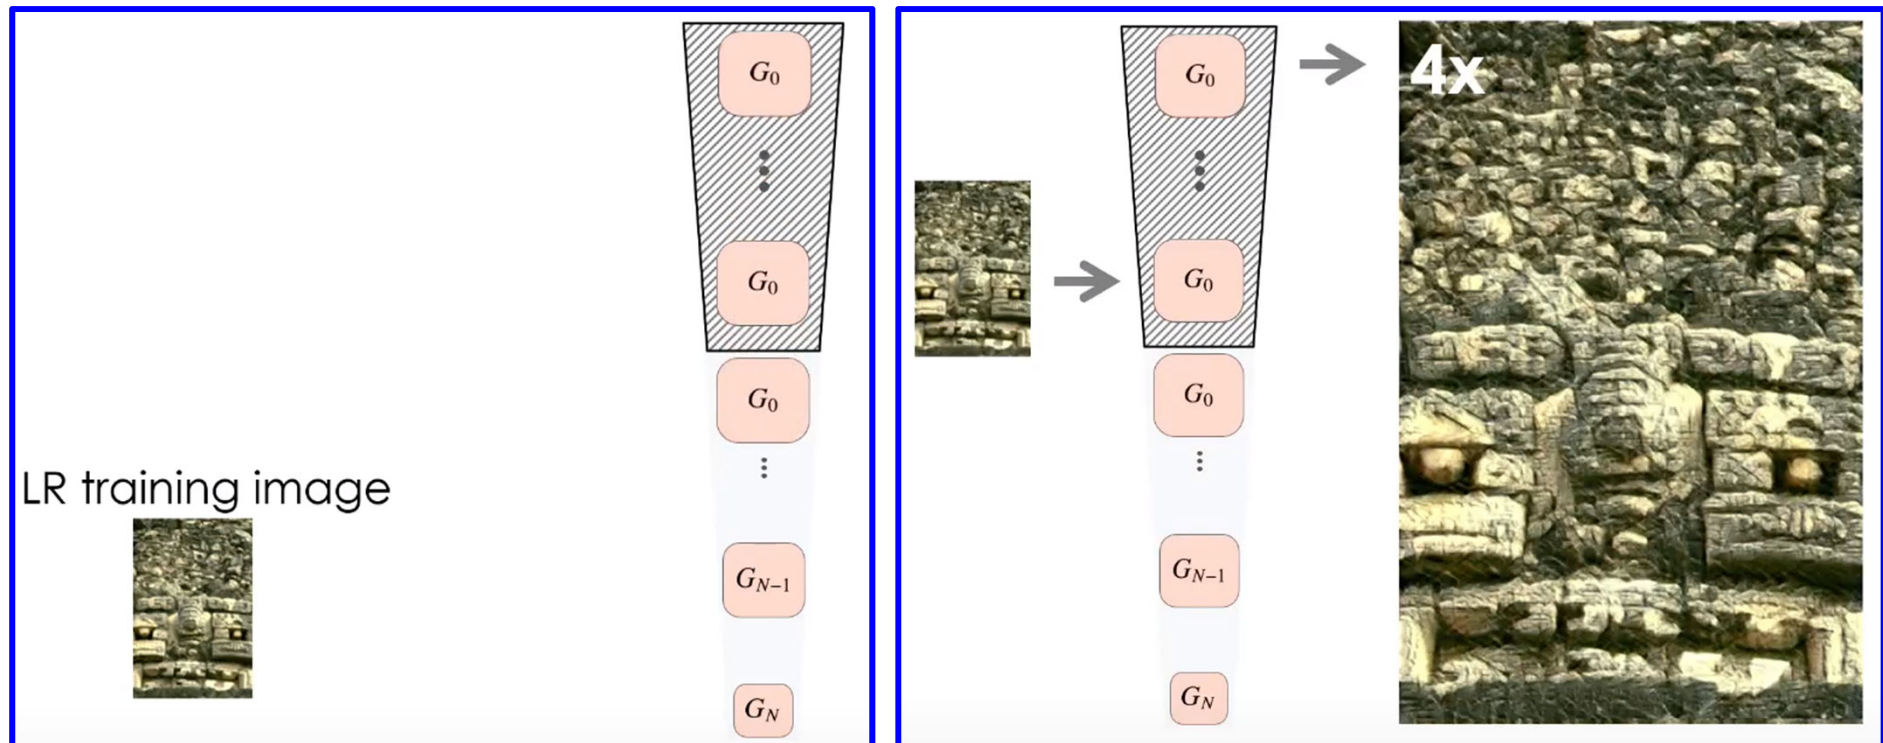
- Random image generation





# SinGAN: Learning a Generative Model from a Single Natural Image

- Super-Resolution





# SinGAN: Learning a Generative Model from a Single Natural Image


- Super-Resolution

	External methods		Internal Methods		
	SRGAN	EDSR	DIP	ZSSR	SinGAN
NIQE	3.4	6.5	6.3	7.1	3.7



*trained on a dataset*

*trained on a single image*



# SinGAN: Learning a Generative Model from a Single Natural Image

- Paint-to-Image

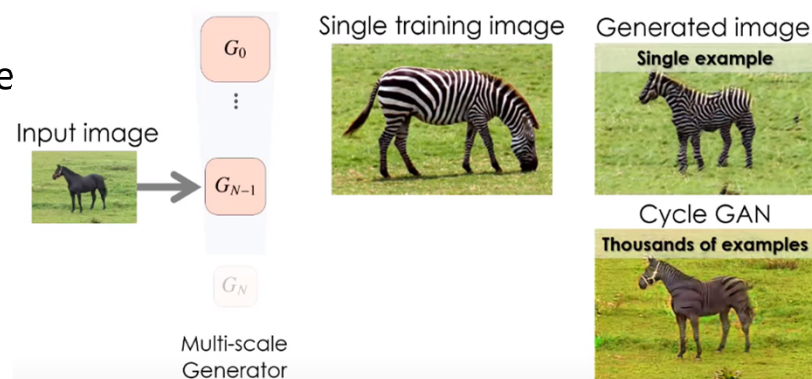



Image	Injection scale	Total number of scales
Balloons1	$n = 7$	$N = 9$
Balloons2	$n = 5$	$N = 9$
Starry night	$n = 6$	$N = 8$
Rock	$n = 6$	$N = 8$
Tree	$n = 6$	$N = 8$
Birds	$n = 5$	$N = 7$
View (Fig. 2, main text)	$n = 7$	$N = 8$
Pyramids (Fig. 11, main text)	$n = 6$	$N = 8$
cows (Fig. 11, main text)	$n = 5$	$N = 7$





# SinGAN: Learning a Generative Model from a Single Natural Image

- Harmonization

Image	Injection scale	Total number of scales
Tree (also Fig. 2, main text)	$n = 1$	$N = 9$
Two Dolphins (also Fig. 13, main text)	$n = 3$	$N = 9$
Single Dolphin	$n = 3$	$N = 9$
Fox	$n = 2$	$N = 8$
Airplane	$n = 2$	$N = 8$
Butterfly	$n = 2$	$N = 8$
Eagle	$n = 2$	$N = 8$
Spaceship (also Fig. 13, main text)	$n = 3$	$N = 8$
Hat	$n = 4$	$N = 9$
Lemon	$n = 3$	$N = 7$
Cat	$n = 2$	$N = 8$



# SinGAN: Learning a Generative Model from a Single Natural Image

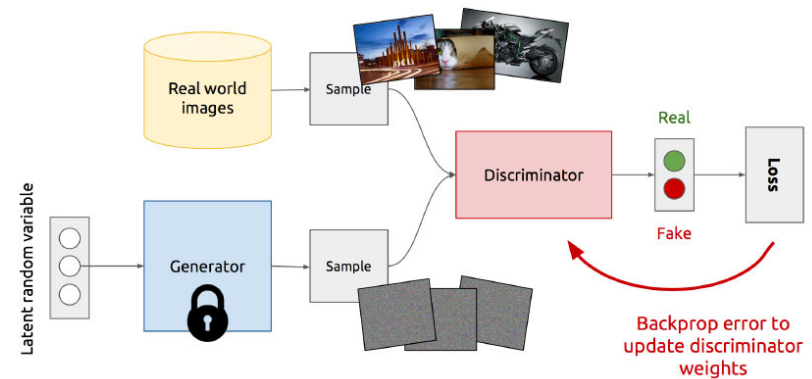


- Editing

Image	Injection scale	Total number of scales
Rock1	$n = 5$	$N = 7$
Rock2	$n = 5$	$N = 7$
Rock3 (also Fig. 12, main text)	$n = 5$	$N = 7$
Tree	$n = 7$	$N = 9$
Mountain	$n = 4$	$N = 8$
Red cliff	$n = 5$	$N = 9$
Hay	$n = 6$	$N = 9$



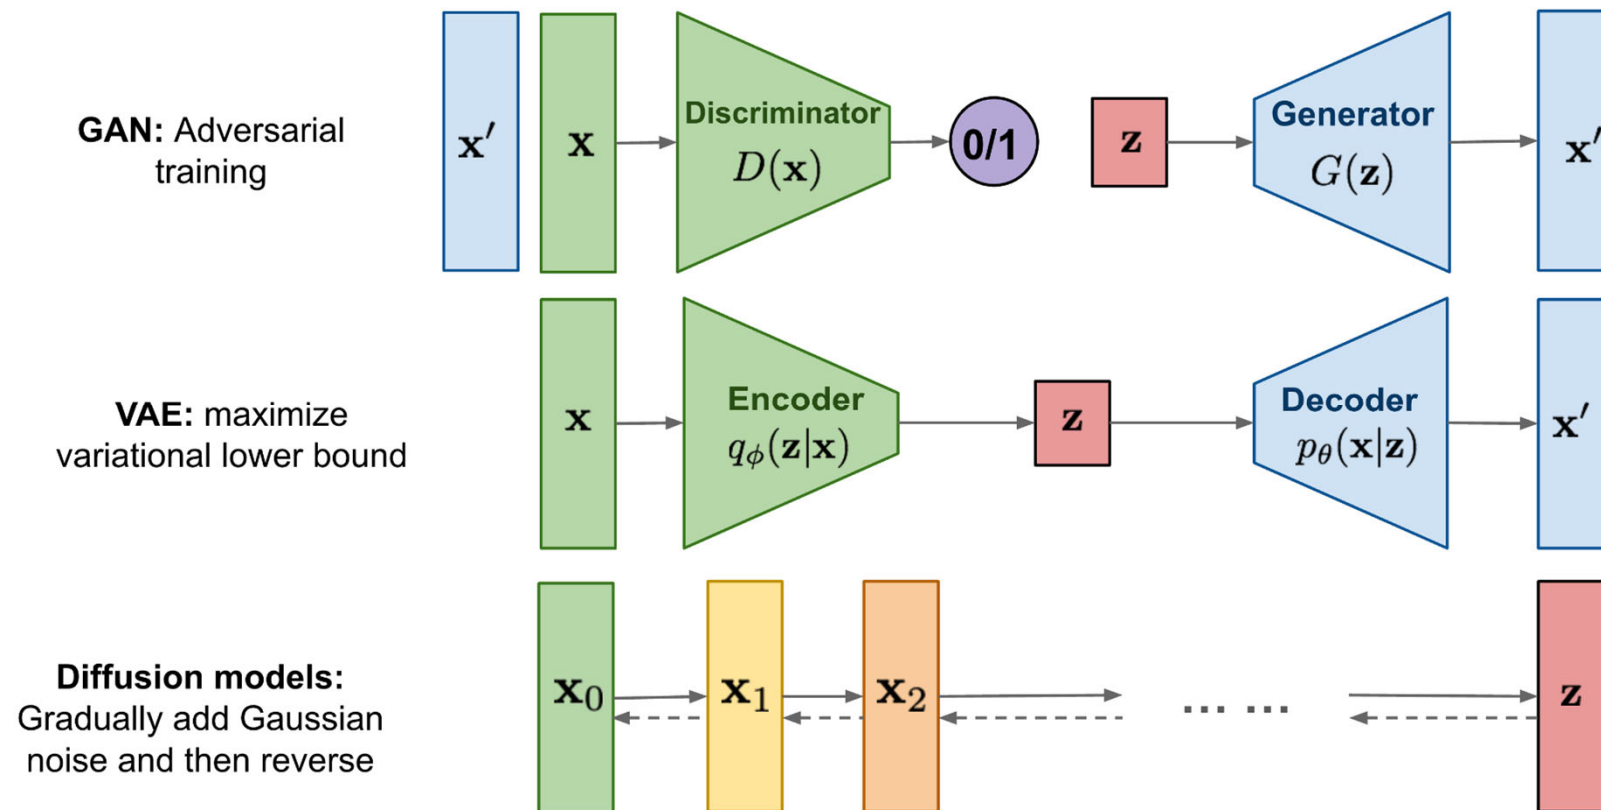
# What's to Be Covered Today...

- Generative Models
  - Auto-Encoder vs. Variational Auto-Encoder
  - Generative Adversarial Network (GAN)
  - Diffusion Model
- HW #1 is due Oct. 10<sup>th</sup> Mon 23:59
- HW #2 will be out next week...






# From VAE to Diffusion Model



# Denoising Diffusion Models

- Emerging as powerful generative models
  - Unconditional image synthesis
  - Conditional image synthesis
  - Outperforms GANs



Diffusion Models Beat GANs on Image Synthesis,  
Dhariwai & Nochoi, OpenAI, 2021



Cascaded Diffusion Models for High Fidelity Image  
Generation, Ho et al., Google, 2021



# Denoising Diffusion Models

- Emerging as powerful generative models
  - Unconditional image synthesis
  - Conditional image synthesis
  - Outperforms GANs

## DALL·E 2

"a teddy bear on a skateboard in times square"



Diffusion Models Beat GANs on Image Synthesis, Dhariwai & Nochoi, OpenAI, 2021

## Imagen

A group of teddy bears in suit in a corporate office celebrating the birthday of their friend. There is a pizza cake on the desk.

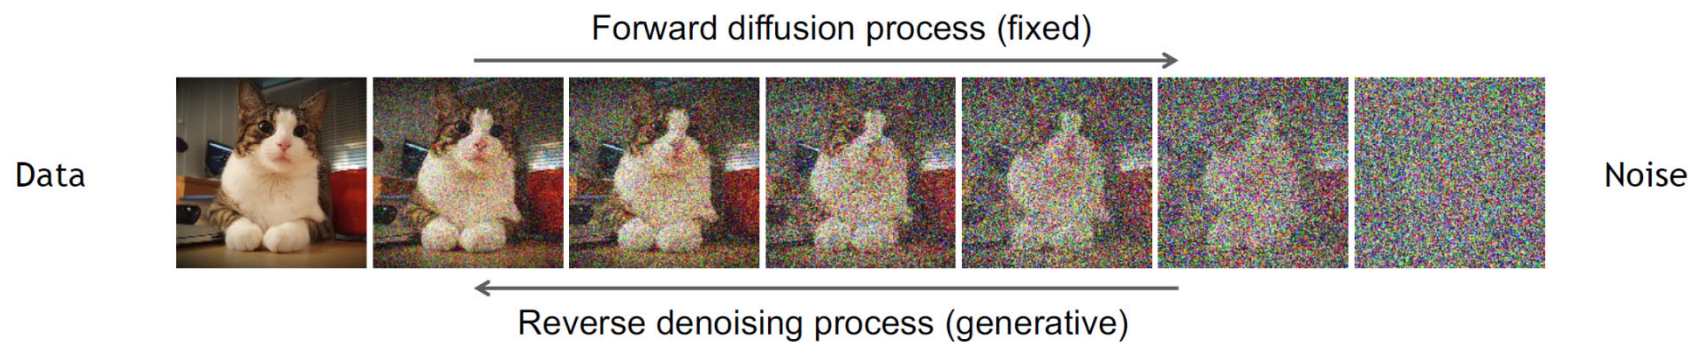



Cascaded Diffusion Models for High Fidelity Image Generation, Ho et al., Google, 2021

# Denoising Diffusion Models:

Learning to generate by denoising

- 2 processes required for training:
  - Forward diffusion process – gradually add noise to input
  - Reverse diffusion process – learns to generate/restore data by denoising (typically implemented via a U-net)
  - Comments about noise scheduling (see next slide)



[Ho et al., Denoising Diffusion Probabilistic Models, NeurIPS 2020](#)

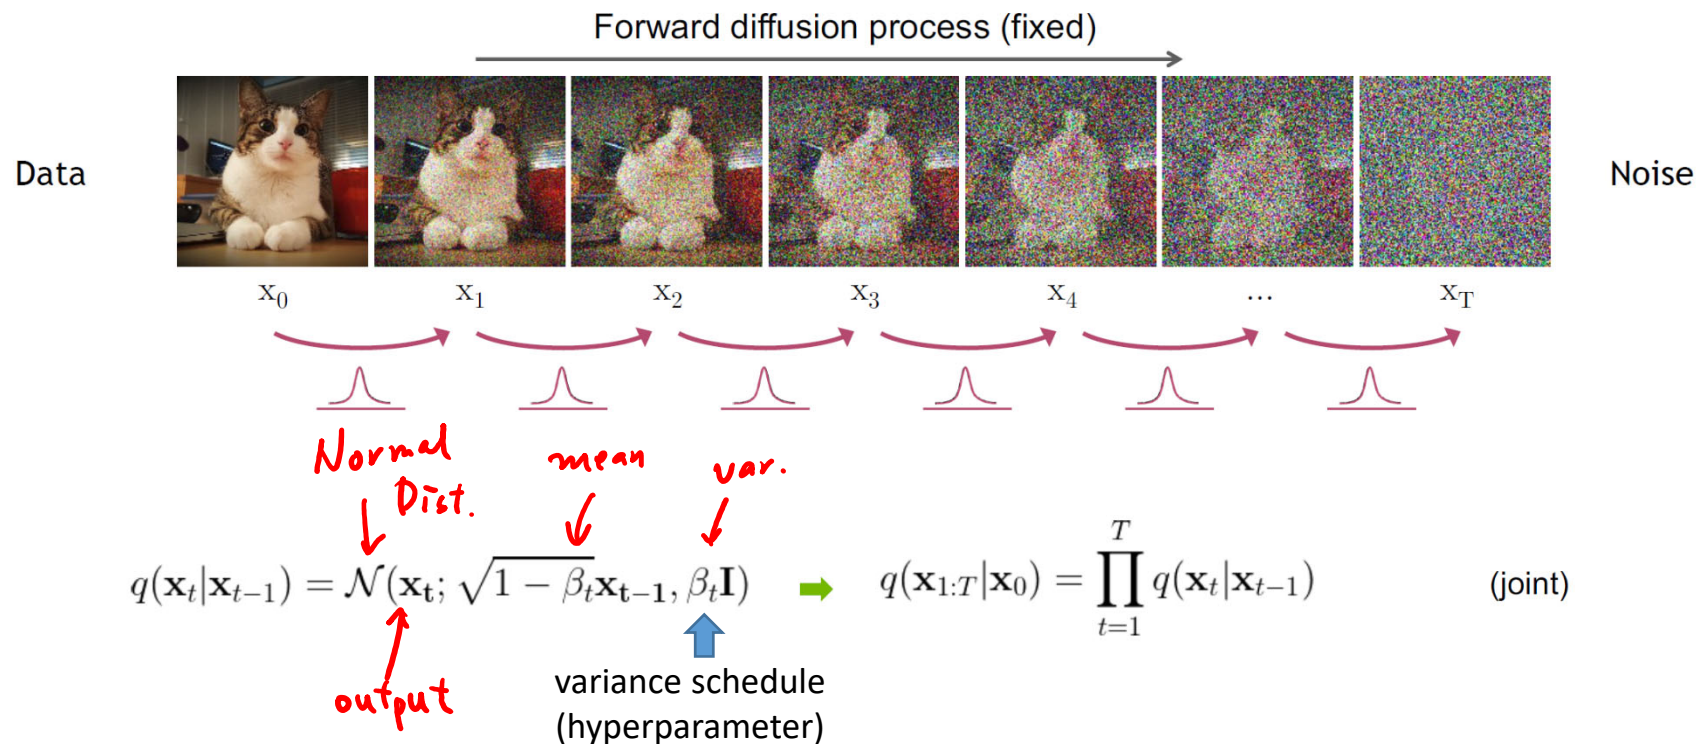

[Song et al., Score-Based Generative Modeling through Stochastic Differential Equations, ICLR 2021](#)

# Denoising Diffusion Models:

Learning to generate by denoising (cont'd)

- Forward diffusion process


- Gradually add noise to the input in  $T$  steps
- Recall that  $x_0$  denotes clean input image, and  $x_T$  is the final noisy one.
- Comments on  $q(x_t|x_{t-1})$



# Denoising Diffusion Models:

Learning to generate by denoising (cont'd)

- Forward diffusion process
  - Gradually add noise to the input in T steps (cont'd)
  - Diffusion kernel
  - So what happens to data distribution during this process?



$$q(\mathbf{x}_t) = \underbrace{\int}_{\text{Diffused data dist.}} \underbrace{q(\mathbf{x}_0, \mathbf{x}_t)}_{\text{Joint dist.}} d\mathbf{x}_0 = \int \underbrace{q(\mathbf{x}_0)}_{\text{Input data dist.}} \underbrace{q(\mathbf{x}_t | \mathbf{x}_0)}_{\text{Diffusion kernel}} d\mathbf{x}_0$$

The diffusion kernel is Gaussian convolution.

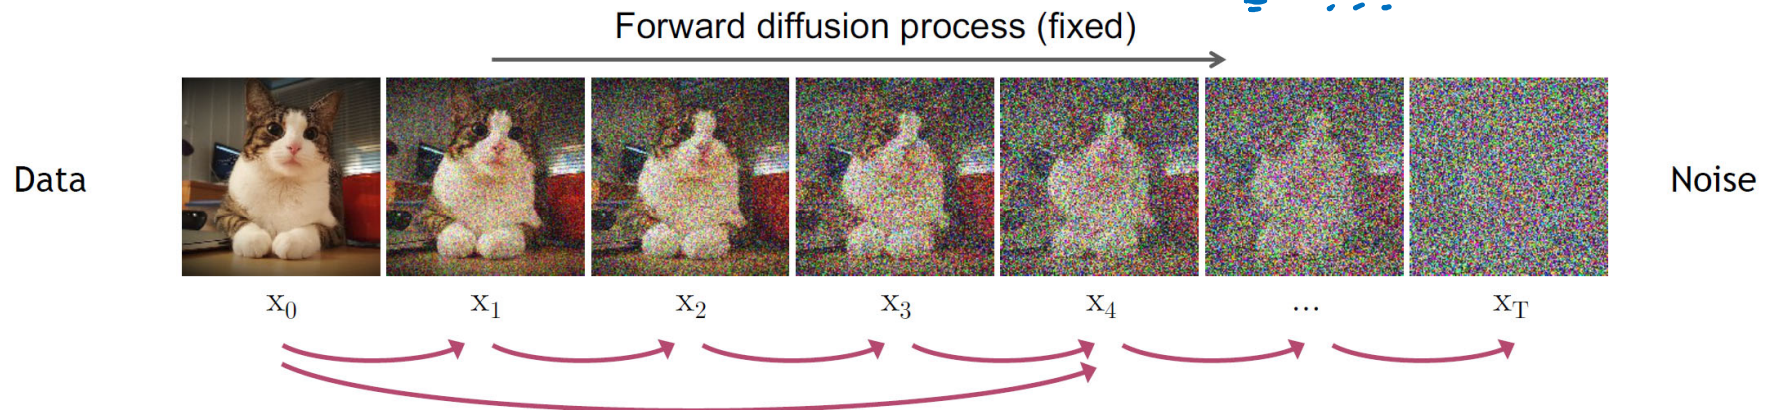


# Denoising Diffusion Models:

Learning to generate by denoising (cont'd)

- Forward diffusion process
  - Gradually add noise to the input in T steps
  - Diffusion kernel:

$$\begin{aligned}
 q(\mathbf{x}_t | \mathbf{x}_{t-1}) &= \mathcal{N}(\mathbf{x}_t, \sqrt{1-\beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}) \\
 &= \sqrt{1-\beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \boldsymbol{\epsilon} \\
 &= \sqrt{\alpha_t} \mathbf{x}_{t-1} + \sqrt{1-\alpha_t} \boldsymbol{\epsilon} \\
 &= \sqrt{\alpha_t \alpha_{t-1}} \mathbf{x}_{t-2} \\
 &\quad + \sqrt{1-\alpha_t \alpha_{t-1}} \boldsymbol{\epsilon} \\
 &= \dots
 \end{aligned}$$



Define  $\alpha_t = \prod_{s=1}^t (1 - \beta_s)$   $\alpha_t = 1 - \beta_t$   $\prod_{s=1}^t \alpha_s$  →  $q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) \mathbf{I})$  (Diffusion Kernel)


For sampling:  $\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{(1 - \alpha_t)} \boldsymbol{\epsilon}$  where  $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

$\beta_t$  values schedule (i.e., the noise schedule) is designed such that  $\bar{\alpha}_T \rightarrow 0$  and  $q(\mathbf{x}_T | \mathbf{x}_0) \approx \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$

# Denoising Diffusion Models:

Learning to generate by denoising (cont'd)

- Generative learning by denoising
  - Diffusion parameters are designed such that:  $q(\mathbf{x}_T) \approx \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$



**Generation:**

Sample  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$

Iteratively sample  $\mathbf{x}_{t-1} \sim \underbrace{q(\mathbf{x}_{t-1}|\mathbf{x}_t)}$

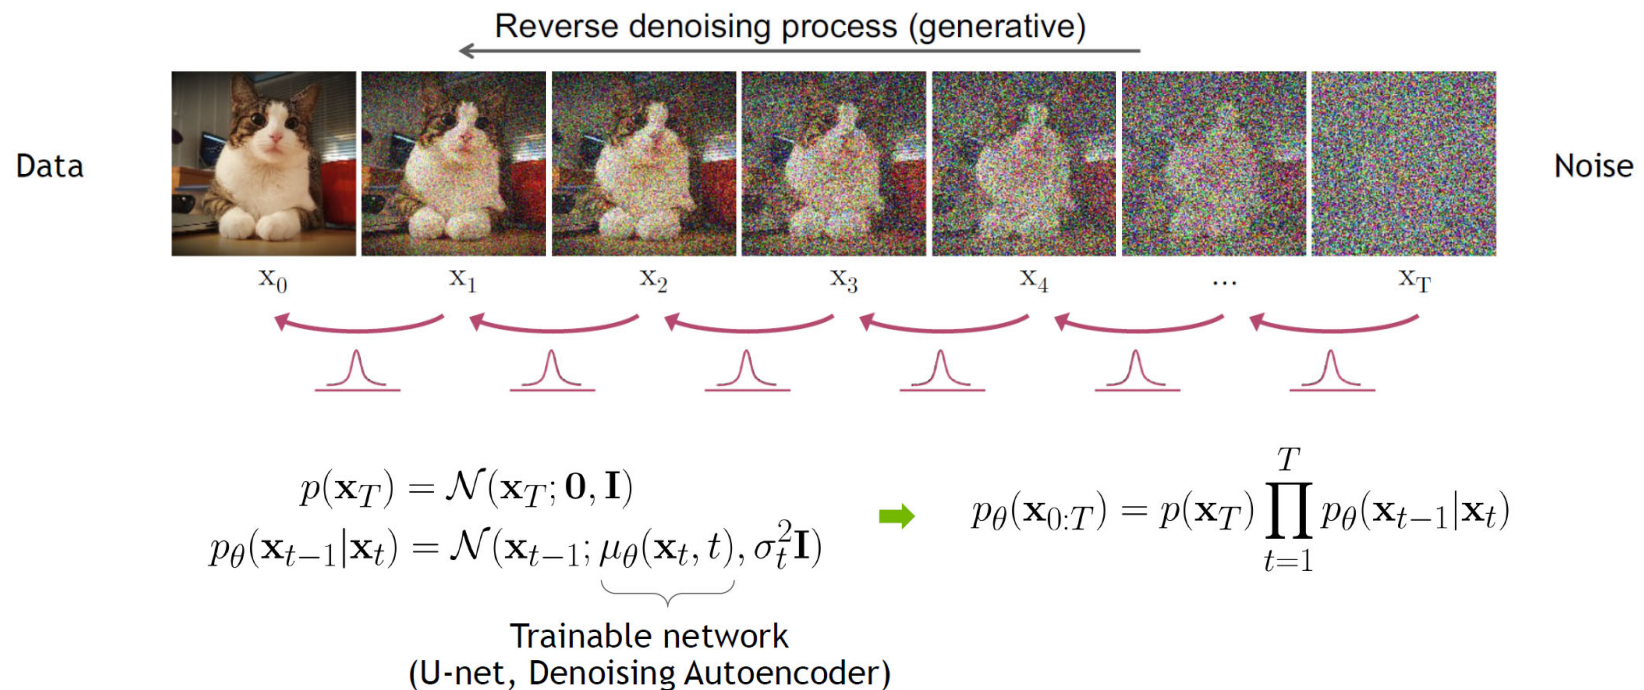
True Denoising Dist.

- Unfortunately,  $q(\mathbf{x}_{t-1}|\mathbf{x}_t) \propto q(\mathbf{x}_{t-1})q(\mathbf{x}_t|\mathbf{x}_{t-1})$  is intractable.  
We approximate  $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$  by Normal distribution by setting small  $\beta_t$  in each step

# Denoising Diffusion Models:

Learning to generate by denoising (cont'd)

- Reverse diffusion process
  - Learn to denoise in T steps
  - Let the model  $\theta$  predict  $p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_{\theta}(\mathbf{x}_t, t), \boldsymbol{\Sigma}_{\theta}(\mathbf{x}_t, t))$
  - To conclude the learning process first, we need to **predict the noise in image**.






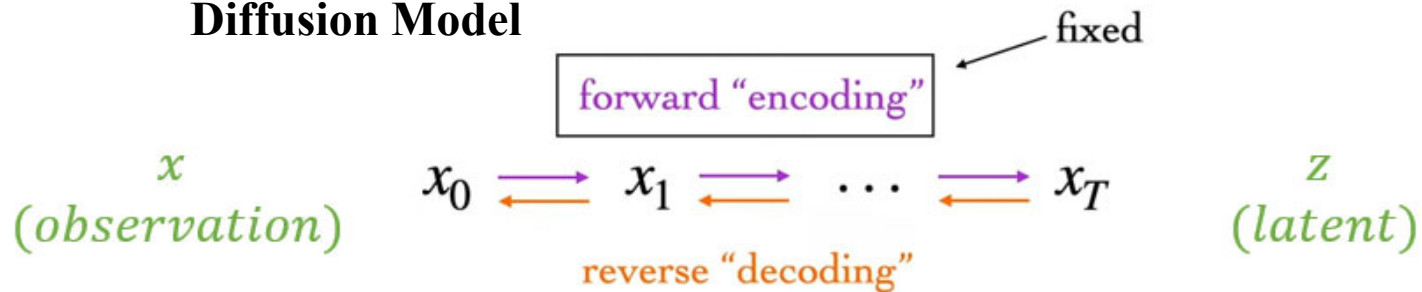
# Learning of Diffusion Models

- $\log p_\theta(x) \geq$  variational lower bound

VAE



Diffusion Model



$$\log p_\theta(x) \geq \text{variational lower bound}$$

# Learning of Diffusion Models

- $\log p_\theta(x) \geq$  variational lower bound

- Recall that we exploit variational bound for optimizing VAE models

$$\log p_\theta(x) \geq E_{z \sim q_\phi(z|x)} [\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x), p(z))$$

$$\text{vs. } \mathbb{E}_{q(\mathbf{x}_0)} [-\log p_\theta(\mathbf{x}_0)] \leq \mathbb{E}_{q(\mathbf{x}_0)q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \left[ -\log \frac{p_\theta(\mathbf{x}_{0:T})}{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \right] =: L$$

- In Ho et al. NeurIPS'20, it is shown that


$$L = \mathbb{E}_q \left[ \underbrace{D_{KL}(q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T))}_{L_T} + \sum_{t>1} \underbrace{D_{KL}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))}_{L_{t-1}} - \underbrace{\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)}_{L_0} \right]$$

*Can ignore. Why? ↗*


*Handwritten notes:*

- $N(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \beta_t \mathbf{I})$  (pointing to  $L_{t-1}$ )
- $N(\mathbf{x}_{t-1}; \tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I})$  (pointing to  $L_{t-1}$ )
- $\tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t} \mathbf{x}_t + \frac{\sqrt{\alpha_{t-1}}\beta_t}{1-\bar{\alpha}_t} \mathbf{x}_0$  (pointing to  $L_{t-1}$ )
- $= \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{\beta_t}{1-\bar{\alpha}_t} \epsilon)$  (pointing to  $L_{t-1}$ )
- $\tilde{\beta}_t$  is fixed (pointing to  $L_{t-1}$ )

VAE



Diffusion model



# Learning of Diffusion Models (cont'd)

- Recall that  $L = \mathbb{E}_q \left[ \underbrace{D_{\text{KL}}(q(\mathbf{x}_T|\mathbf{x}_0) \parallel p(\mathbf{x}_T))}_{L_T} + \sum_{t>1} \underbrace{D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \parallel p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))}_{L_{t-1}} - \underbrace{\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)}_{L_0} \right]$  {0, 1, ..., 255} → [-1,1]

- Still working on it...

- Only care about KL divergence between two Gaussian distributions

$$\begin{cases} q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0) \frac{q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)} := \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}) \\ p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \underbrace{\mu_\theta(\mathbf{x}_t, t)}_{\text{learned}}, \underbrace{\Sigma_\theta(\mathbf{x}_t, t)}_{\text{fixed}}) \end{cases}$$

$\frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon) \quad (1)$   
 (actual  $\mu$ )

$\frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(\mathbf{x}_t, t)) \quad (2)$   
 (predicted  $\mu$ )

- As a result,

$$\frac{1}{2\sigma_t^2} \| (1) - (2) \|^2 \rightarrow \mathbb{E}_{\mathbf{x}_0, \epsilon} \left[ \frac{\beta_t^2}{2\sigma_t^2 \alpha_t (1 - \bar{\alpha}_t)} \left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) \right\|^2 \right]$$

- For simplicity, we calculate

$$L_{\text{simple}}(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \epsilon} \left[ \left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) \right\|^2 \right]$$

# Learning of Diffusion Models

- Summary
  - Training and sample generation

## Algorithm 1 Training

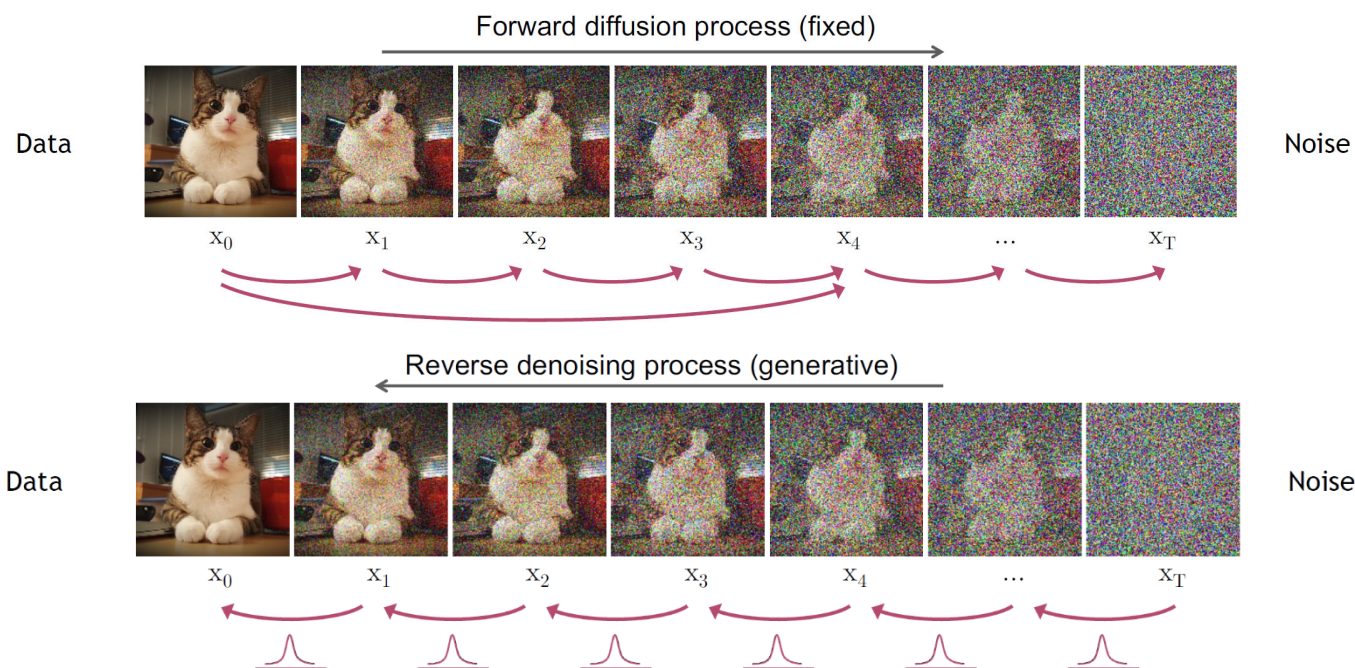
```

1: repeat
2:    $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ 
3:    $t \sim \text{Uniform}(\{1, \dots, T\})$ 
4:    $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
5:   Take gradient descent step on
        $\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t)\|^2$ 
6: until converged
    
```

## Algorithm 2 Sampling

```

1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
4:    $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_{\theta}(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$ 
5: end for
6: return  $\mathbf{x}_0$ 
    
```



# Learning of Diffusion Models

- Summary
  - Training and sample generation

## Algorithm 1 Training


```

1: repeat
2:    $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ 
3:    $t \sim \text{Uniform}(\{1, \dots, T\})$ 
4:    $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
5:   Take gradient descent step on
        $\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|^2$ 
6: until converged
    
```

## Algorithm 2 Sampling

```

1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
4:    $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_{\theta}(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$ 
5: end for
6: return  $\mathbf{x}_0$ 
    
```



*More steps* →

Slide credit: Kreis, Gao, & Vahdat

<https://medium.com/ai-blog-tw/%E9%82%8A%E5%AF%A6%E4%BD%9C%E9%82%8A%E5%AD%B8%E7%BF%92diffusion-model-%E5%BE%9Eddpm%E7%9A%84%E7%B0%A1%E5%8C%96%E6%A6%82%E5%BF%B5%E7%90%86%E8%A7%A3-4c565a1c09c>



# What We've Covered Today...

- Generative Models
  - Auto-Encoder vs. Variational Auto-Encoder
  - Generative Adversarial Network (GAN)
  - Diffusion Model
- HW #1 is due Oct. 10<sup>th</sup> Mon 23:59
- HW #2 will be out next week...

