
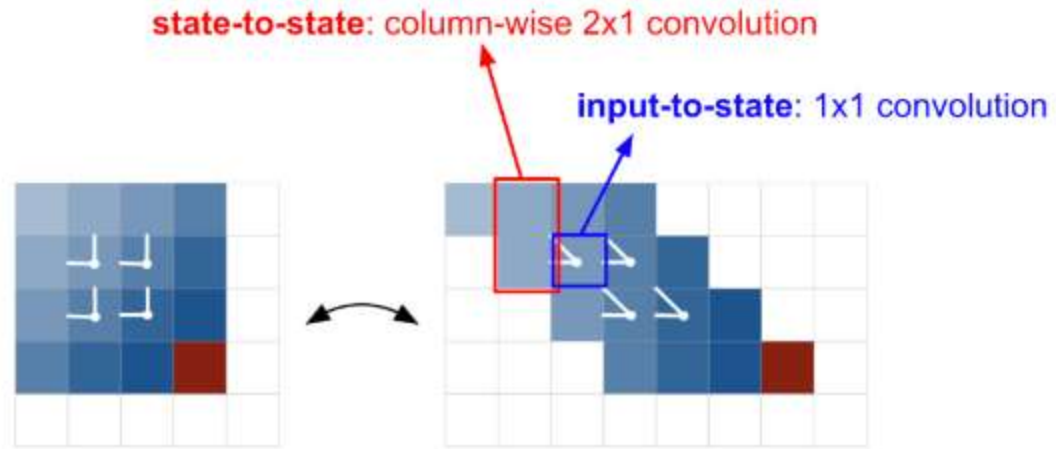


Recap.

PixelRNN [van der Oord et al. 2016]

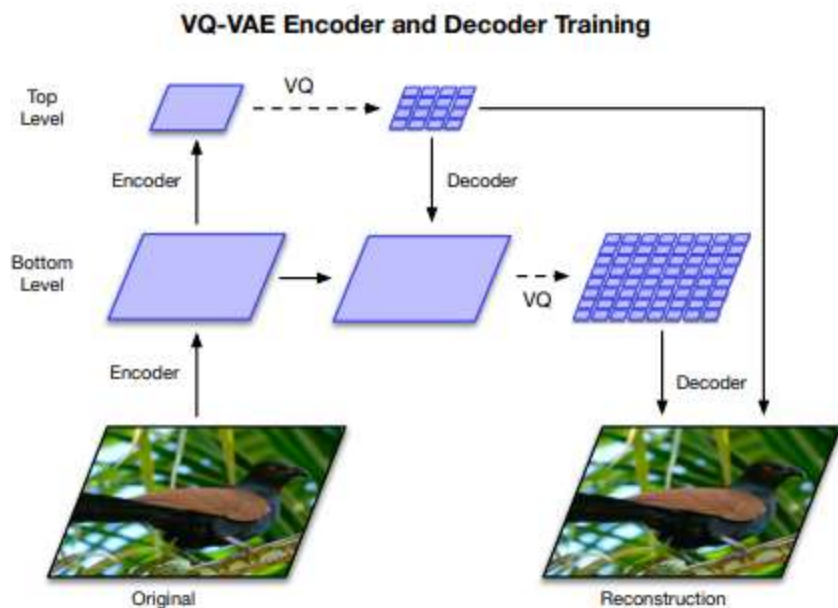


(a) Diagonal BiLSTM




(b) Skewing operation

VQ-VAE and VQ-VAE2



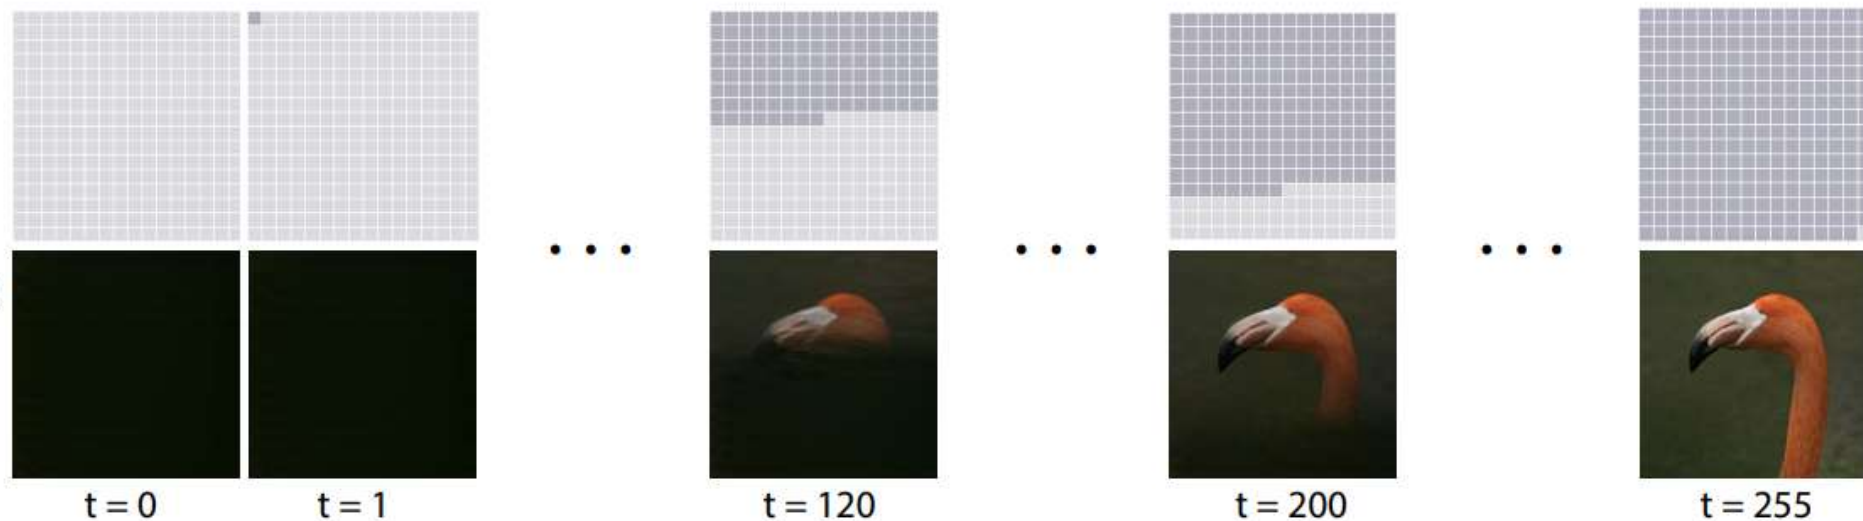
(a) Overview of the architecture of our hierarchical VQ-VAE. The encoders and decoders consist of deep neural networks. The input to the model is a 256×256 image that is compressed to quantized latent maps of size 64×64 and 32×32 for the *bottom* and *top* levels, respectively. The decoder reconstructs the image from the two latent maps.



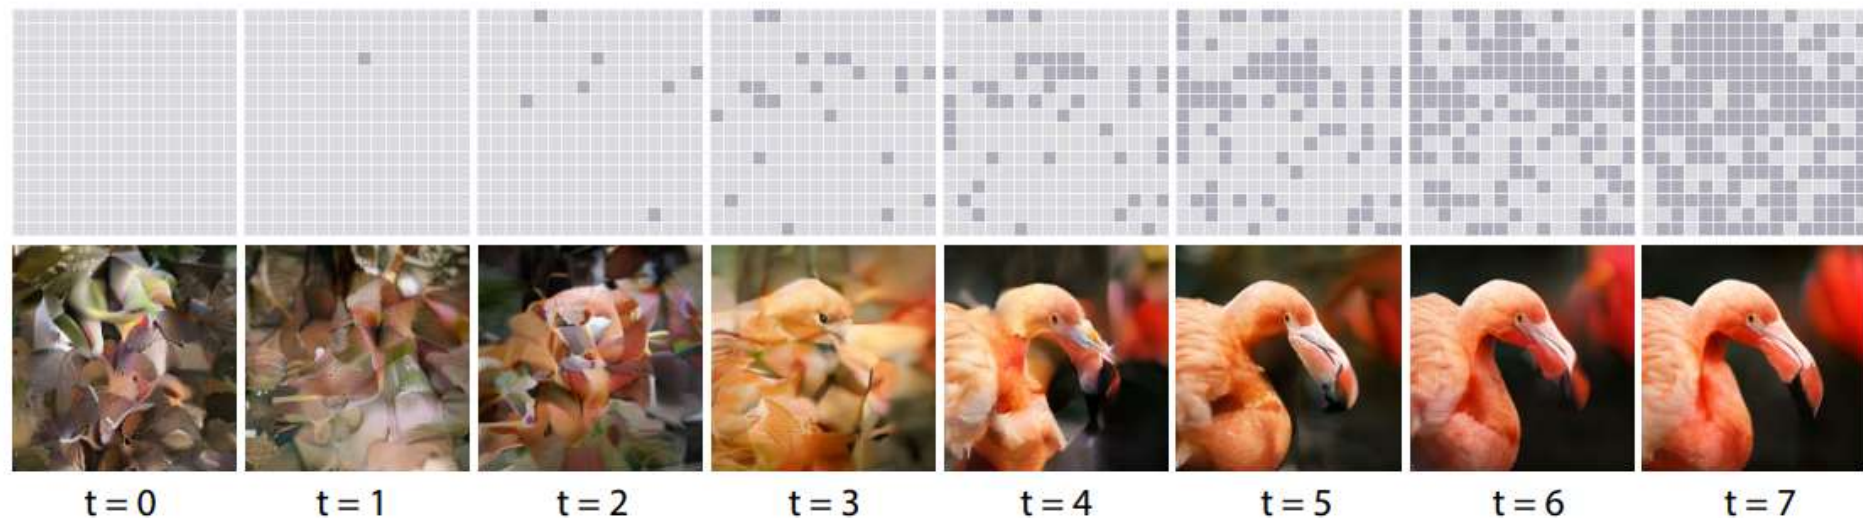
(b) Multi-stage image generation. The top-level PixelCNN prior is conditioned on the class label, the bottom level PixelCNN is conditioned on the class label as well as the first level code. Thanks to the feed-forward decoder, the mapping between latents to pixels is fast. (The example image with a parrot is generated with this model).

MaskGIT: Masked Generative Image Transformer

Sequential
Decoding
with Autoregressive
Transformers

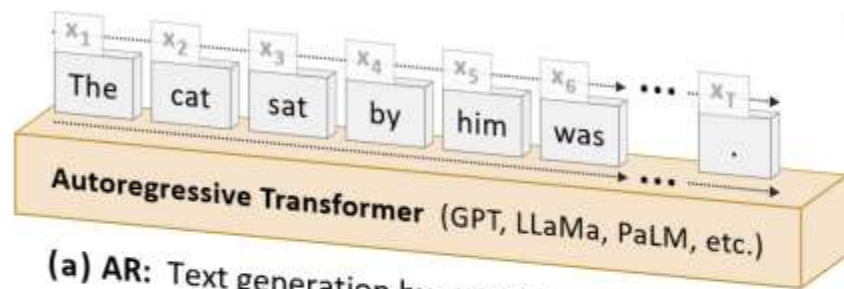


Scheduled
Parallel
Decoding
with MaskGIT

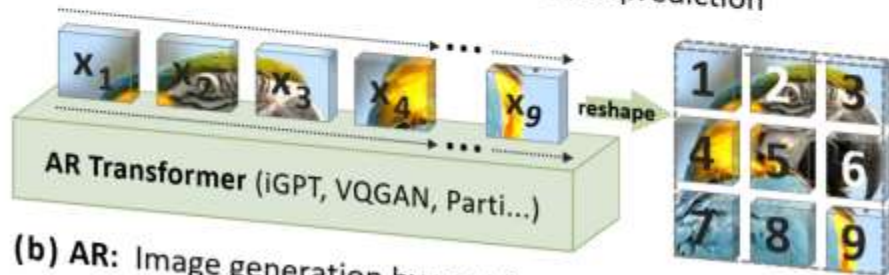


VAR: Visual Autoregressive Modeling

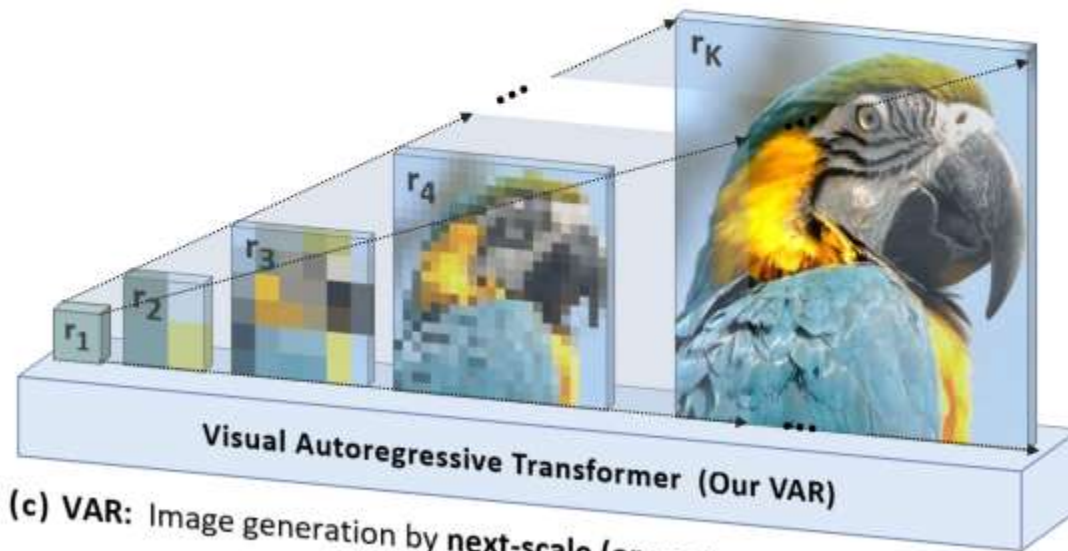
Three Different Autoregressive Generative Models



(a) AR: Text generation by next-token prediction




(b) AR: Image generation by next-image-token prediction



(c) VAR: Image generation by next-scale (or next-resolution) prediction

Generative Adversarial Networks



Overview

- Generative Adversarial Networks (GAN)
- Adversary as a Loss Function
- Wasserstein GAN (W-GAN)

Generative Adversarial Networks (GAN)

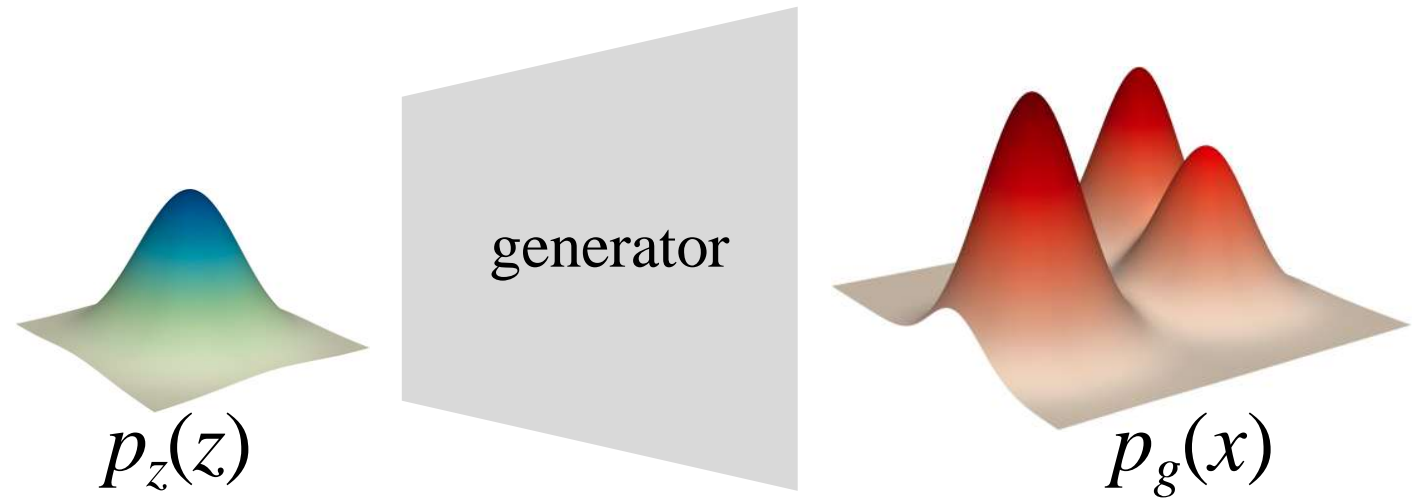
Introduction

- **“Generative”**
 - “Discriminative” was dominant back then
- **“Adversarial”**
 - Generative models w/ discriminative models
 - Min-max process
- **“Networks”**
 - SGD + backprop for problem solving

Recap: Latent Variable Models

Represent a distribution by a neural network:

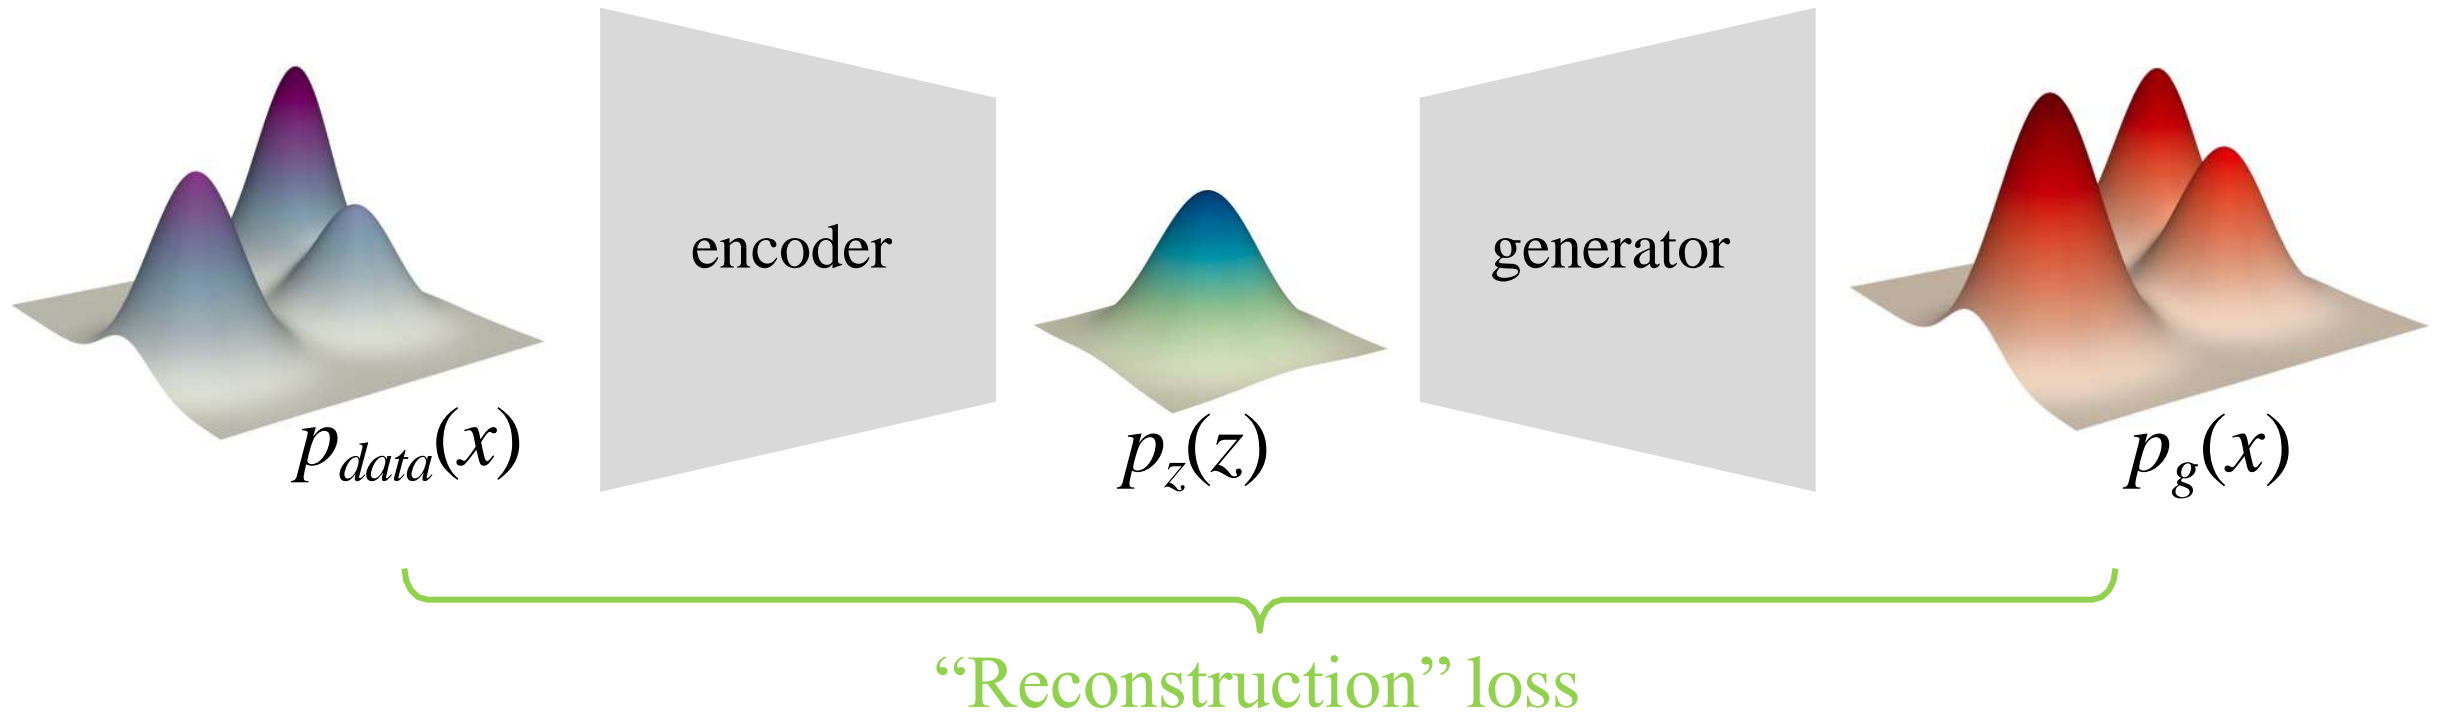
- z - latent variables
- x - observed variables



Recap: Variational Autoencoder (VAE)

Autoencoding distributions:

“Encoding” data distribution p_{data} into latent distribution p_z



What's the implication of a “*reconstruction*” loss?

- Elements (e.g., pixels) are **independently** distributed
- Each element follows a **simple** distribution (Gaussian/Bernoulli/...)

Assumptions are too strict for **high-dim** variables

Can we measure the distribution difference in another way?

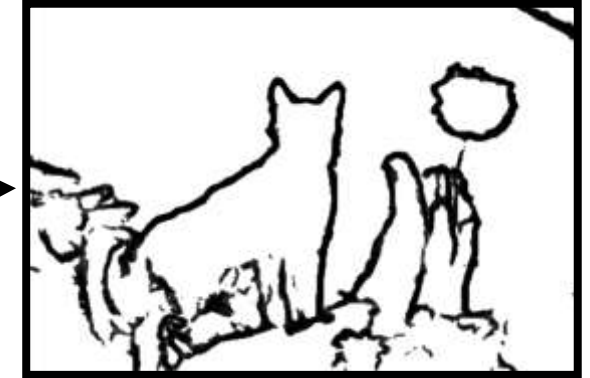
Data translation problems (“structured prediction”)

Semantic segmentation



[Long et al. 2015, ...]

Edge detection



[Xie et al. 2015, ...]

Text-to-photo

“this small bird has a pink
breast and crown...”



[Reed et al. 2014, ...]

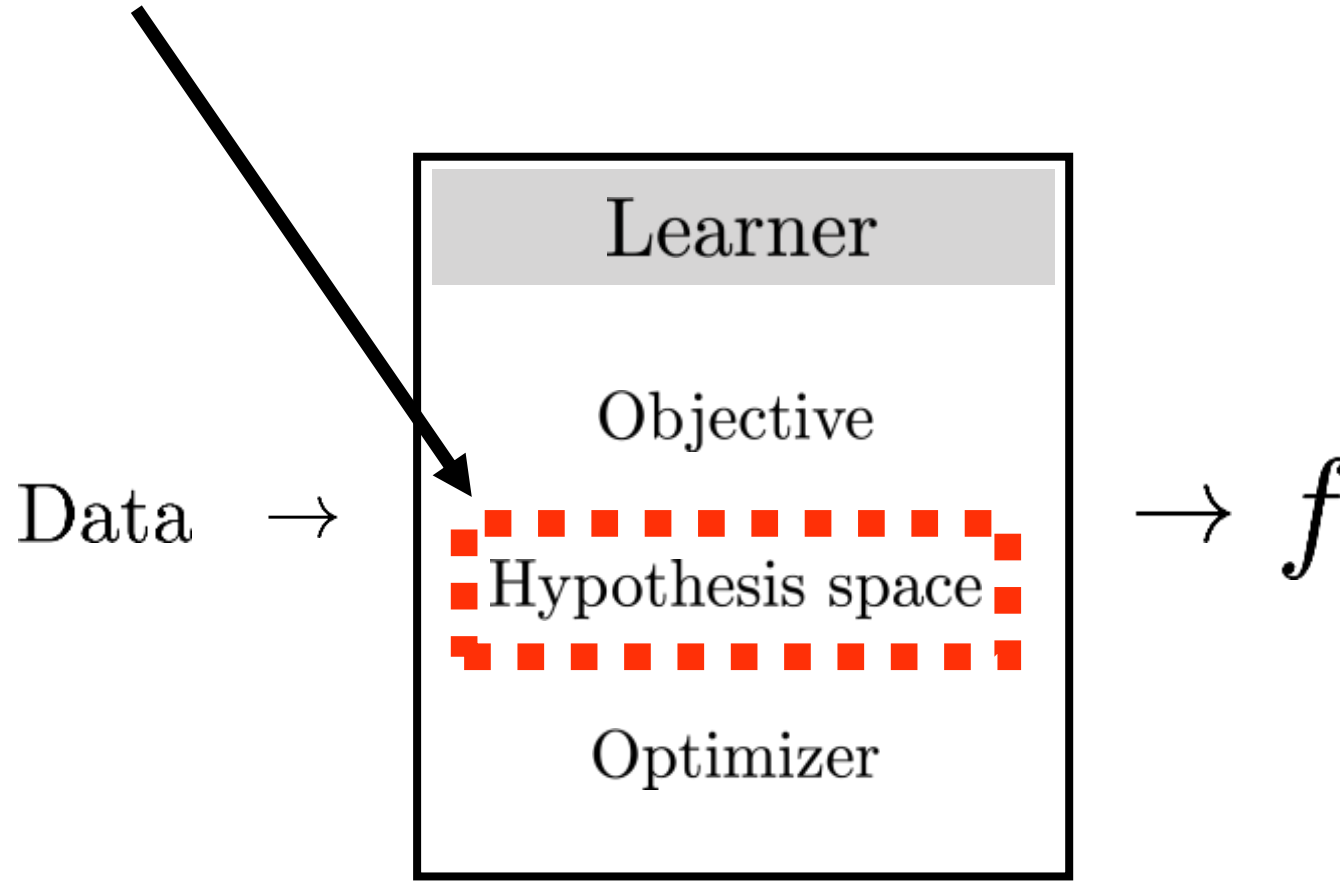
Future frame prediction



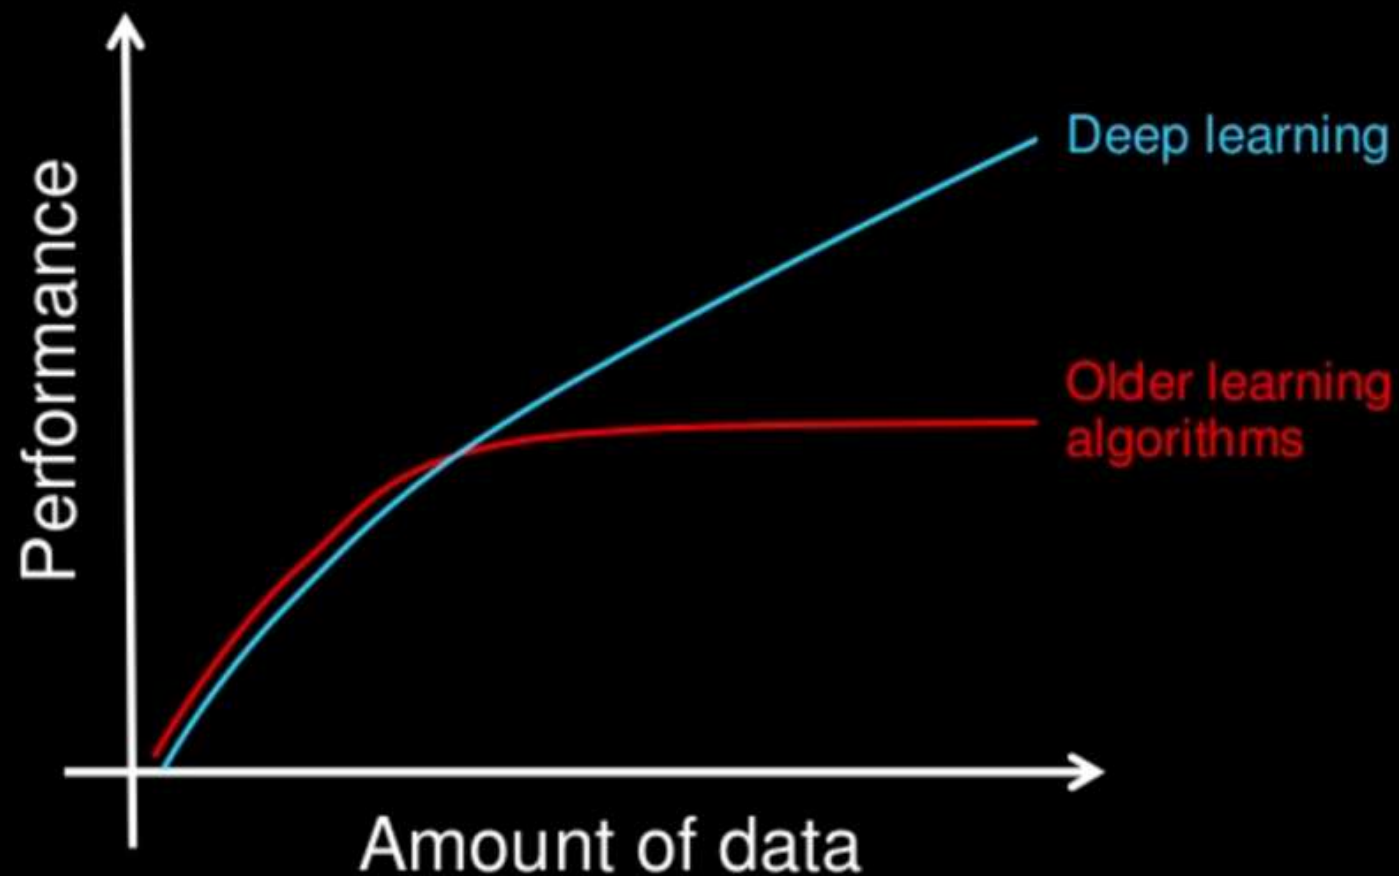
[Mathieu et al. 2016, ...]

Deep learning in 2012

Use a **hypothesis space** that can model complex structure
(e.g., a CNN, nearest-neighbor)

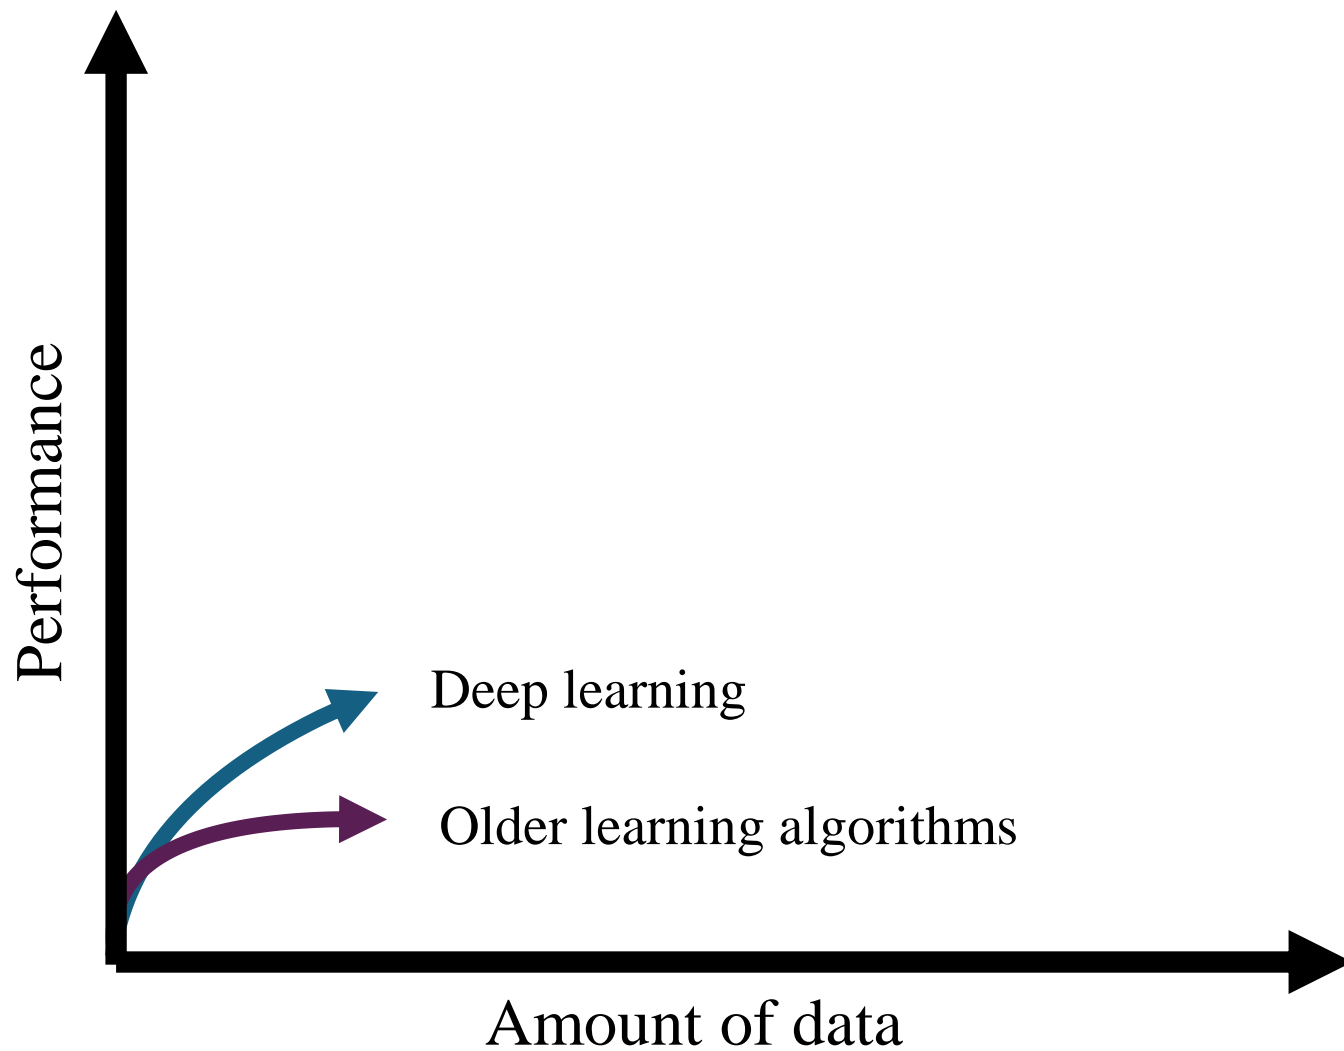


Why deep learning

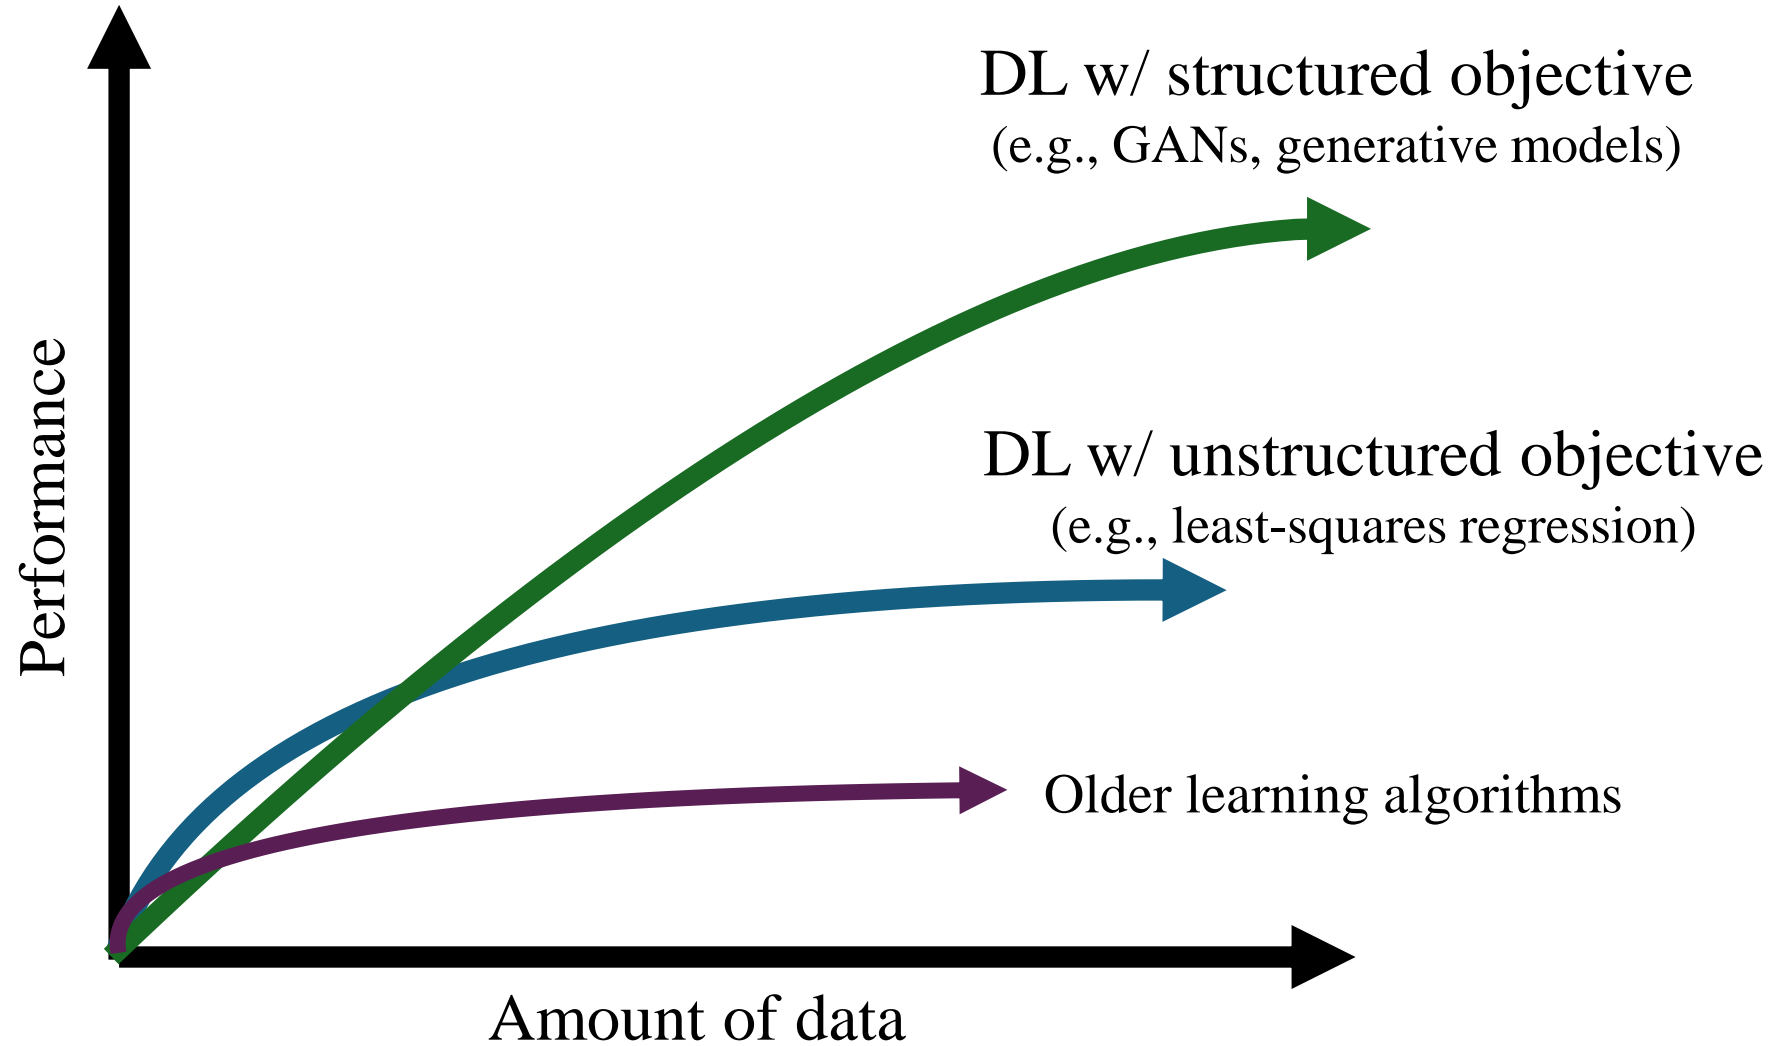


How do data science techniques scale with amount of data?

Why structured objectives





Why structured objectives





[Photo credit: Fredo Durand]



(Colors represent one-hot codes)

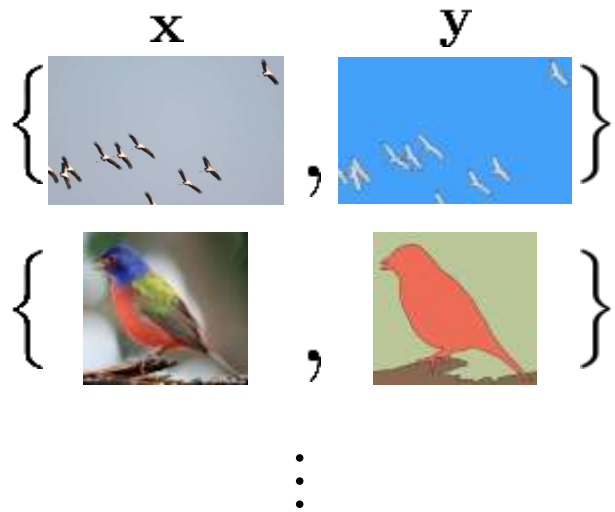
$$\arg \min_{\mathcal{F}} \mathbb{E}_{\mathbf{x}, \mathbf{y}} [L(\mathcal{F}(\mathbf{x}), \mathbf{y})]$$

Hypothesis space

Objective function
(loss)

Semantic Segmentation

Data



$$\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$$

$$\mathbf{y} \in \mathbb{R}^{H \times W \times K}$$



Learner

Objective

$$f^* = \arg \min_{f \in \mathcal{F}} \sum_{i=1}^N H(\mathbf{y}_i, \hat{\mathbf{y}}_i)$$

Hypothesis space

Convolutional neural net

Optimizer


Stochastic gradient descent



f


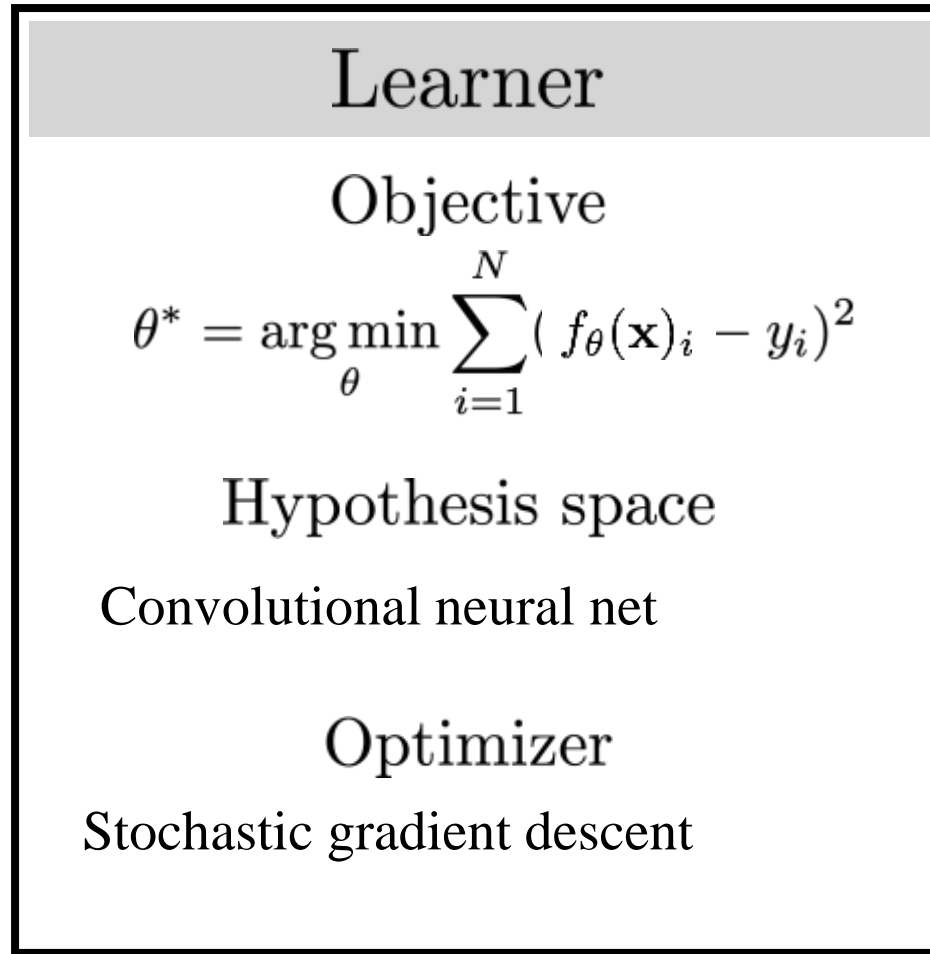
Sat2Map

Data



$$\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$$

$$\mathbf{y} \in \mathbb{R}^{H \times W \times 3}$$



Input

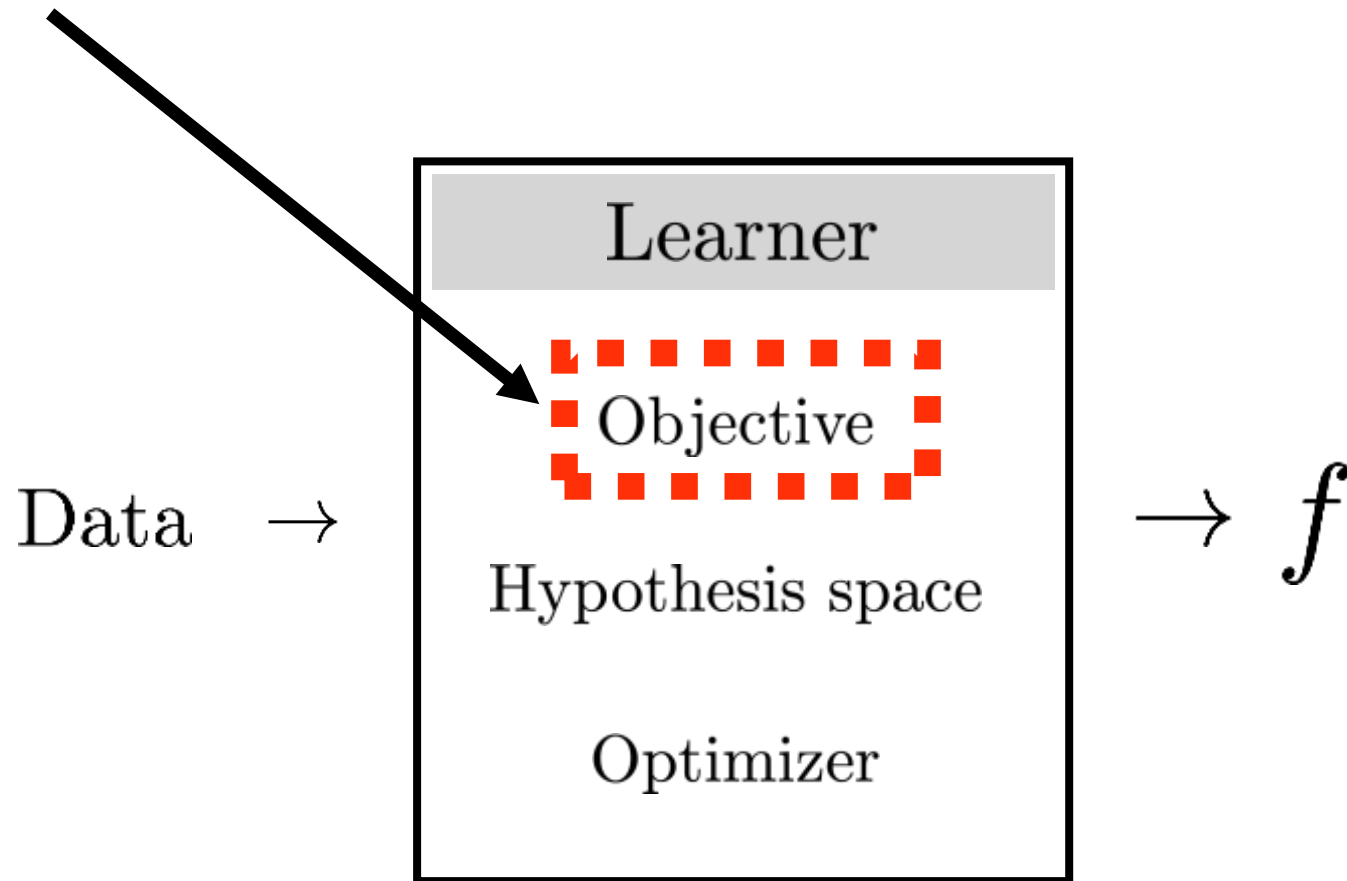


Deep net output



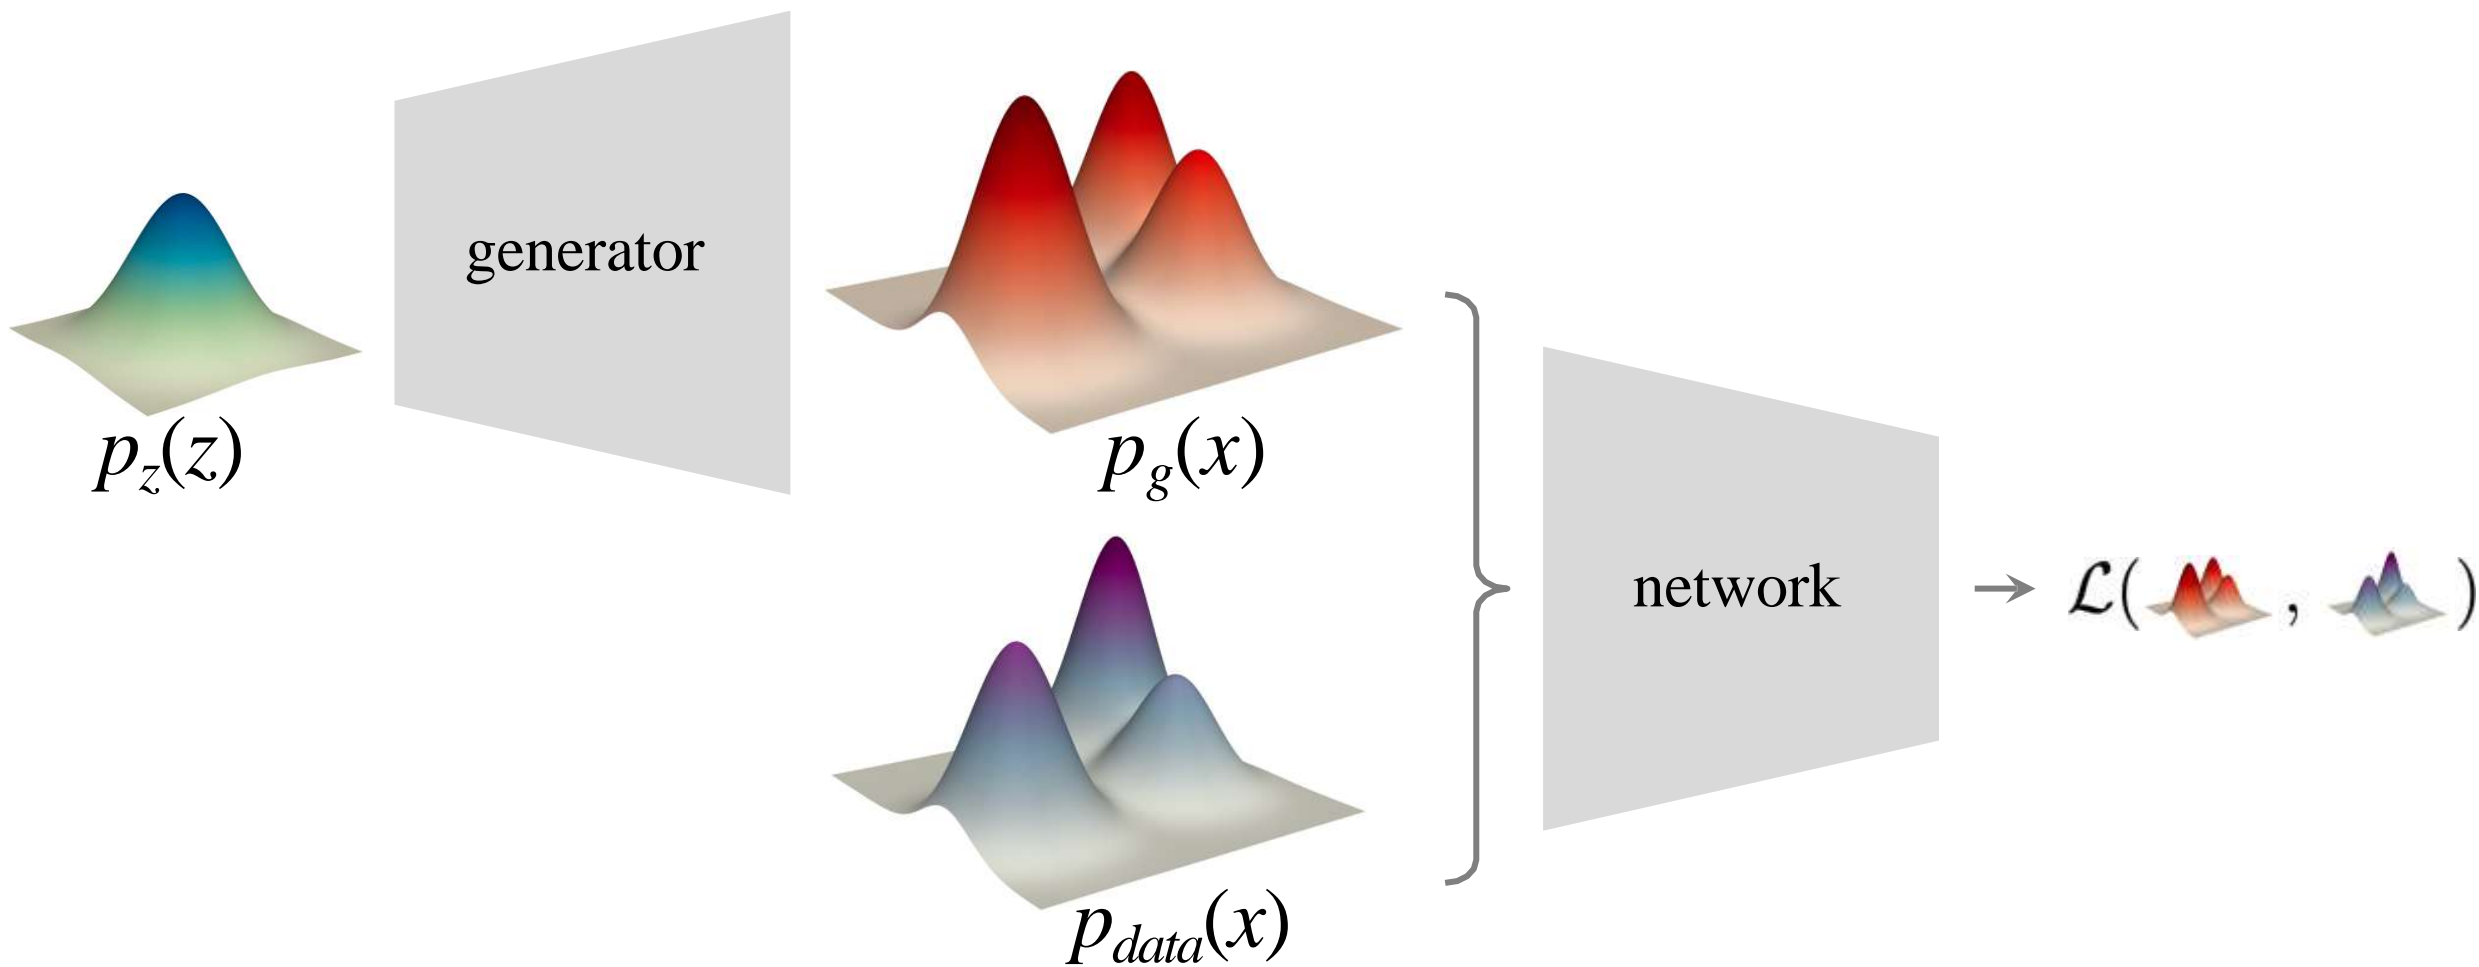
Structured prediction

Use an **objective** that can model structure! (e.g., a graphical model, a GAN, etc)



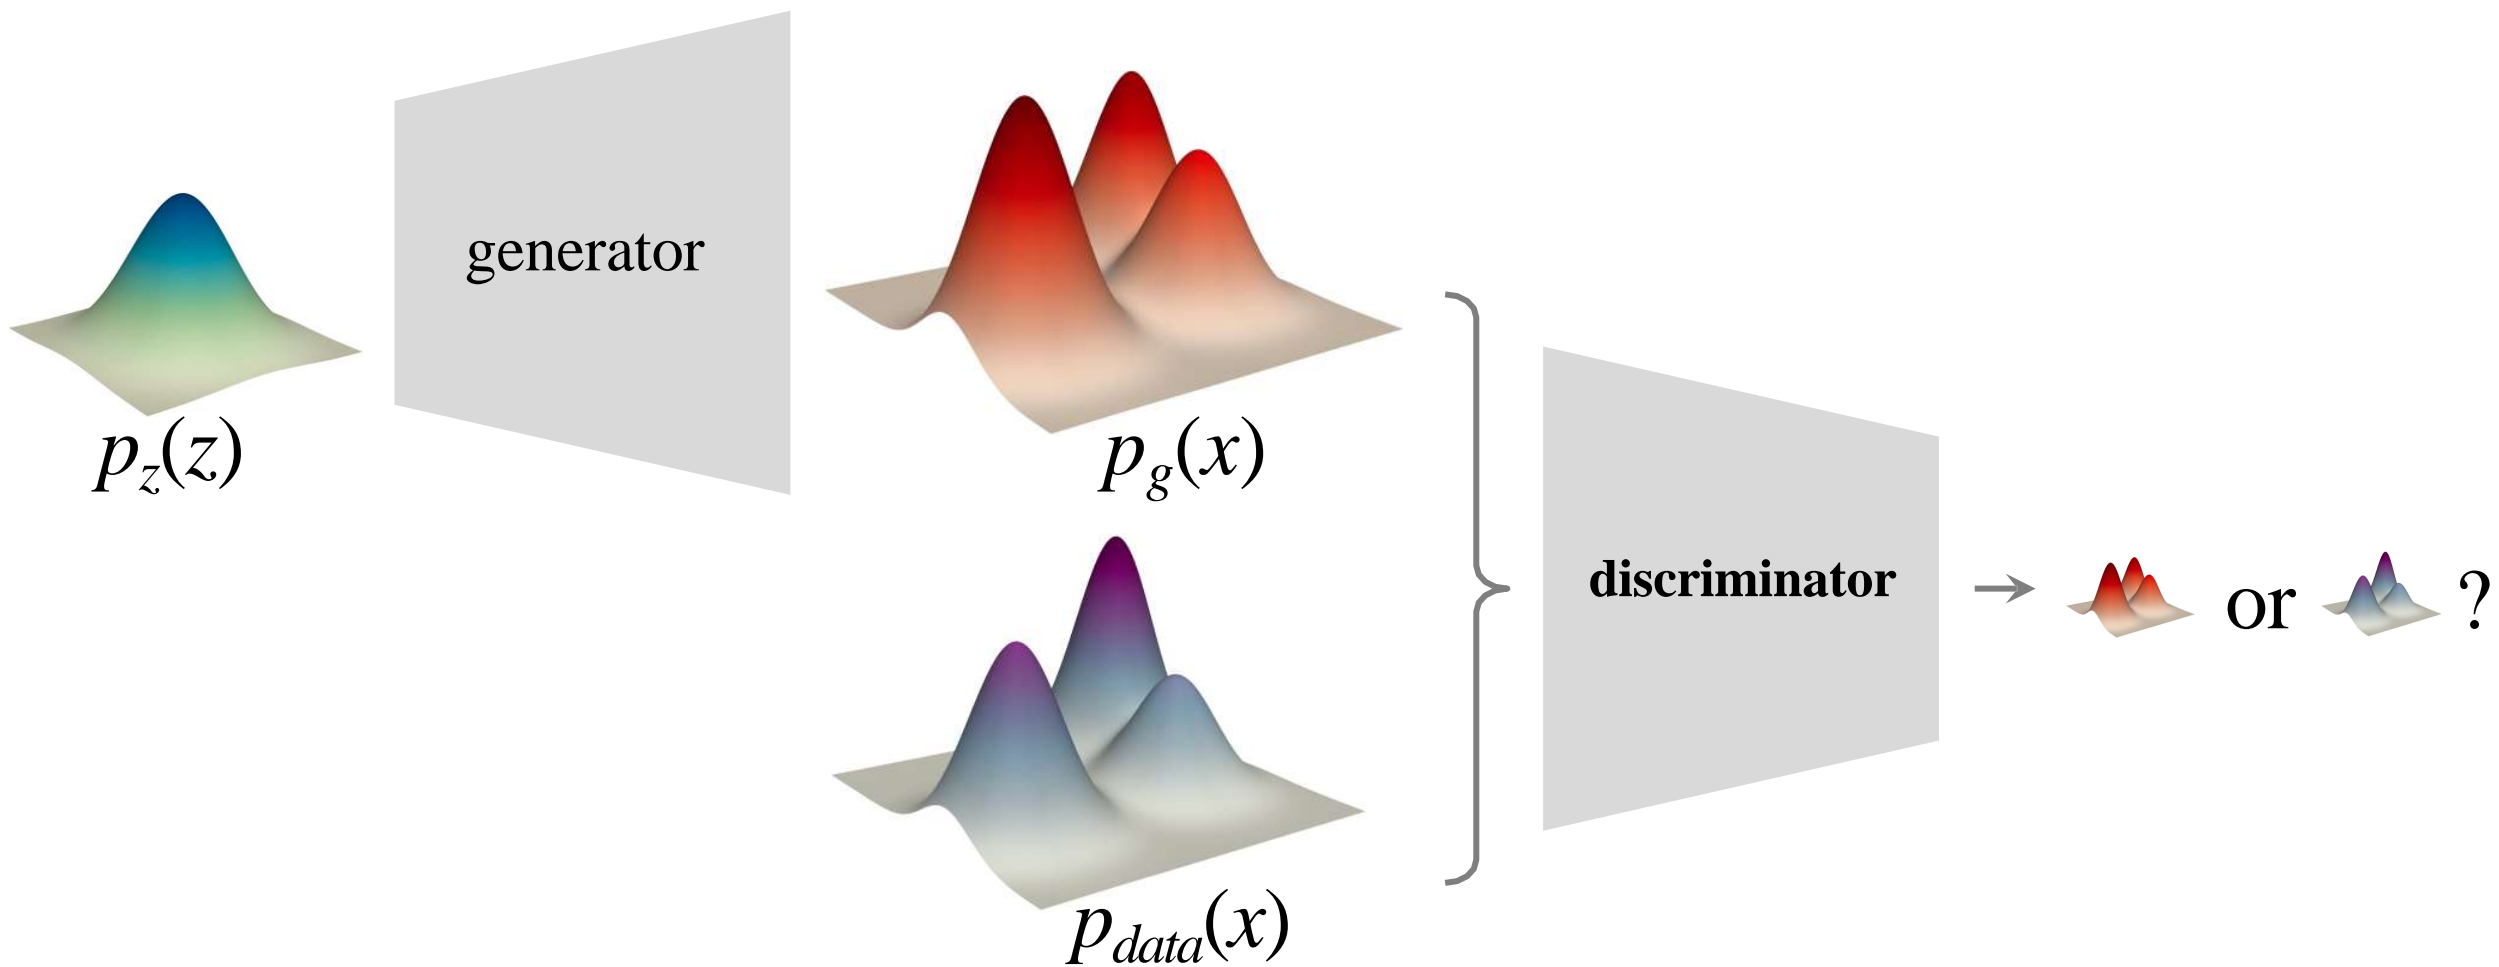
Generative Adversarial Networks

Representing **distribution difference** by a neural network



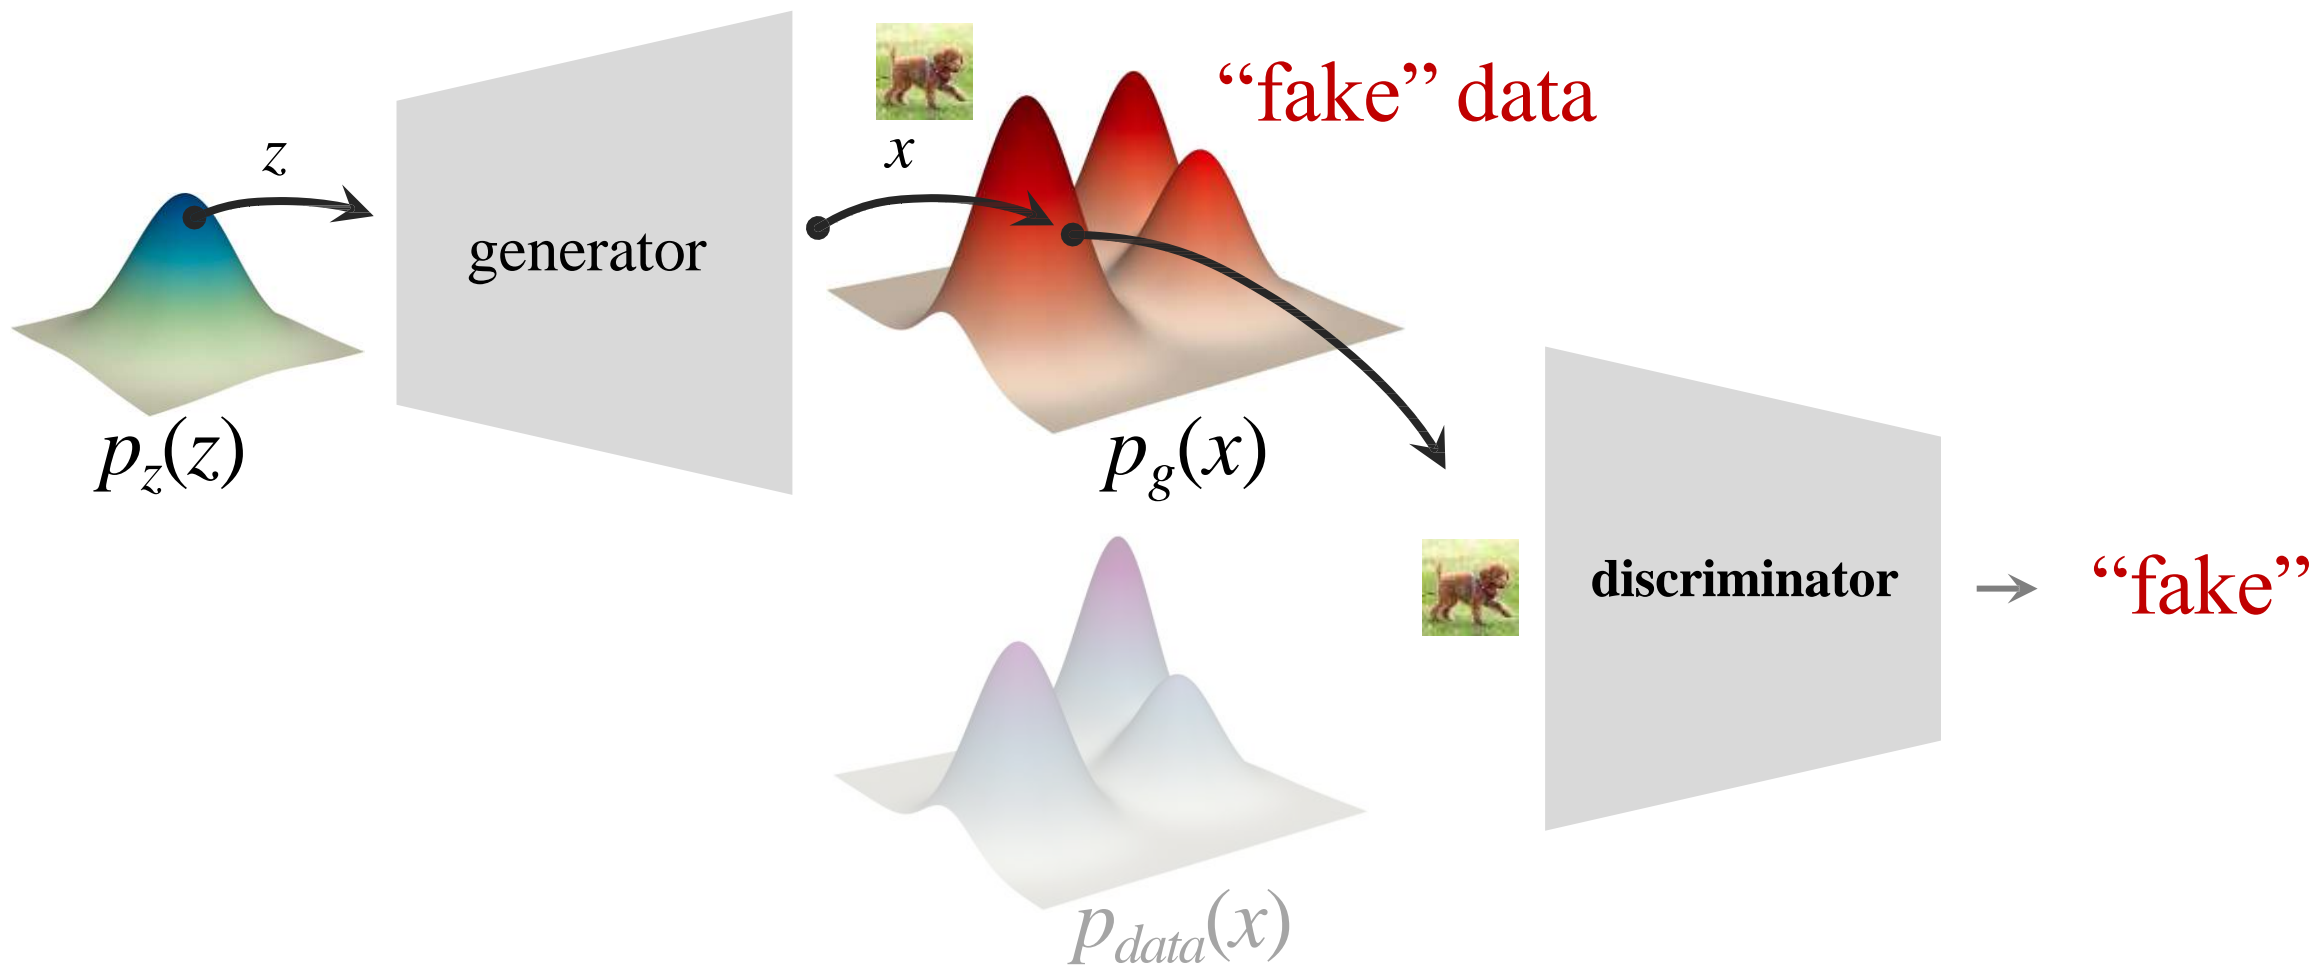
Generative Adversarial Networks

Representing **distribution difference** by a neural network



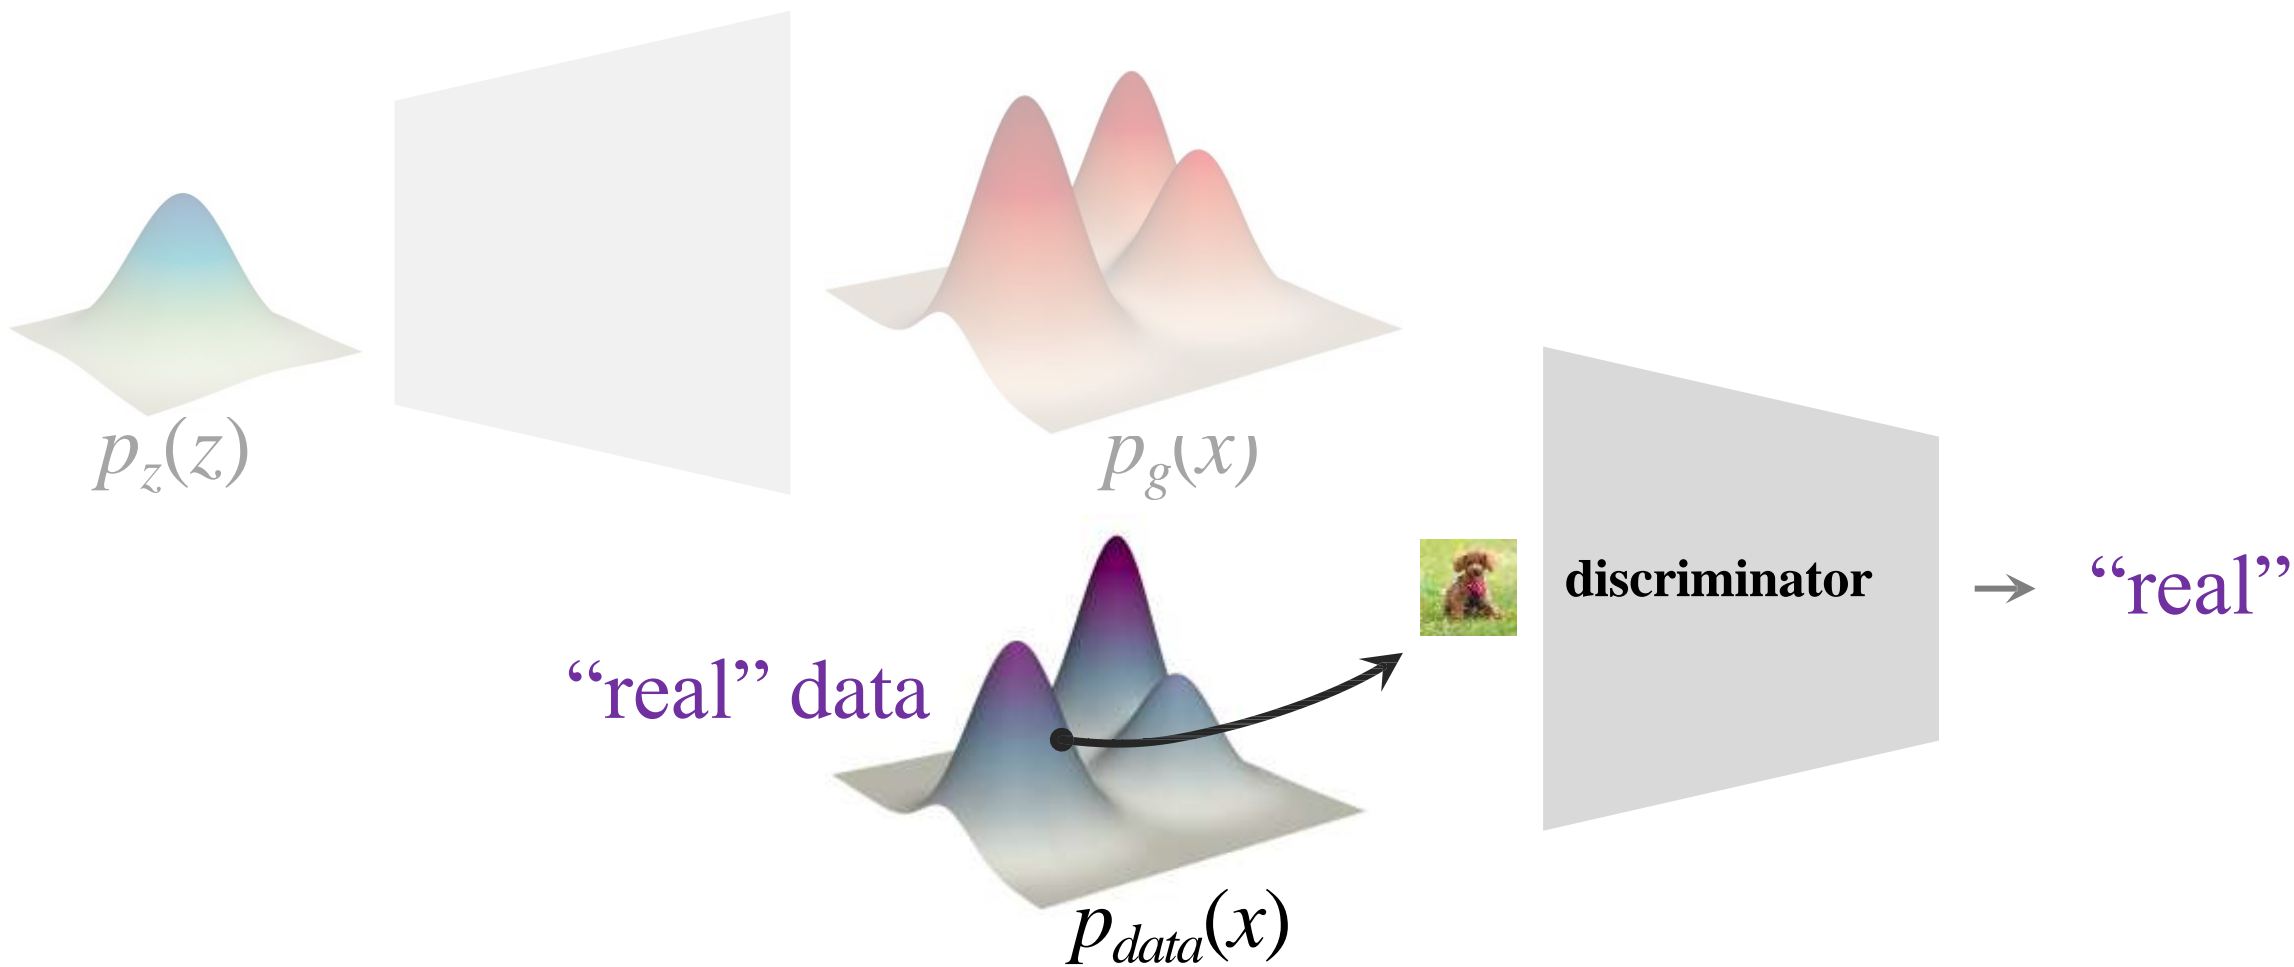
Generative Adversarial Networks

Representing **distribution difference** by a neural network




Generative Adversarial Networks

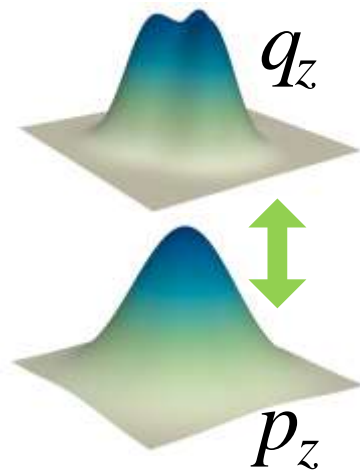
Representing **distribution difference** by a neural network




VAE



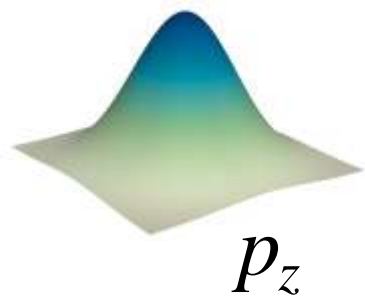
encoder



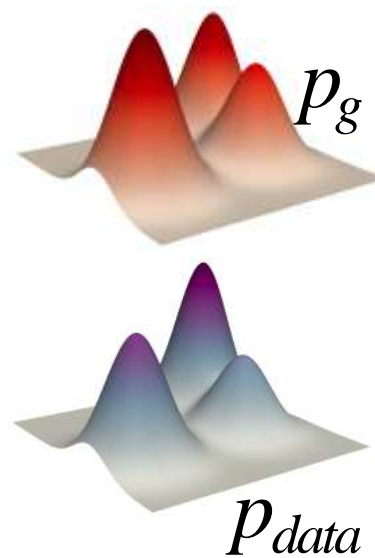
decoder



GAN



generator

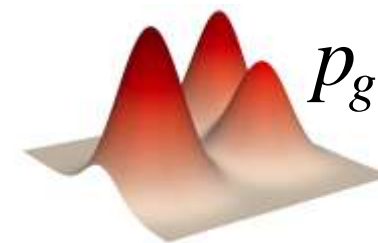
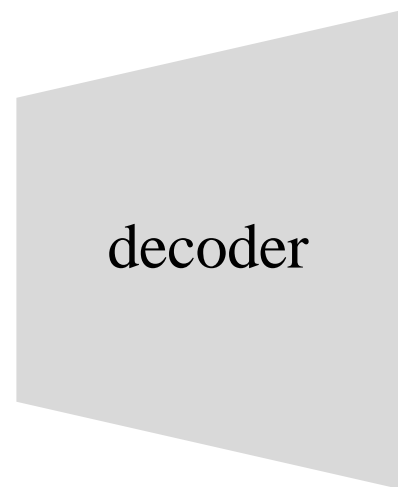




discriminator



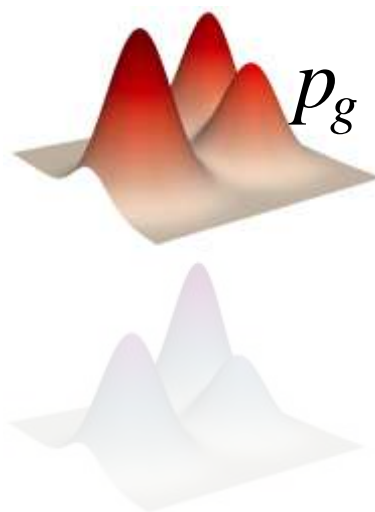


VAE

generation



GAN

generation



discriminator




Adversarial Objective

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

min-max process

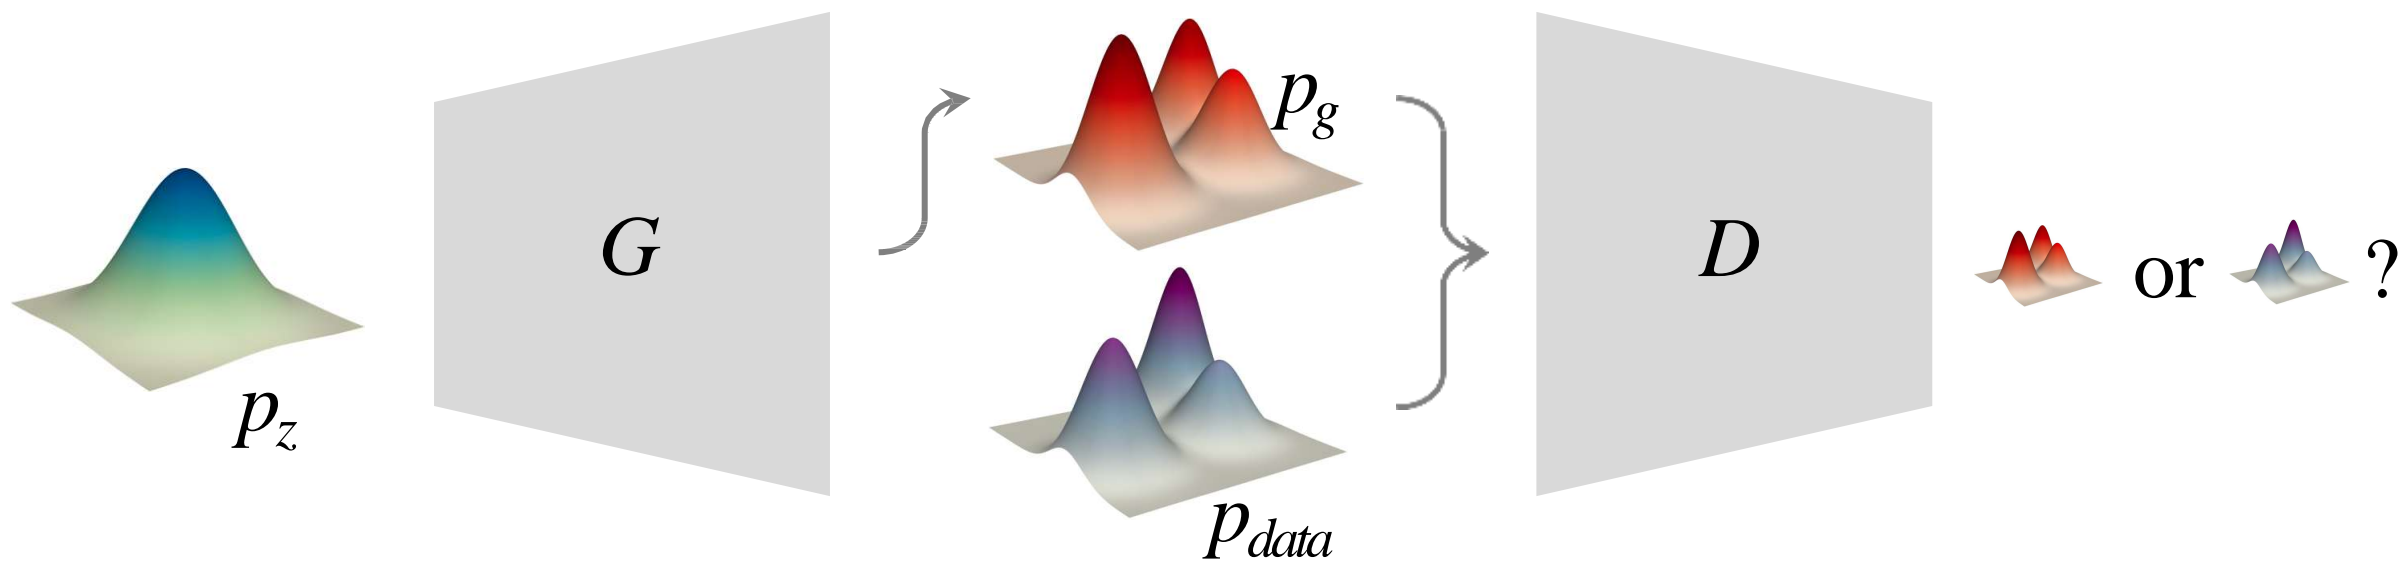
(vs. EM's max-max process)



Adversarial Objective: D-step

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

D-step: fix *G*, optimize *D*

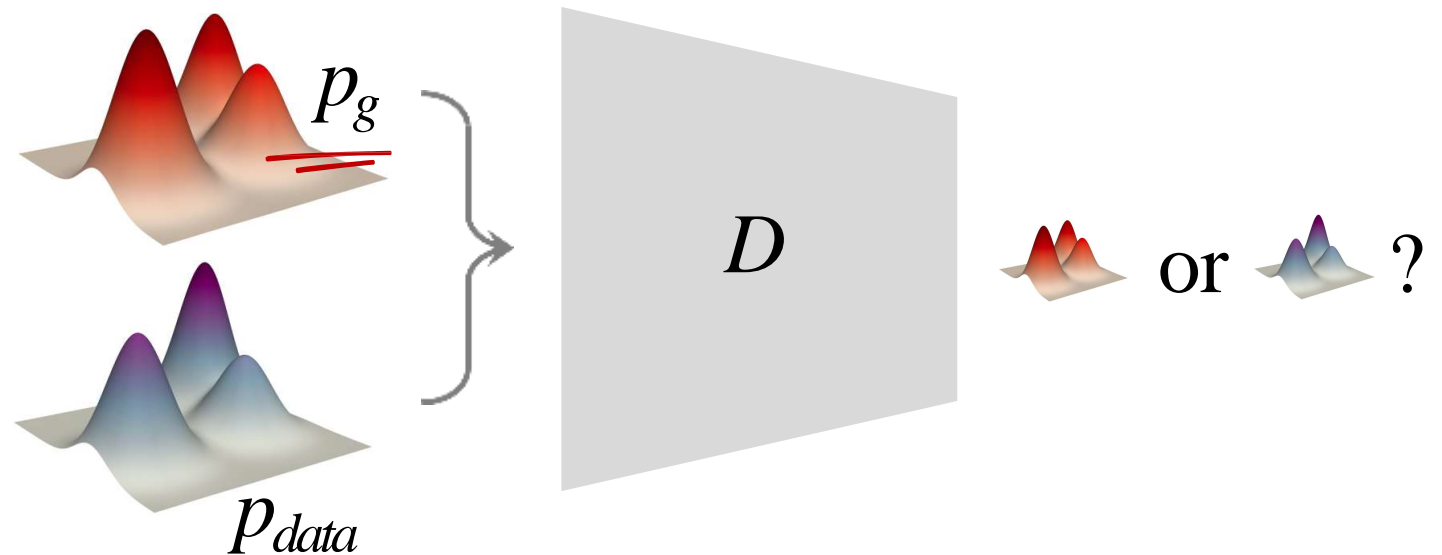


Adversarial Objective: D-step

$$\max_D \mathcal{L}(D) = \mathbb{E}_{x \sim p_{\text{data}}} [\log \underbrace{D(x)}_{\text{push to 1}}] + \mathbb{E}_{x \sim p_g} [\log(1 - \underbrace{D(x)}_{\text{push to 0}})]$$

D-step: fix *G*, optimize *D*


- *D* to classify real or fake
- binary logistic regression (sigmoid + cross-entropy)



Adversarial Objective: G-step

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


G -step: fix D , optimize G



Adversarial Objective: G-step

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

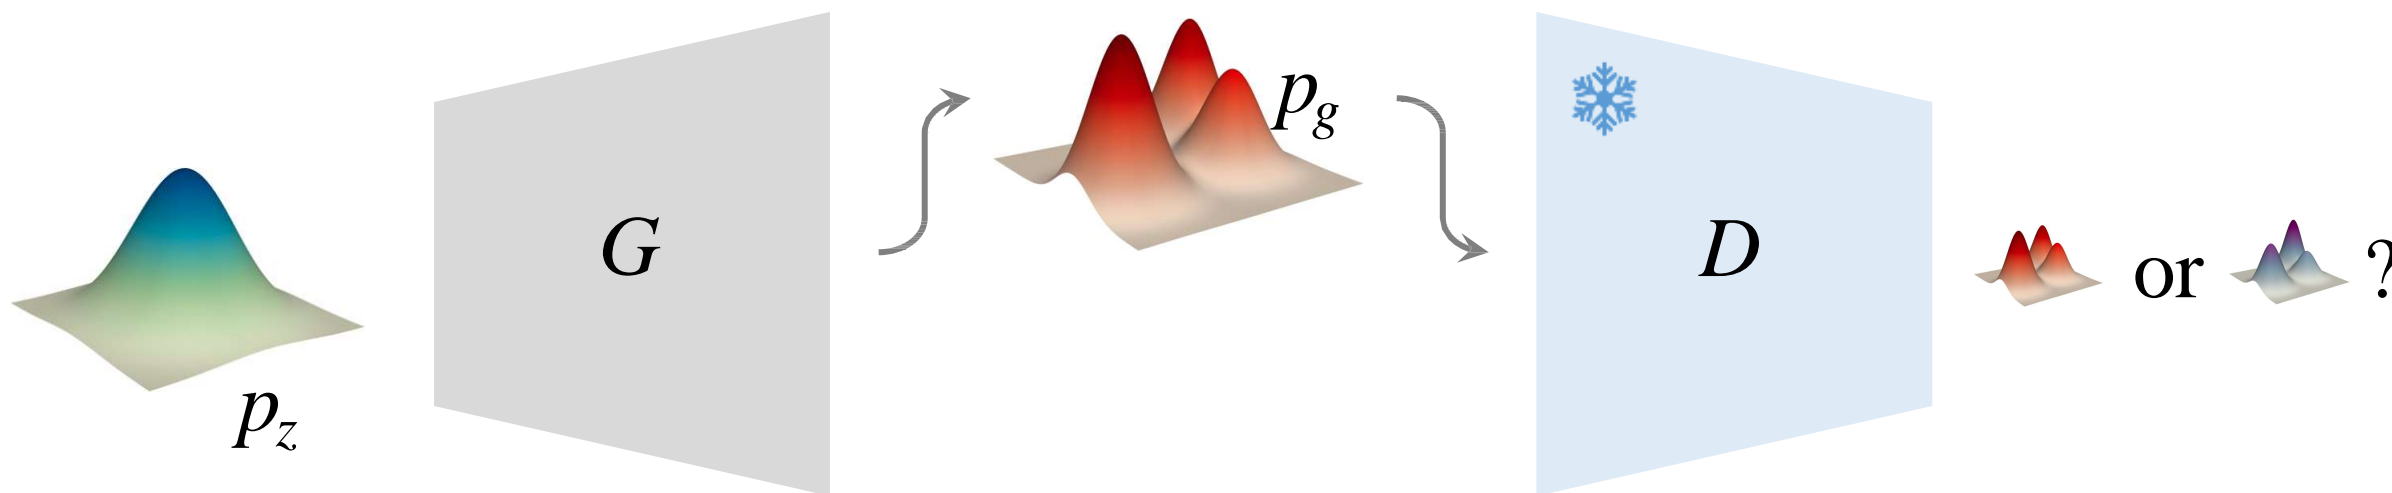
G-step: fix D , optimize G



Adversarial Objective: G-step

$$\min_G \mathcal{L}(G) = \mathbb{E}_{z \sim p_z} [\log(1 - \underbrace{D(G(z))}_{\text{push to 1}})]$$

- G -step: fix D , optimize G
- generate fake data such that D classifies it as “real”
- G to “confuse” D

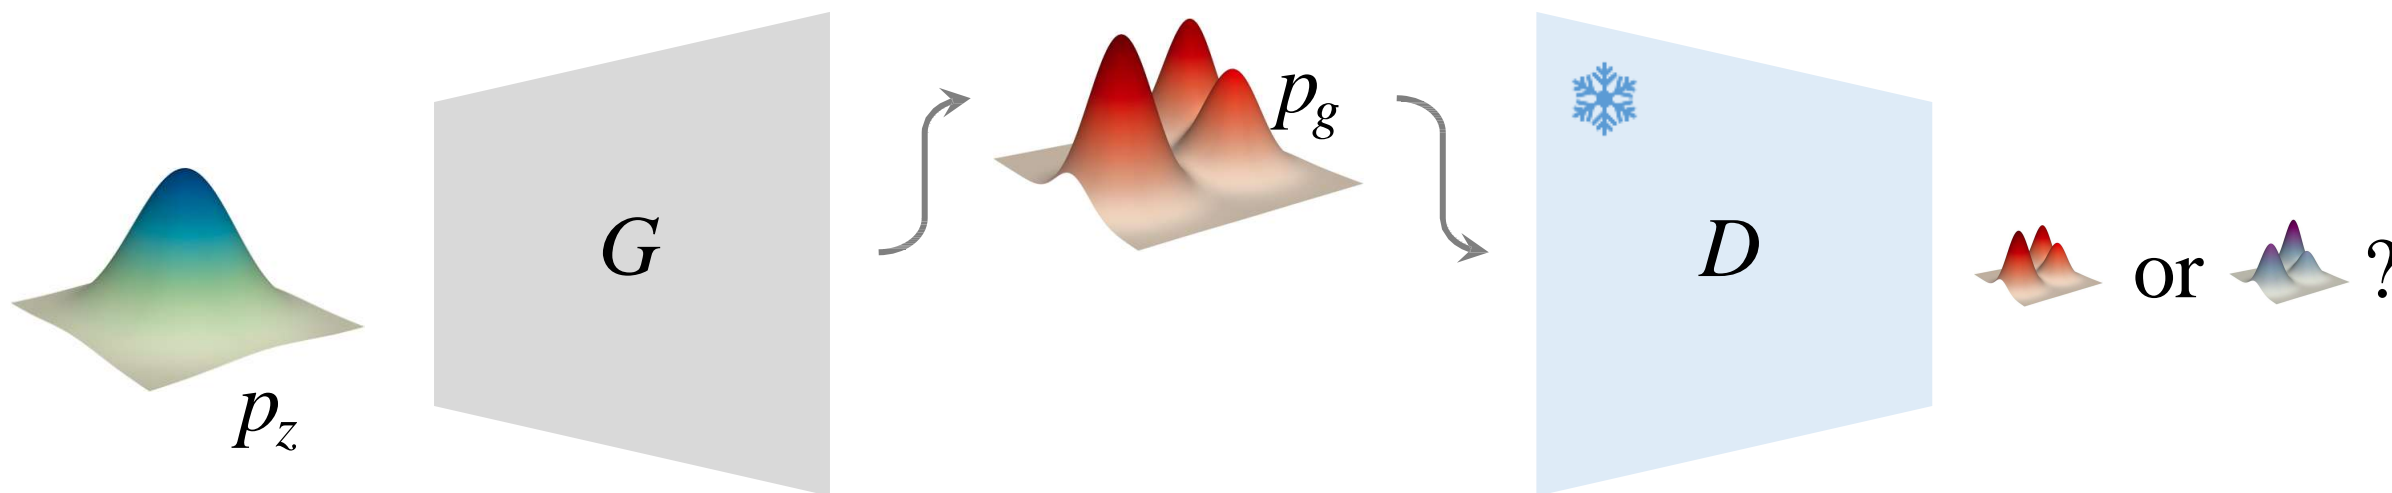


Adversarial Objective: G-step

a “flip” trick:

$$\max_G \mathcal{L}(G) = \mathbb{E}_{z \sim p_z} [\log(\underbrace{1 - D(G(z))}_{\text{push to 1}})]$$

- G -step: fix D , optimize G
- generate fake data such that D classifies it as “real”
- G to “confuse” D



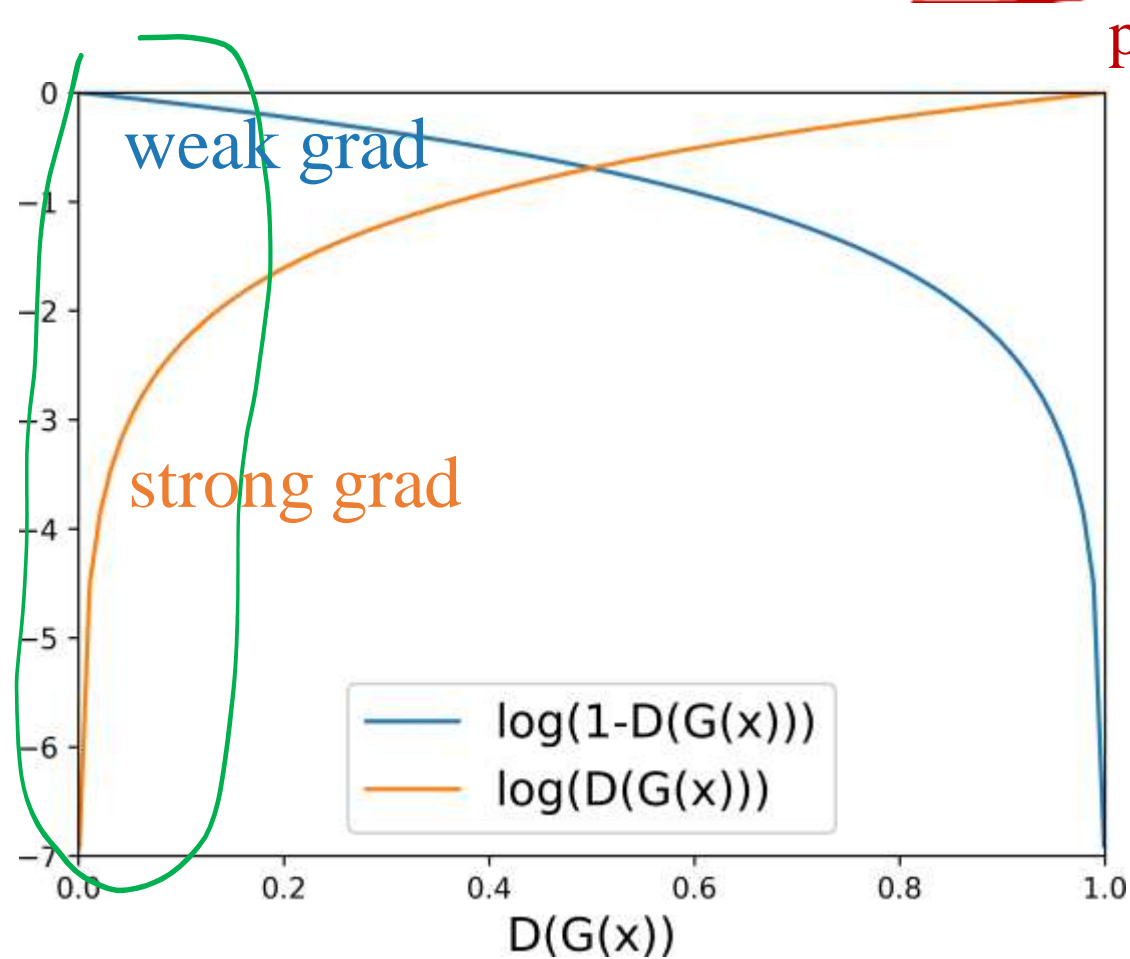
Adversarial Objective: G-step

a “flip” trick:

$$\max_G \mathcal{L}(G) = \mathbb{E}_{z \sim p_z} [\log(\underbrace{1 - D(G(z))}_{\text{push to 1}})]$$

Early in training:

- G is poor
- $D(G)$ is near 0



GAN algorithm annotated

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.

minibatch

for number of training iterations do

for k steps do

SGD

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

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

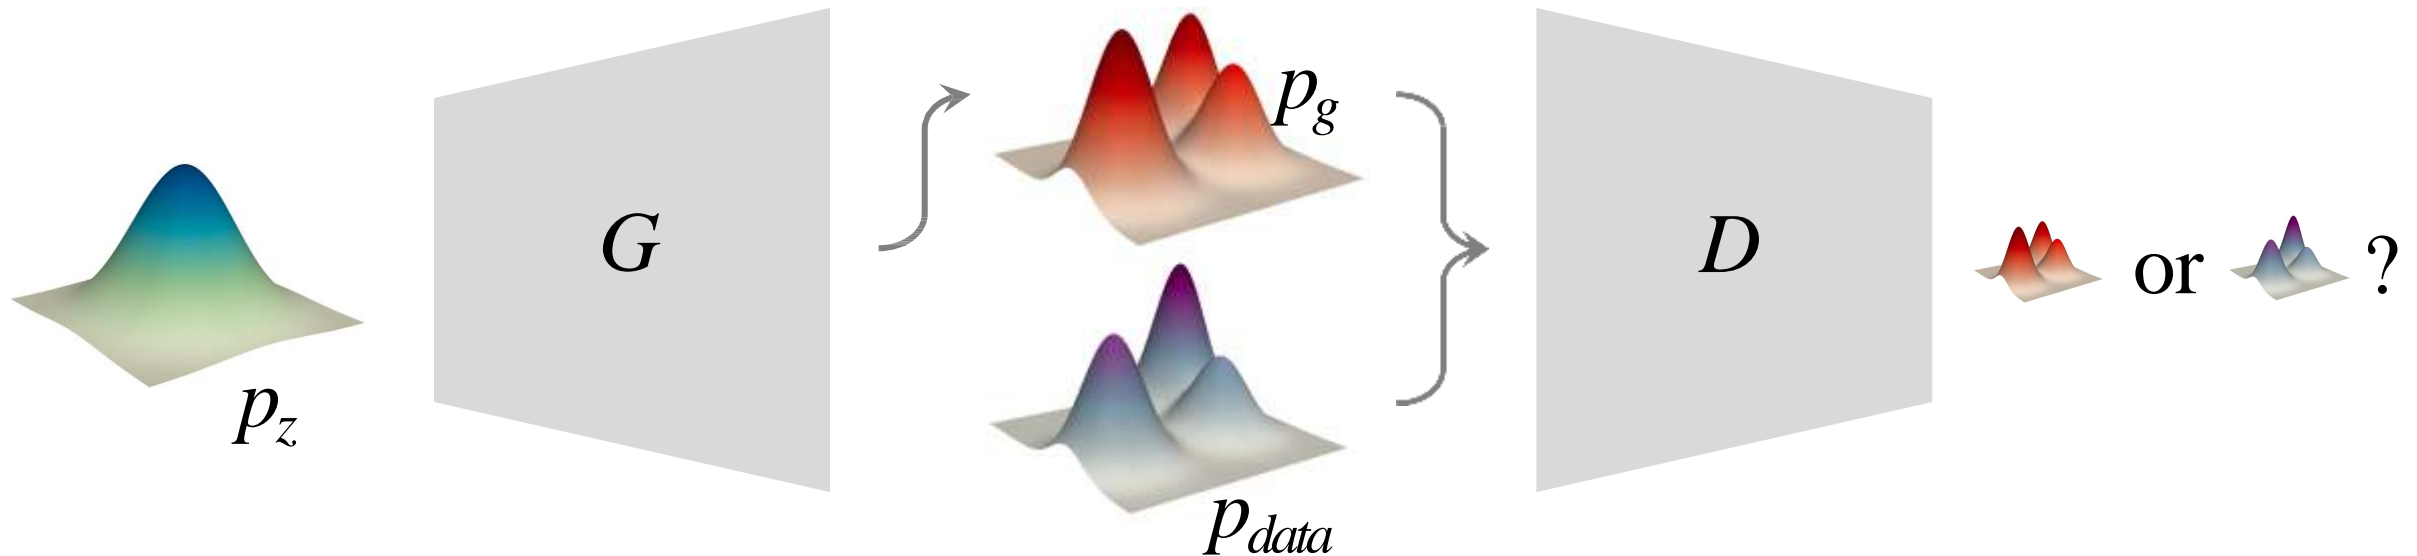
end for

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

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

end for

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



GAN algorithm annotated

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

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

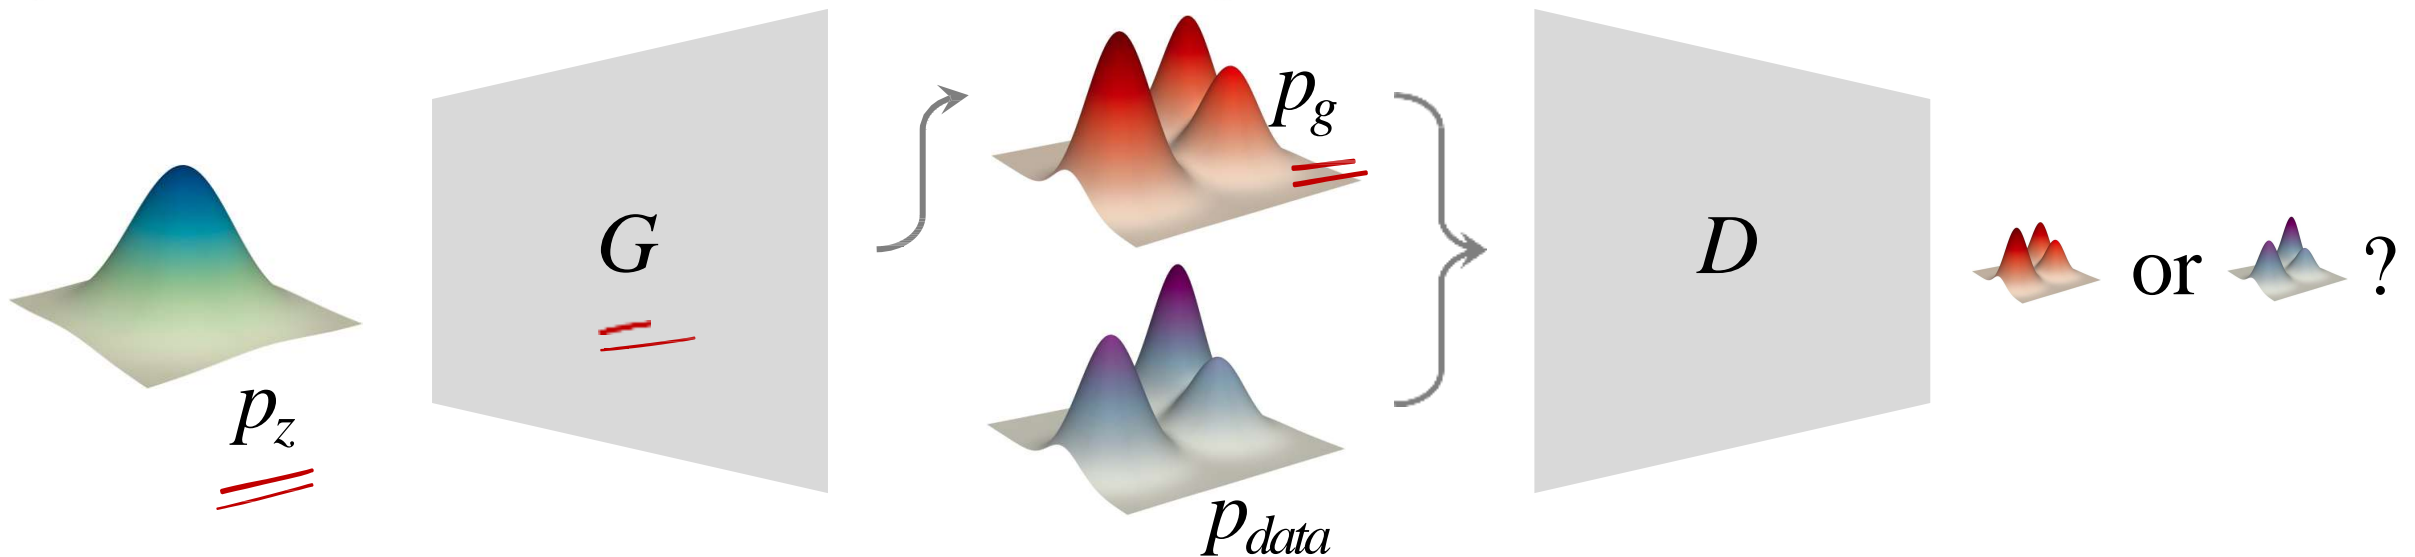
end for

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

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

end for

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



GAN algorithm annotated

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

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

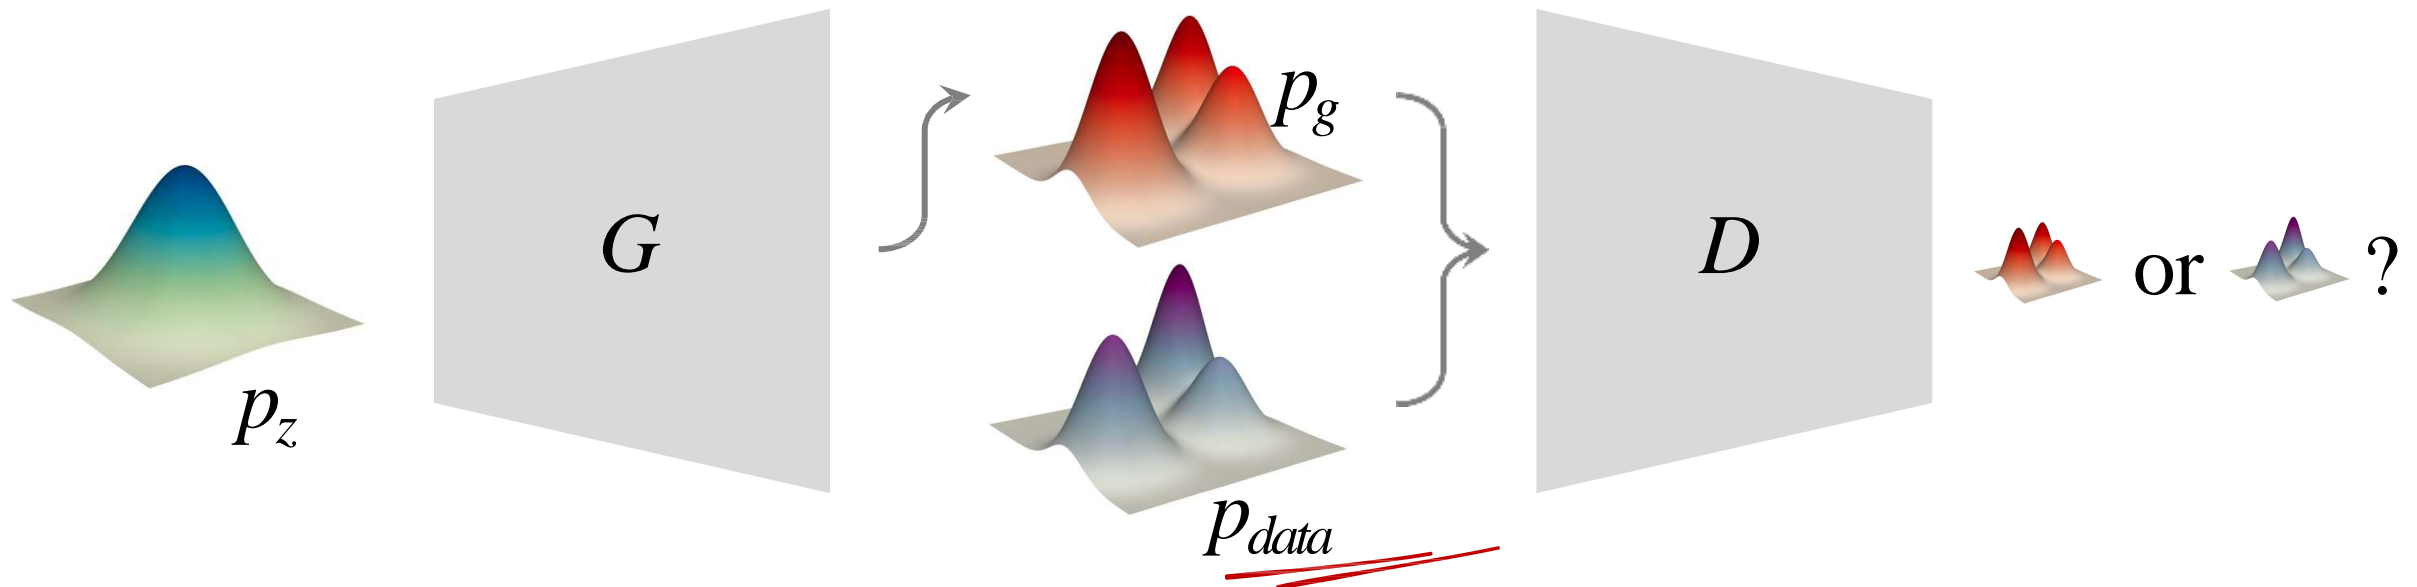
end for

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

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

end for

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



GAN algorithm annotated

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 κ steps **do**

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

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

gradient ascend
(maximize)


end for

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

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

end for

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



GAN algorithm annotated

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

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

end for

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

G-step

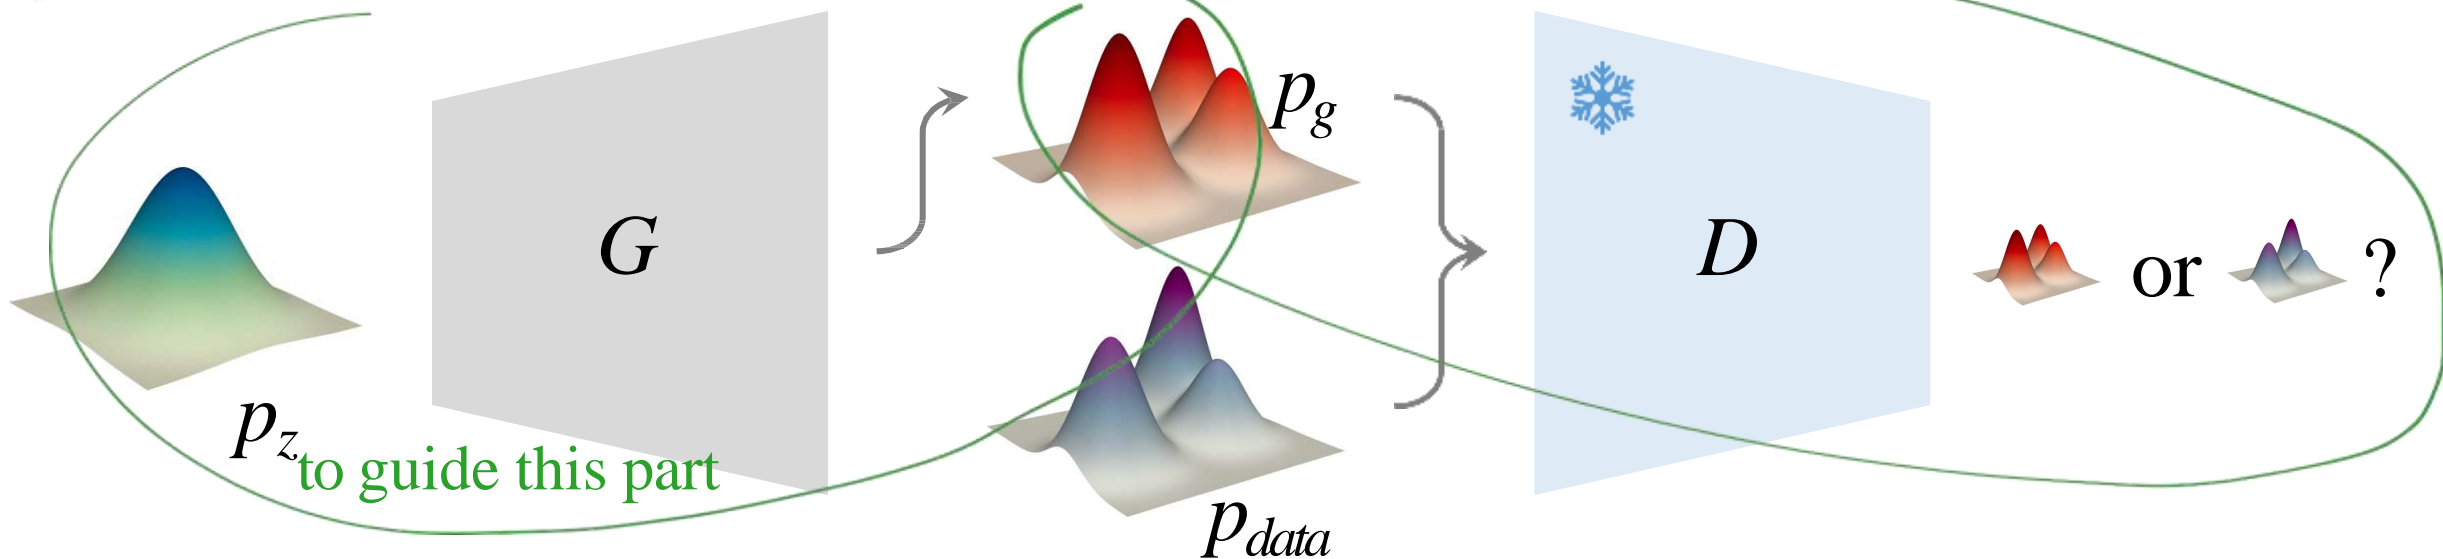
$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log (1 - D(G(z^{(i)}))).$$

**gradient descend
(minimize)**

end for

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

a parameterized
loss function



GAN algorithm annotated

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

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

end for

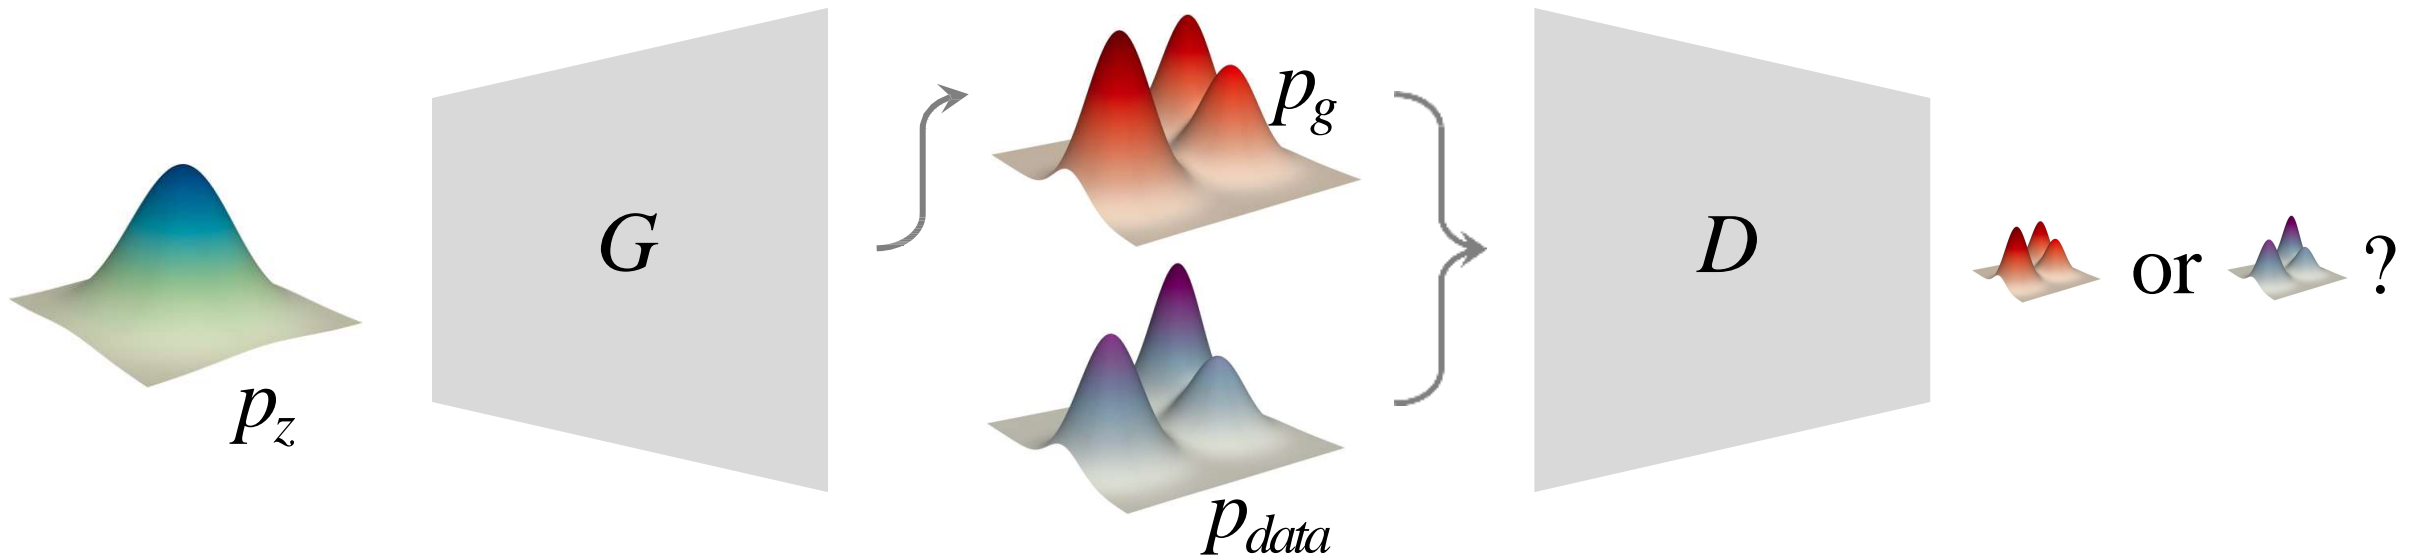
- Sample minibatch of m noise samples $\{z^{(1)}, \dots, z^{(m)}\}$ from noise prior $p_g(z)$.
- Update the generator by descending its stochastic gradient:

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

end for

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

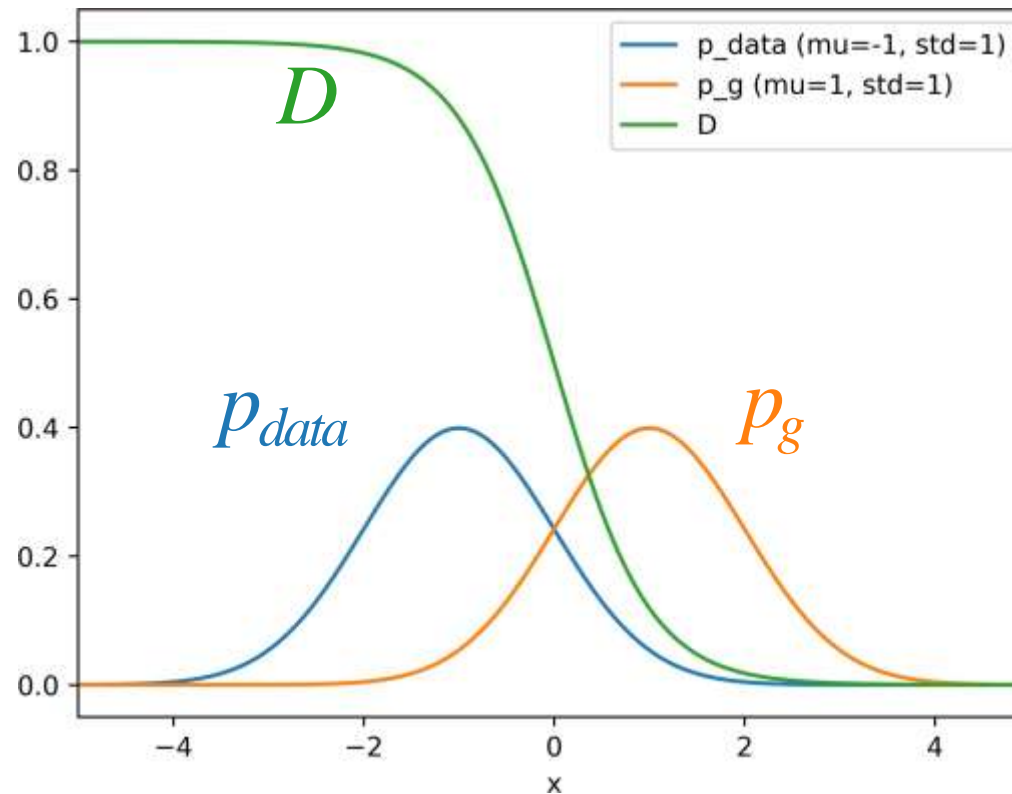
iterating
min-max



Theoretical Results


1. For any given G , the optimal D is:

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

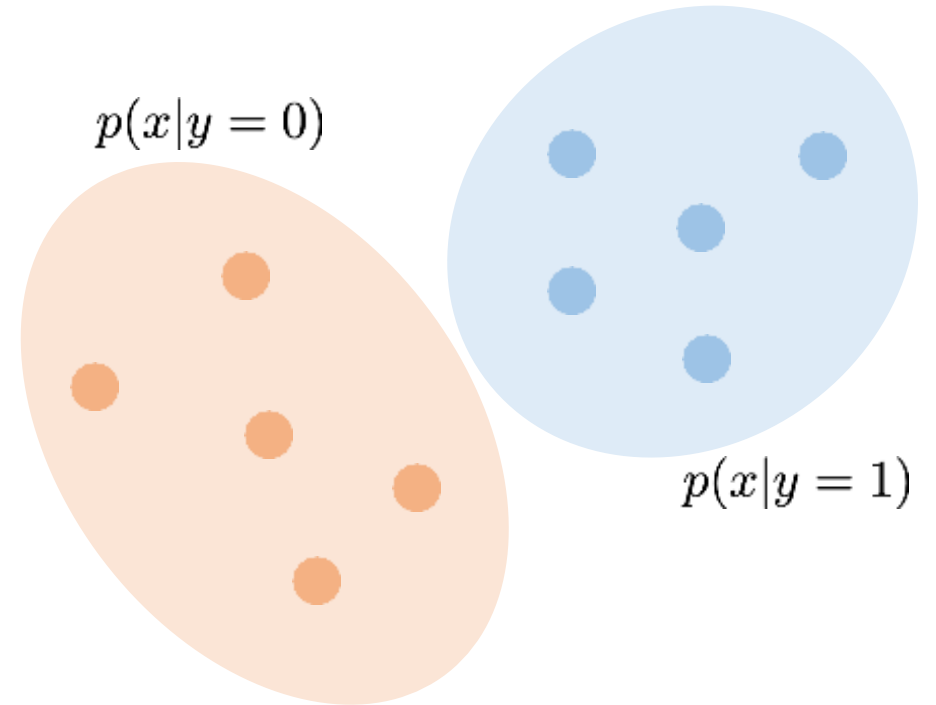


Recap : Discriminative vs. Generative

discriminative



generative



Theoretical Results

2. With the optimal D_G , the objective function is:

$$\mathcal{L}(D^*, G) = 2D_{JS}(p_{\text{data}} || p_g) - 2 \log 2$$

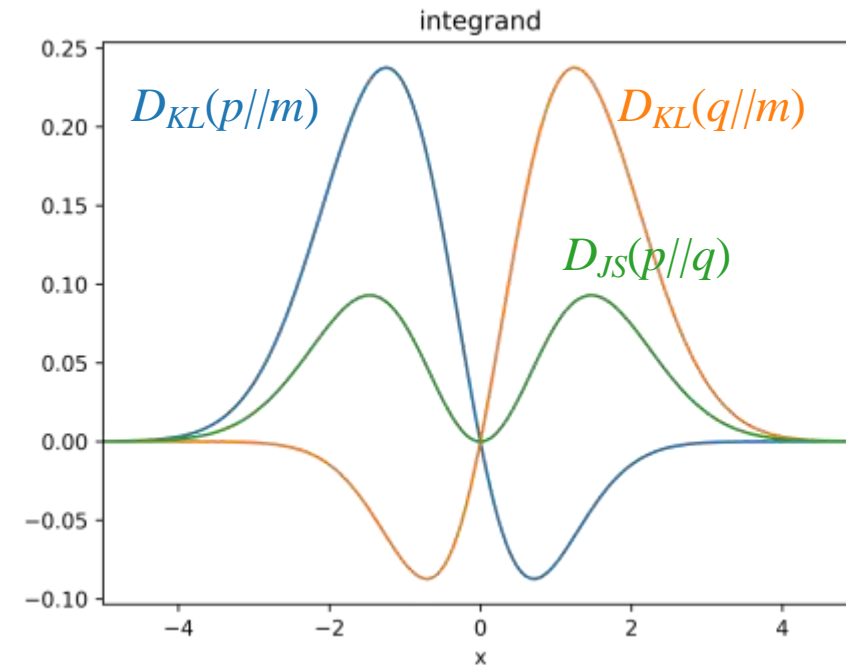
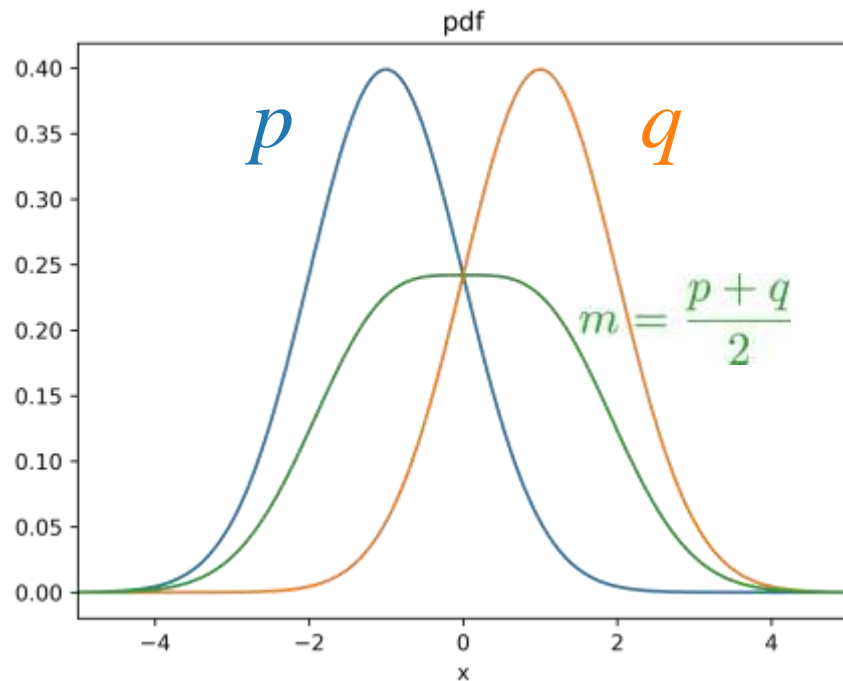
where D_{JS} is Jensen–Shannon divergence

Background: Jensen–Shannon divergence

D_{JS} : “total divergence to the average”

$$D_{KL}(p||q) = \int_x p(x) \log \frac{p(x)}{q(x)} dx$$

$$D_{JS}(p||q) \triangleq \frac{1}{2} D_{KL}(p||\frac{p+q}{2}) + \frac{1}{2} D_{KL}(q||\frac{p+q}{2})$$



Background: Jensen–Shannon divergence

D_{JS} : “total divergence to the average”

$$D_{JS}(p\|q) \triangleq \frac{1}{2}D_{KL}(p\|\frac{p+q}{2}) + \frac{1}{2}D_{KL}(q\|\frac{p+q}{2})$$

Properties:


- D_{JS} is symmetric; D_{KL} is not
- D_{JS} is bounded: $[0, \log 2]$; D_{KL} is unbounded: $[0, \infty)$
- D_{JS} is more stable

Theoretical Results

2. With the optimal D_G , the objective function is:

$$\mathcal{L}(D^*, G) = 2D_{JS}(p_{\text{data}} || p_g) - 2 \log 2$$

GAN optimizes for Jensen–Shannon divergence.




Theoretical Results

3. Global optimality is achieved at $p_g = p_{data}$

$$\mathcal{L}(D^*, G^*) = \cancel{2D_{JS}(p_{data} || p_g)} - 2 \log 2$$

$\Rightarrow \emptyset$



Theoretical Results: Summary

1. For any given G , the optimal D is:

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

2. With optimal D_G , GAN optimizes for Jensen–Shannon divergence:

$$\mathcal{L}(D^*, G) = 2D_{JS}(p_{\text{data}} \| p_g) - 2 \log 2$$

3. Global optimality is achieved at $p_g = p_{\text{data}}$

$$\mathcal{L}(D^*, G^*) = -2 \log 2$$

Theoretical Results: Summary

1. For any given G , the optimal D is:

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

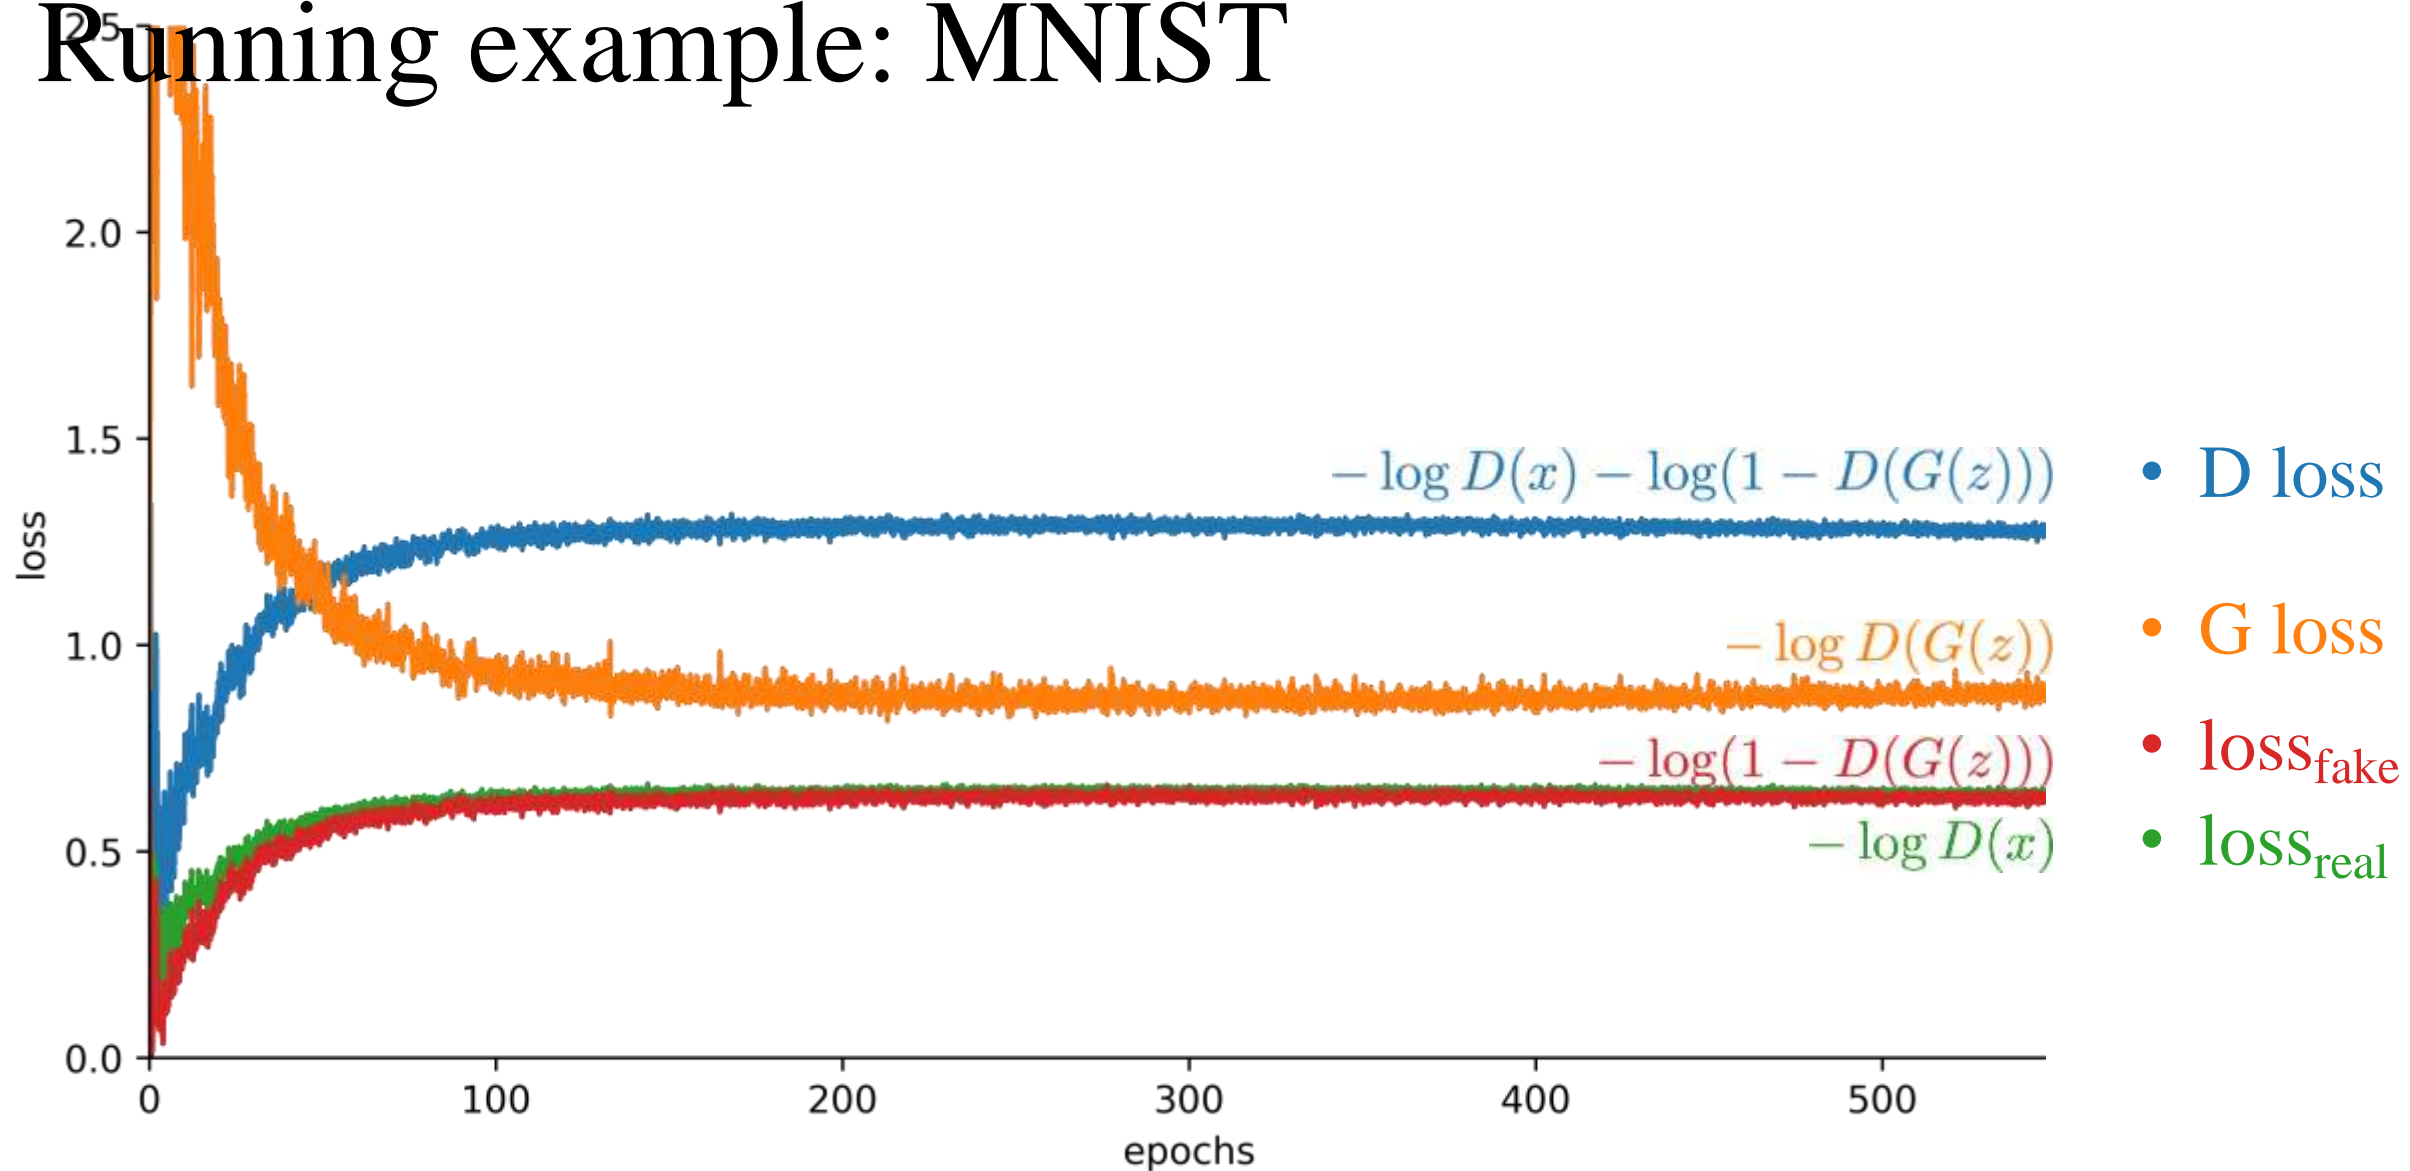
2. With optimal D_G , GAN optimizes for Jensen–Shannon divergence:

$$\mathcal{L}(D^*, G) = 2D_{JS}(p_{\text{data}} || p_g) - 2 \log 2$$

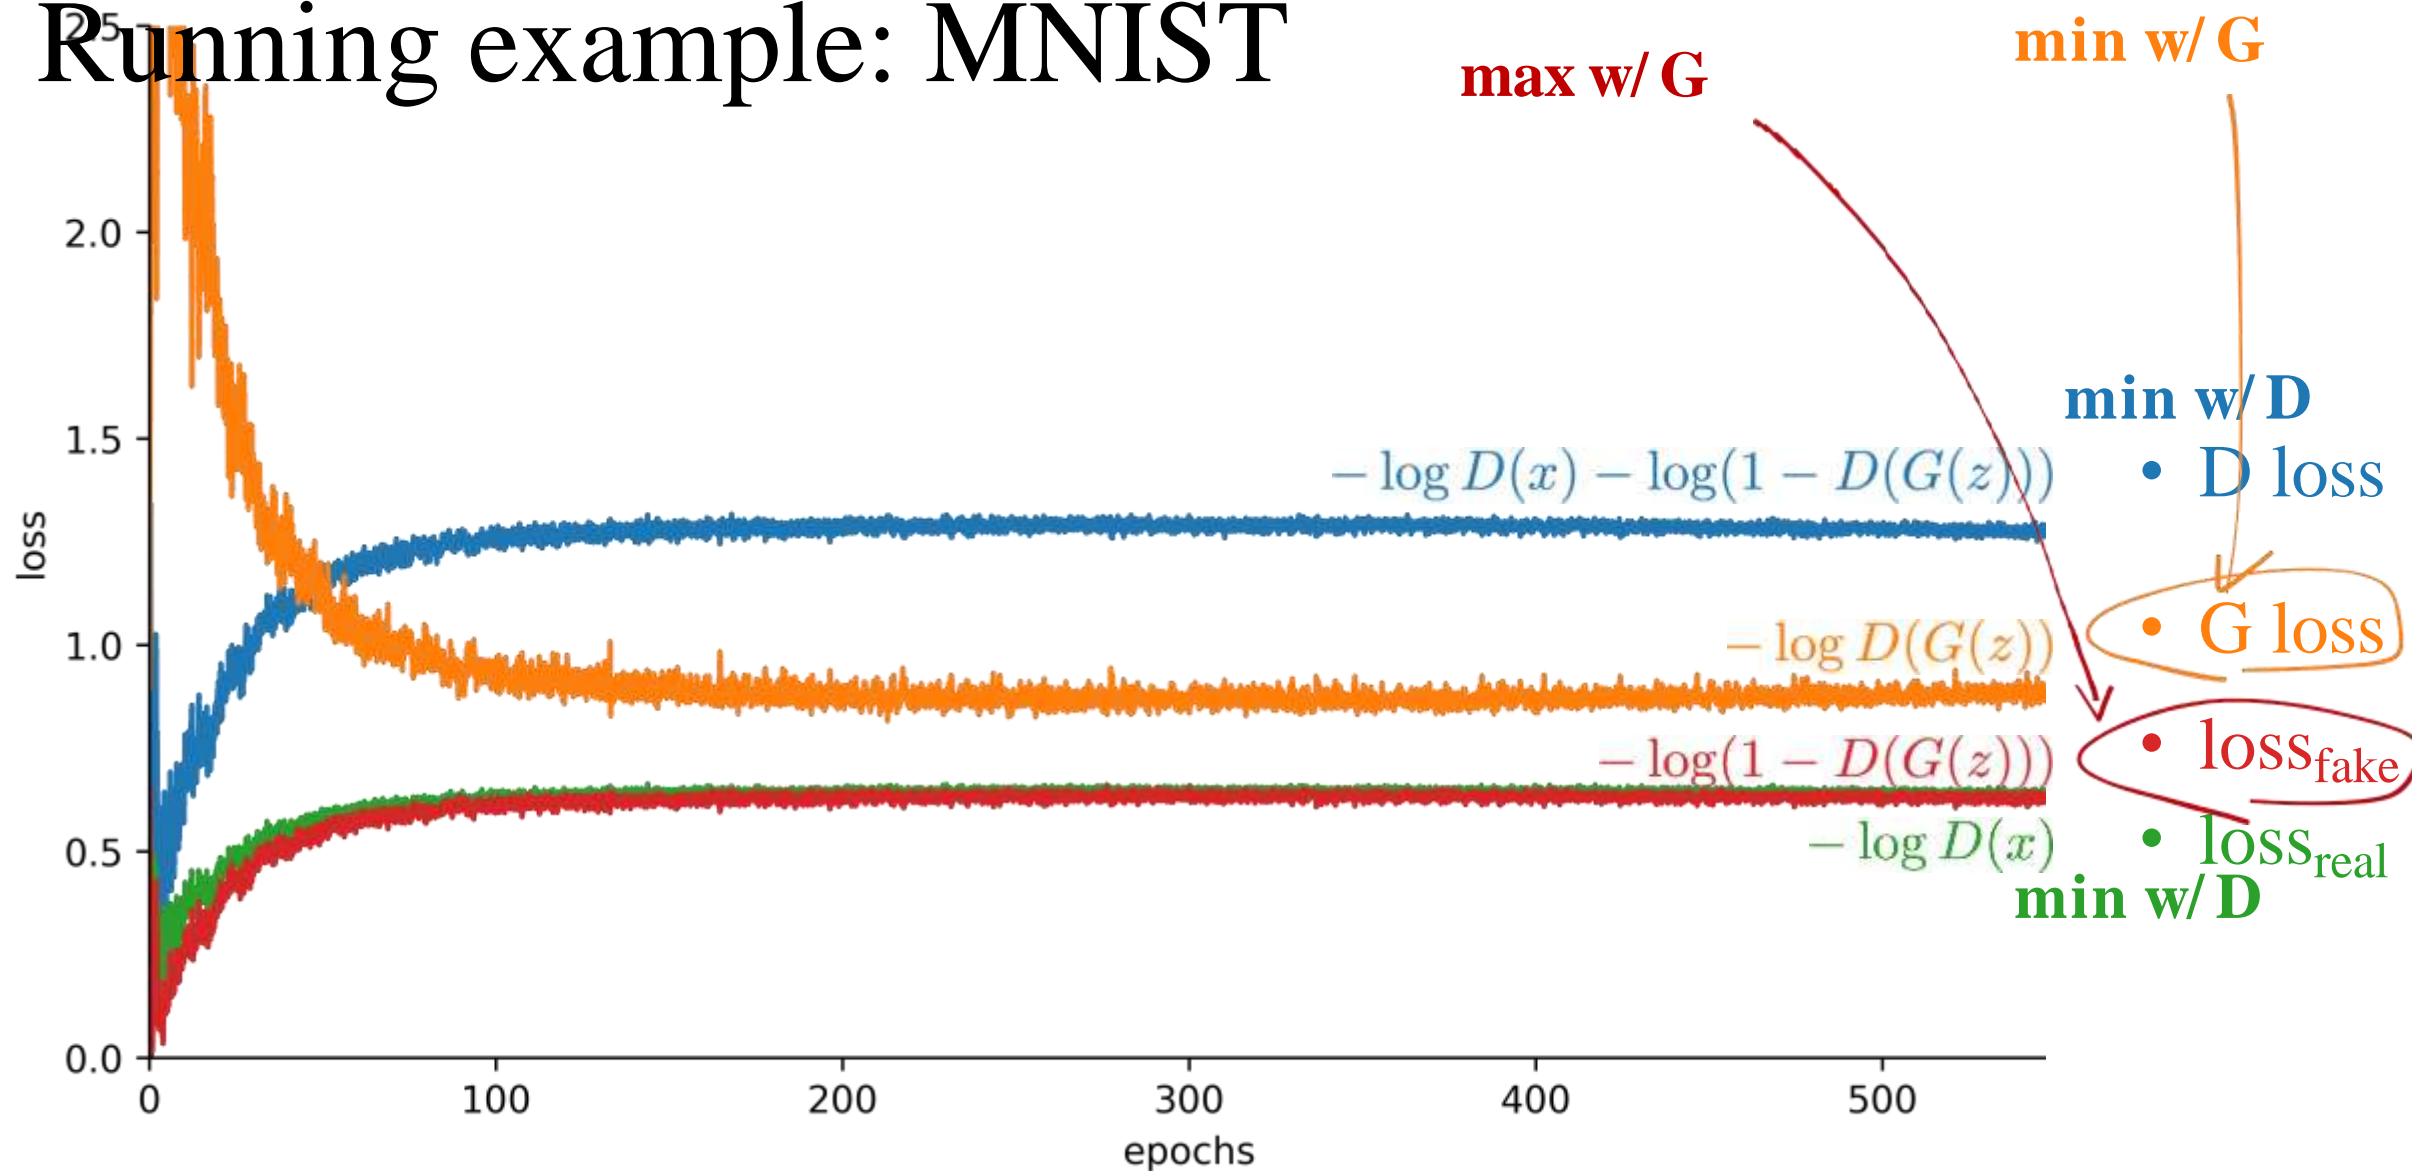
3. Global optimality is achieved at $p_g = p_{\text{data}}$

$$\begin{aligned} L(G, D^*) &= \int_x \left(p_r(x) \log(D^*(x)) + p_g(x) \log(1 - D^*(x)) \right) dx \\ &= \log \frac{1}{2} \int_x p_r(x) dx + \log \frac{1}{2} \int_x p_g(x) dx && D^*(x) = 1/2 \\ &= -2 \log 2 \end{aligned}$$

Running example: MNIST



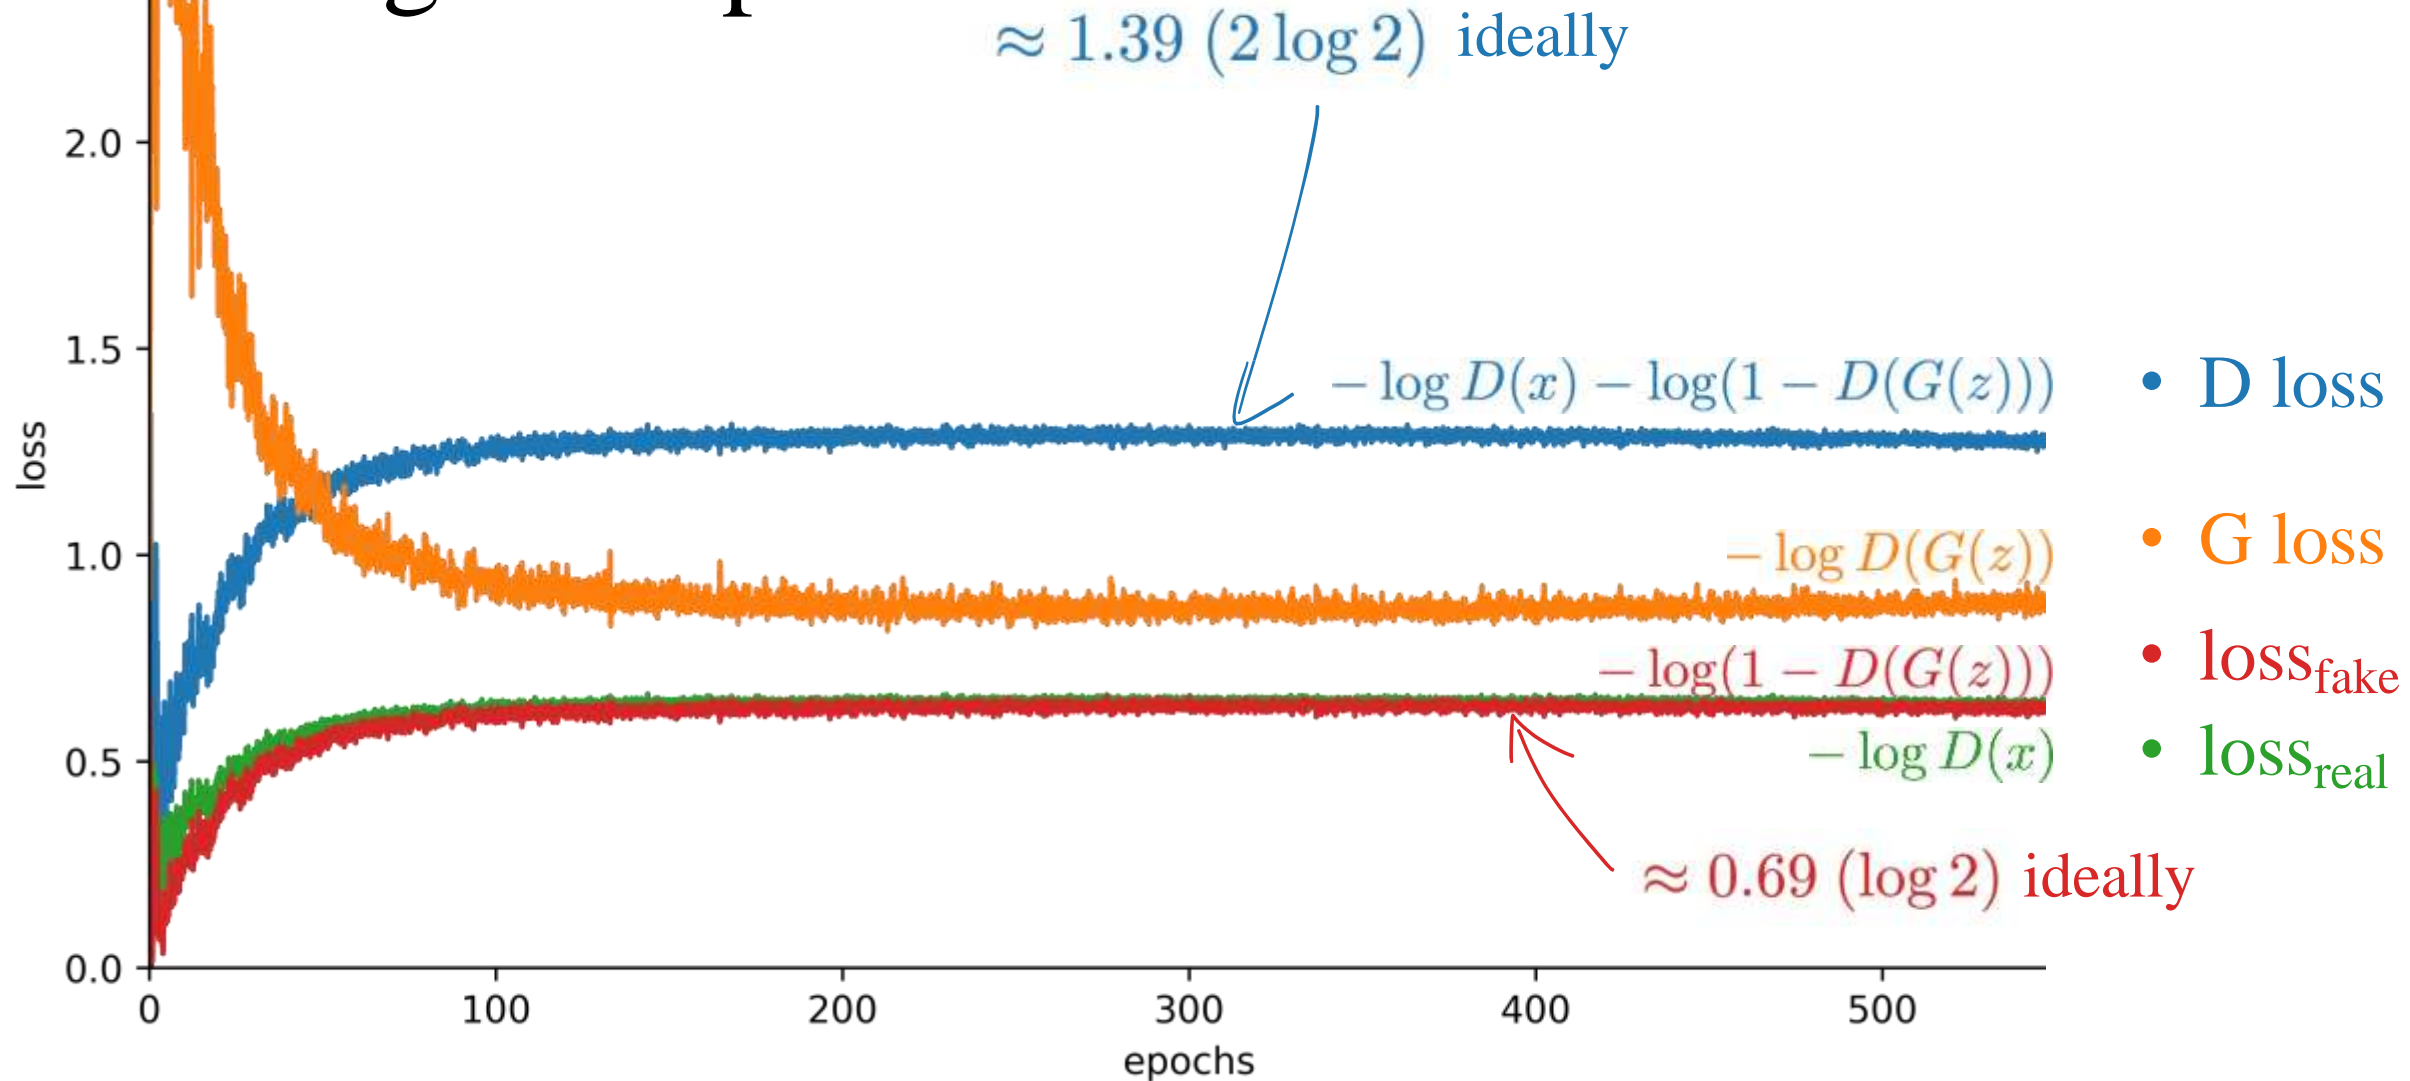
Running example: MNIST



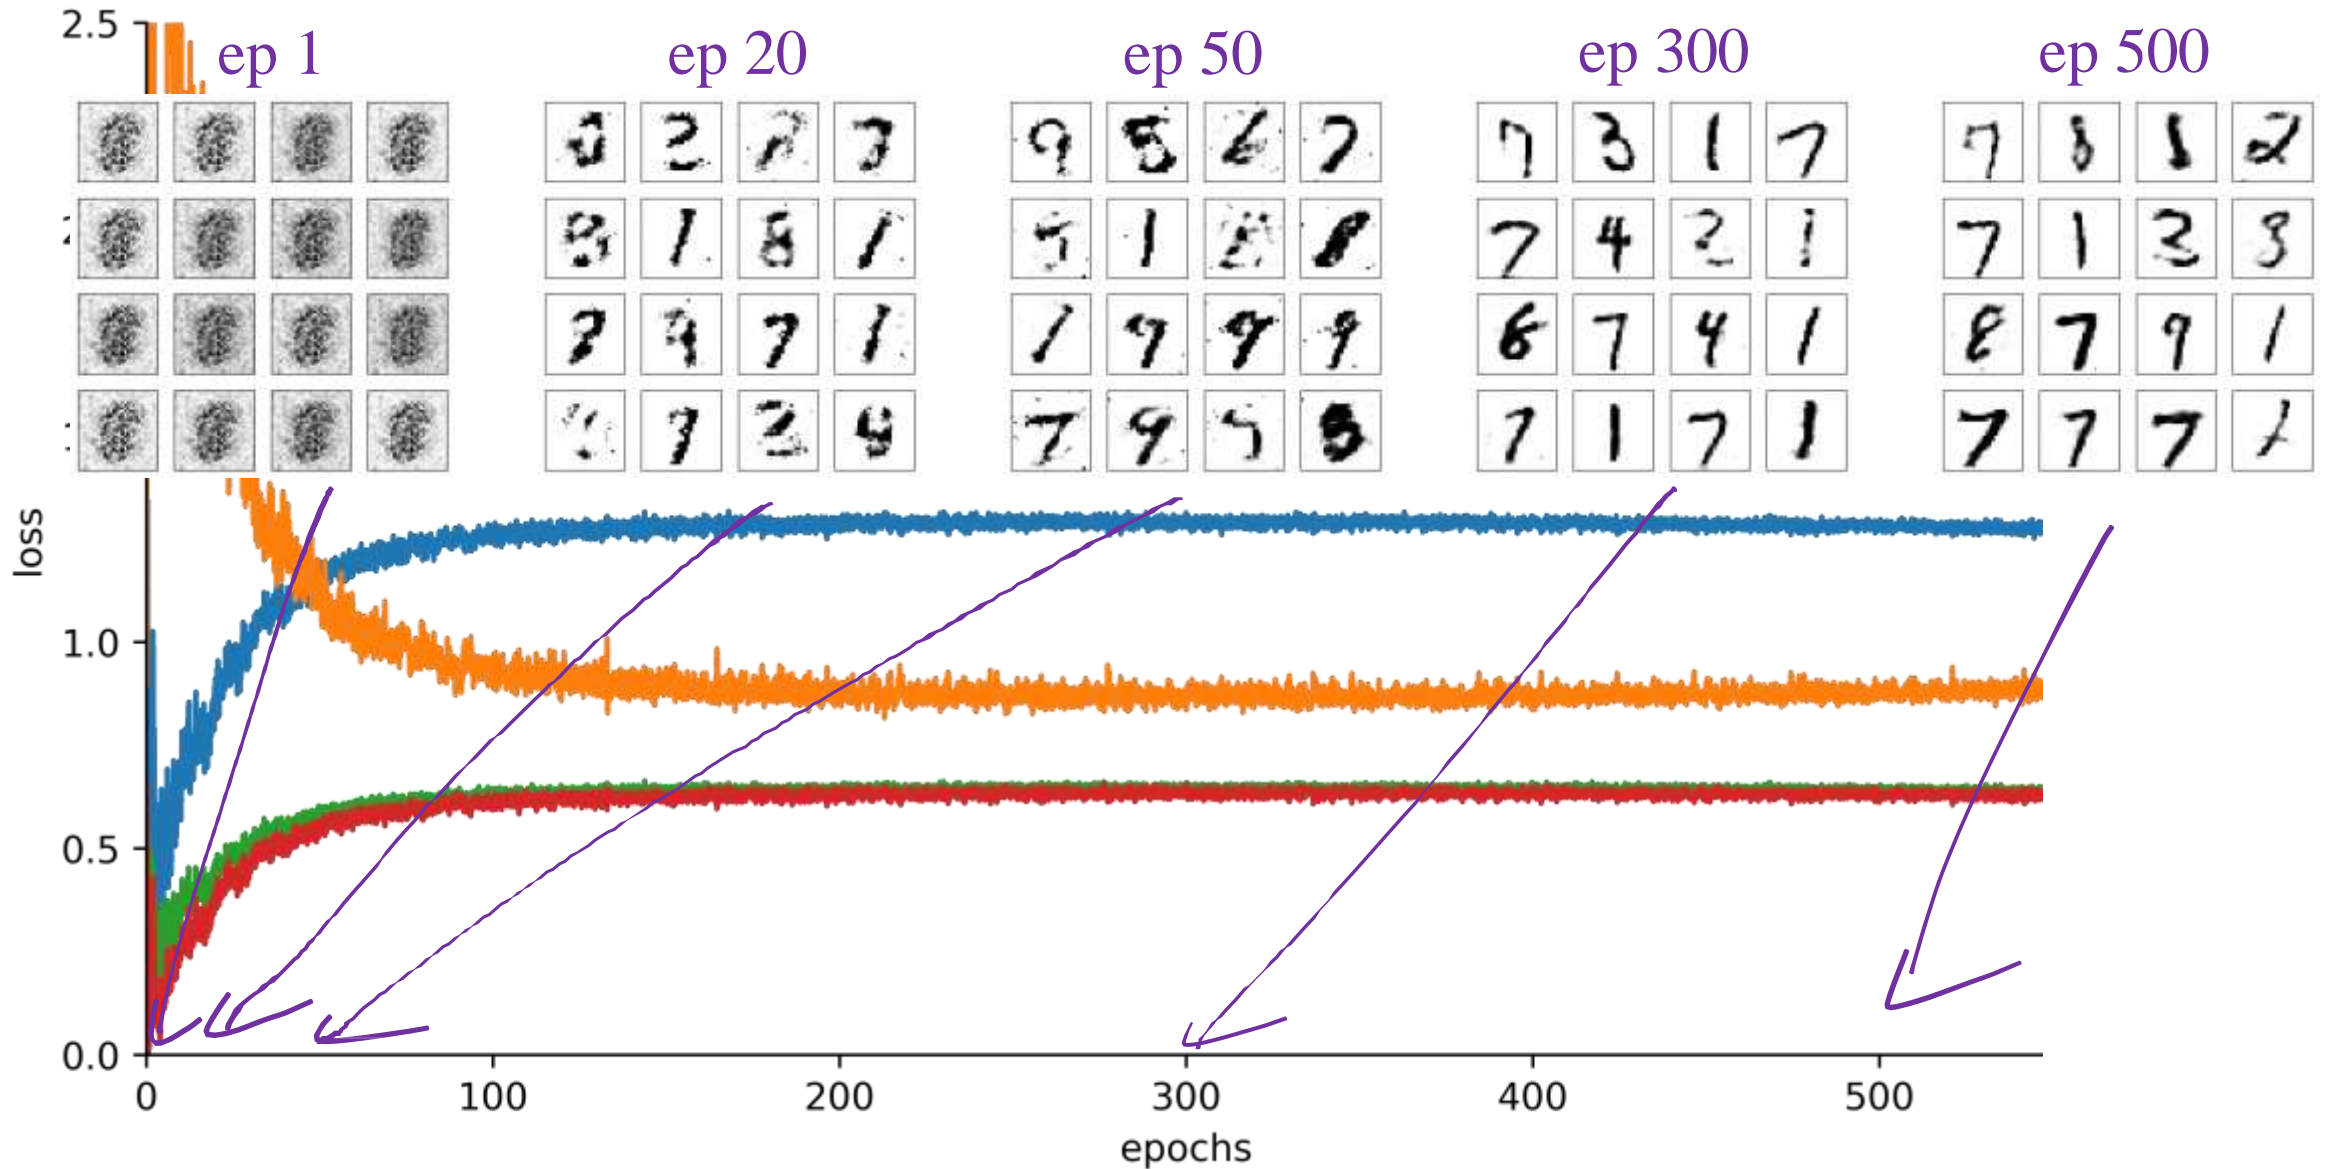
*All objectives are negative of their original form

Code adapted from: <https://github.com/prcastro/pytorch-gan/tree/master>

Running example: MNIST



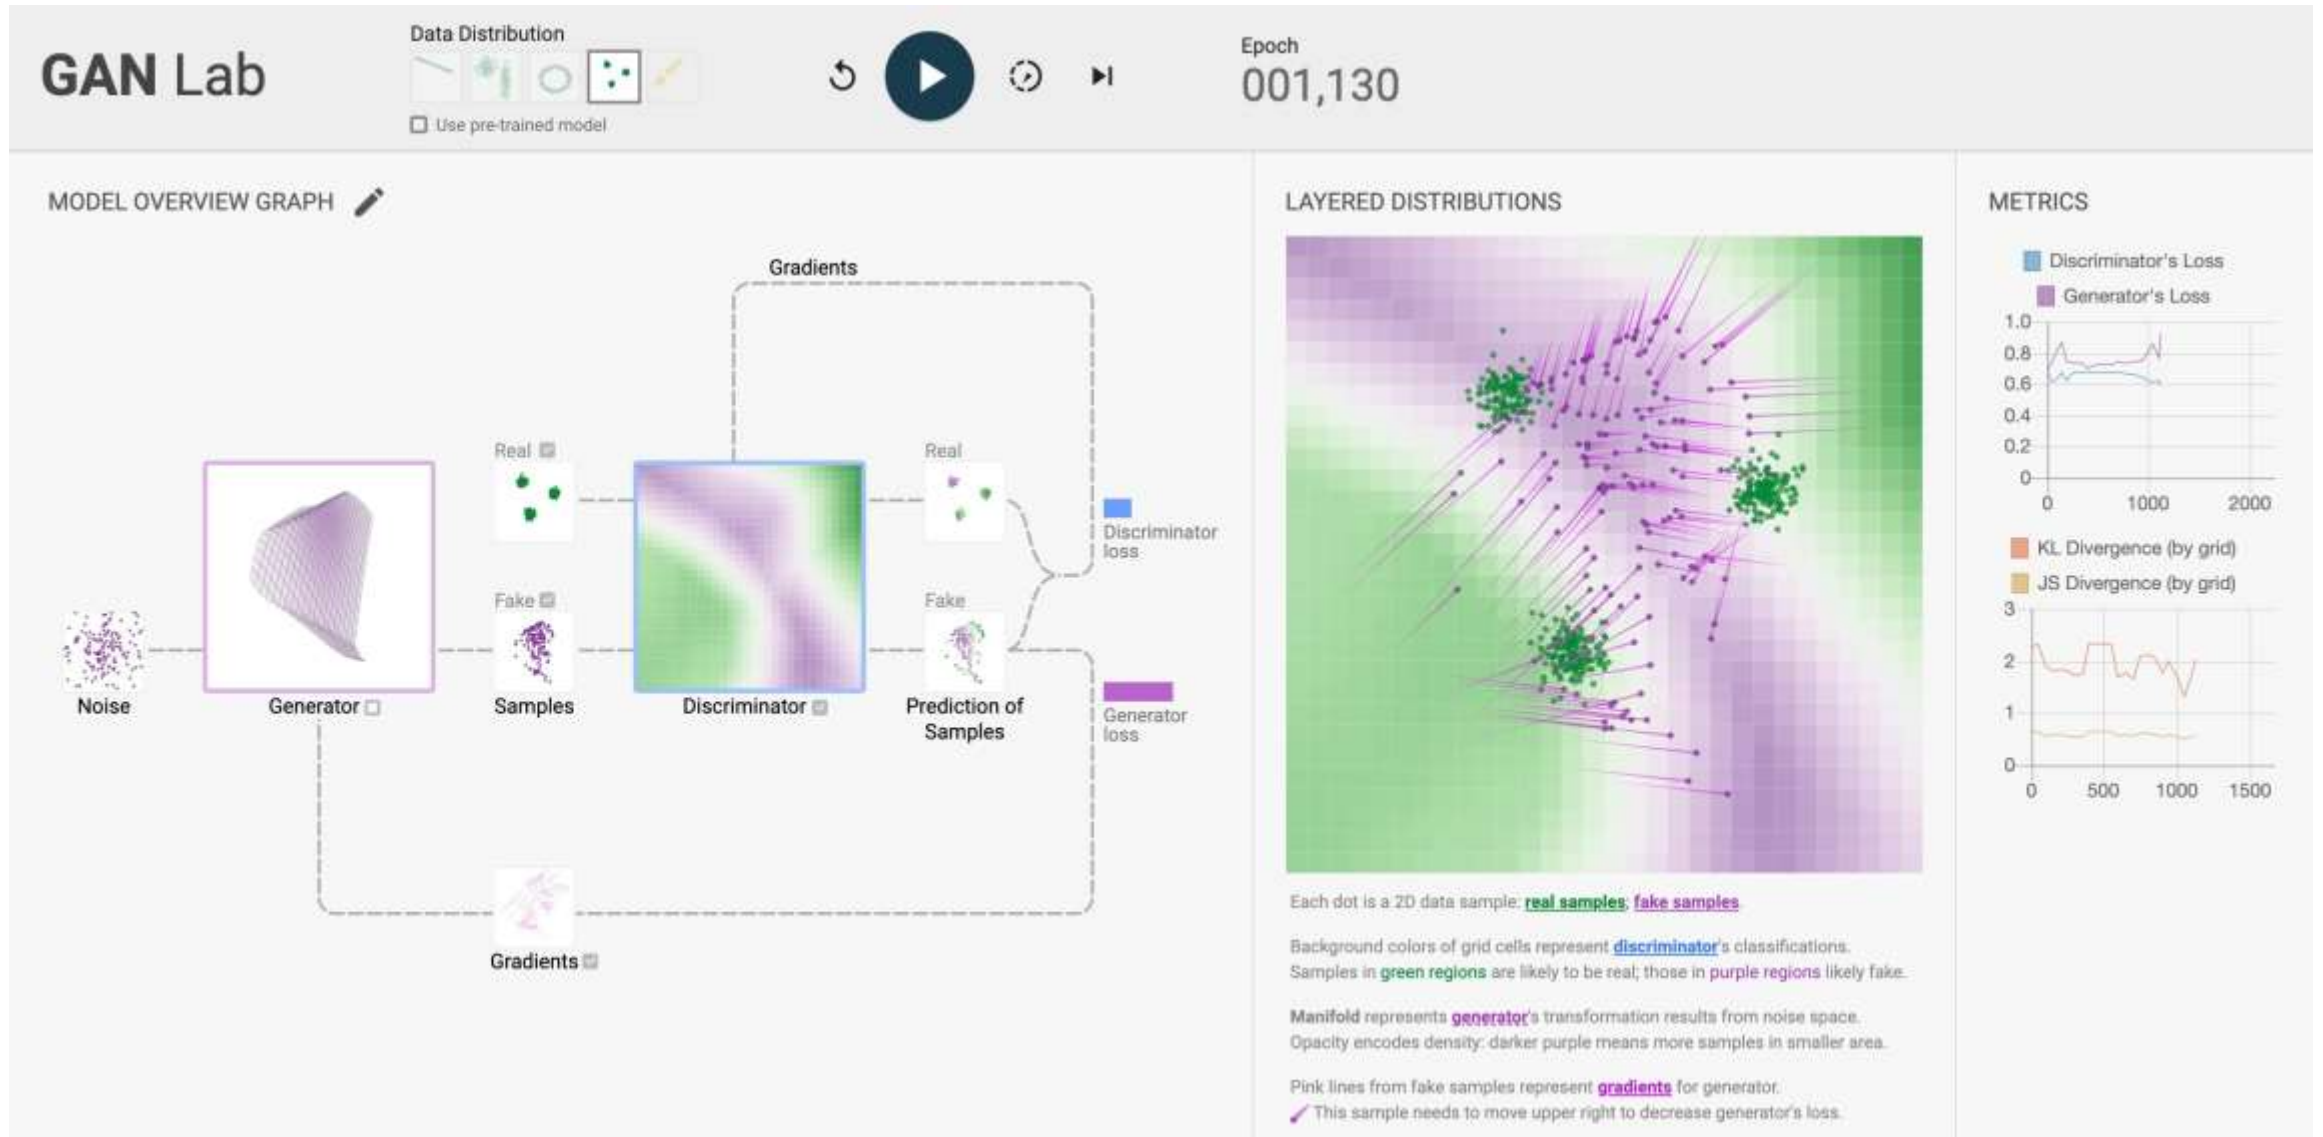
Running example: MNIST



*All objectives are negative of their original form

Code adapted from: <https://github.com/prcastro/pytorch-gan/tree/master>

Running example: GAN Lab <https://poloclub.github.io/ganlab/>



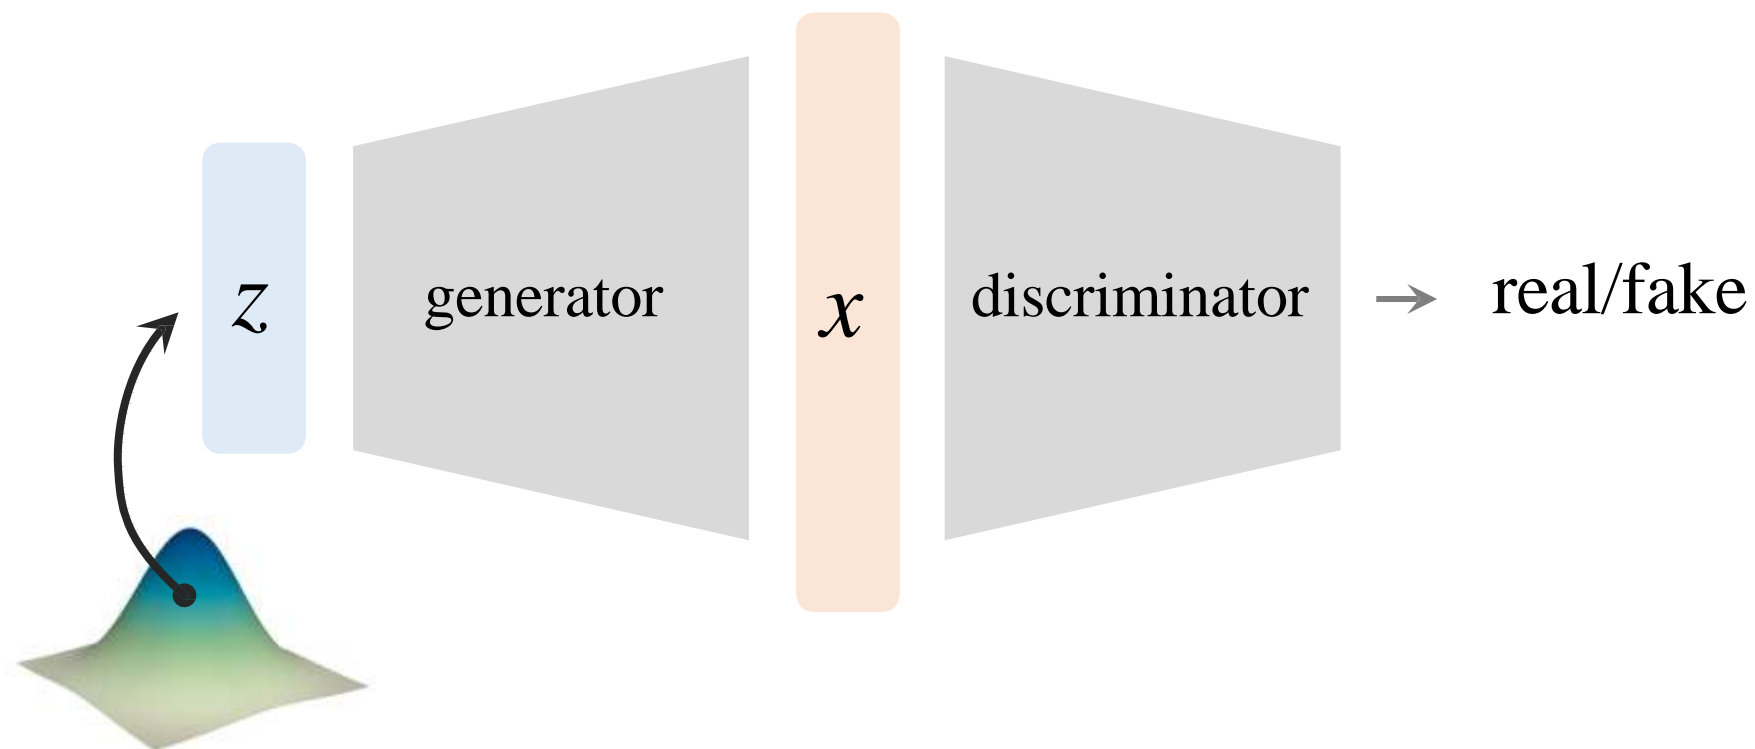
Adversary as a Loss Function

Adversary as a Loss Function

- GAN essentially defines an **adversarial loss** function
- Input to networks is **not** necessarily random/noise
- **Beyond L2/L1**: adversarial loss encourages output to look “realistic”
- **Combined with L2/L1**: reconstruction loss largely stabilizes training

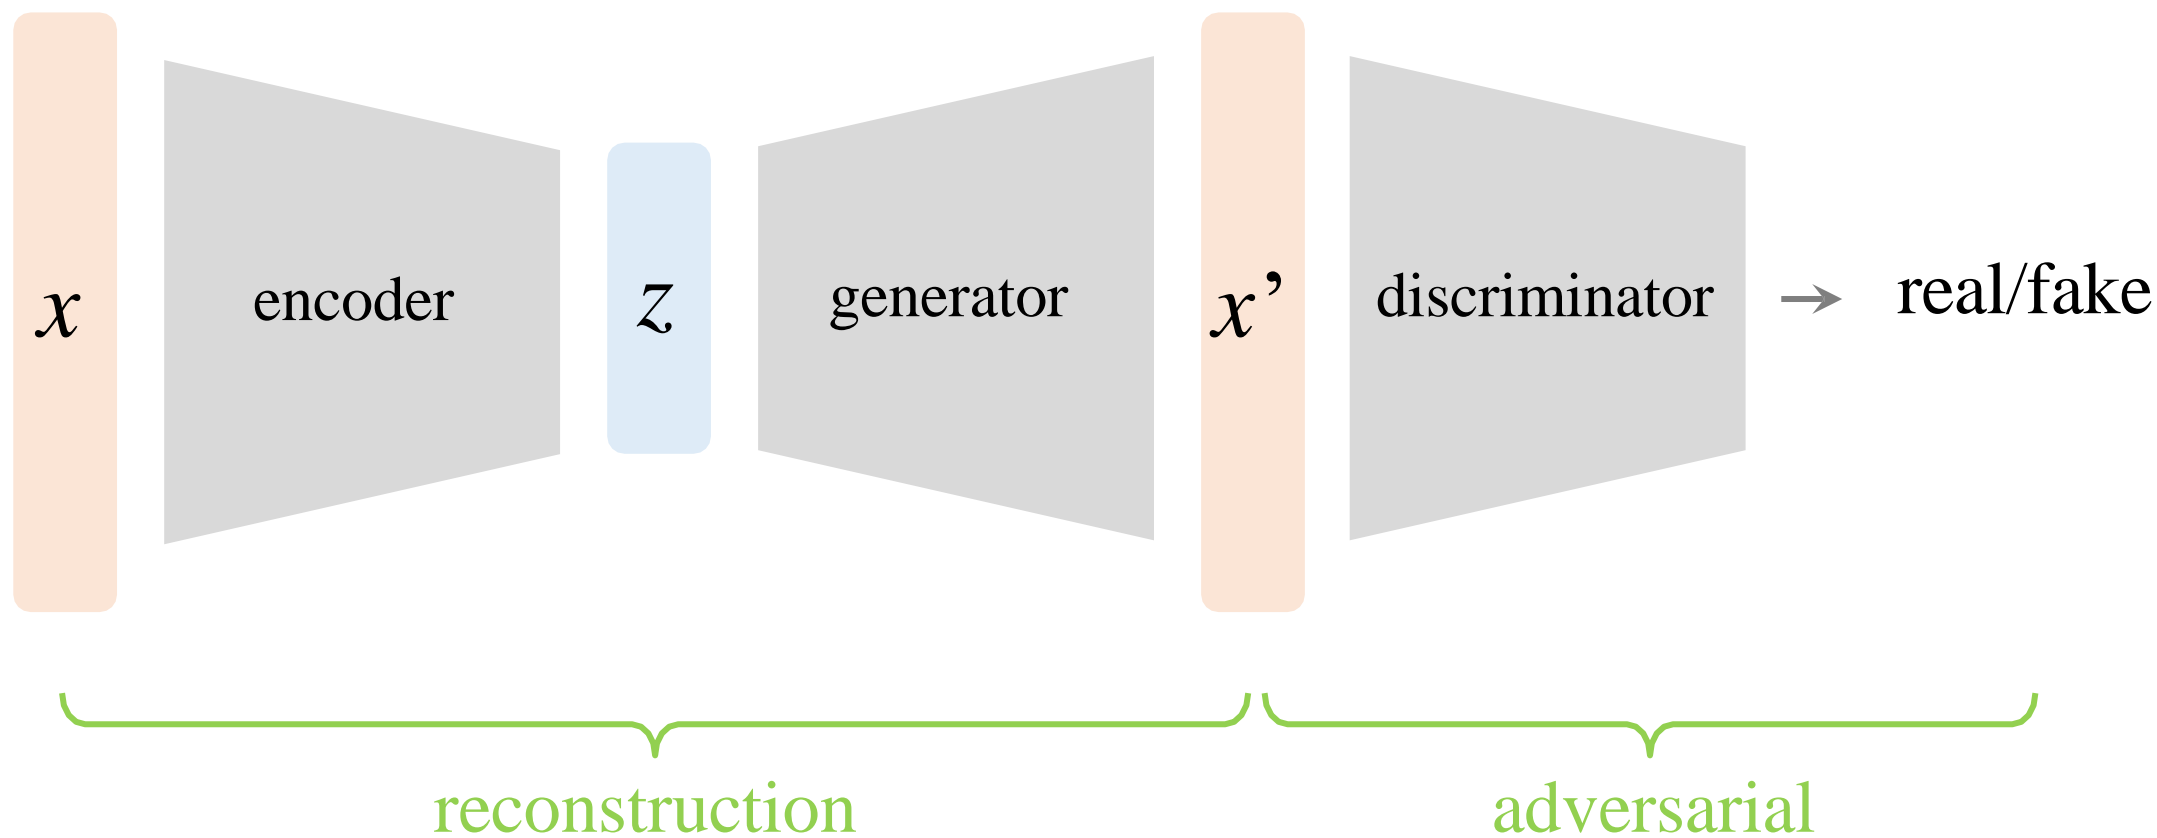
Adversary as a Loss Function

GAN: input is random



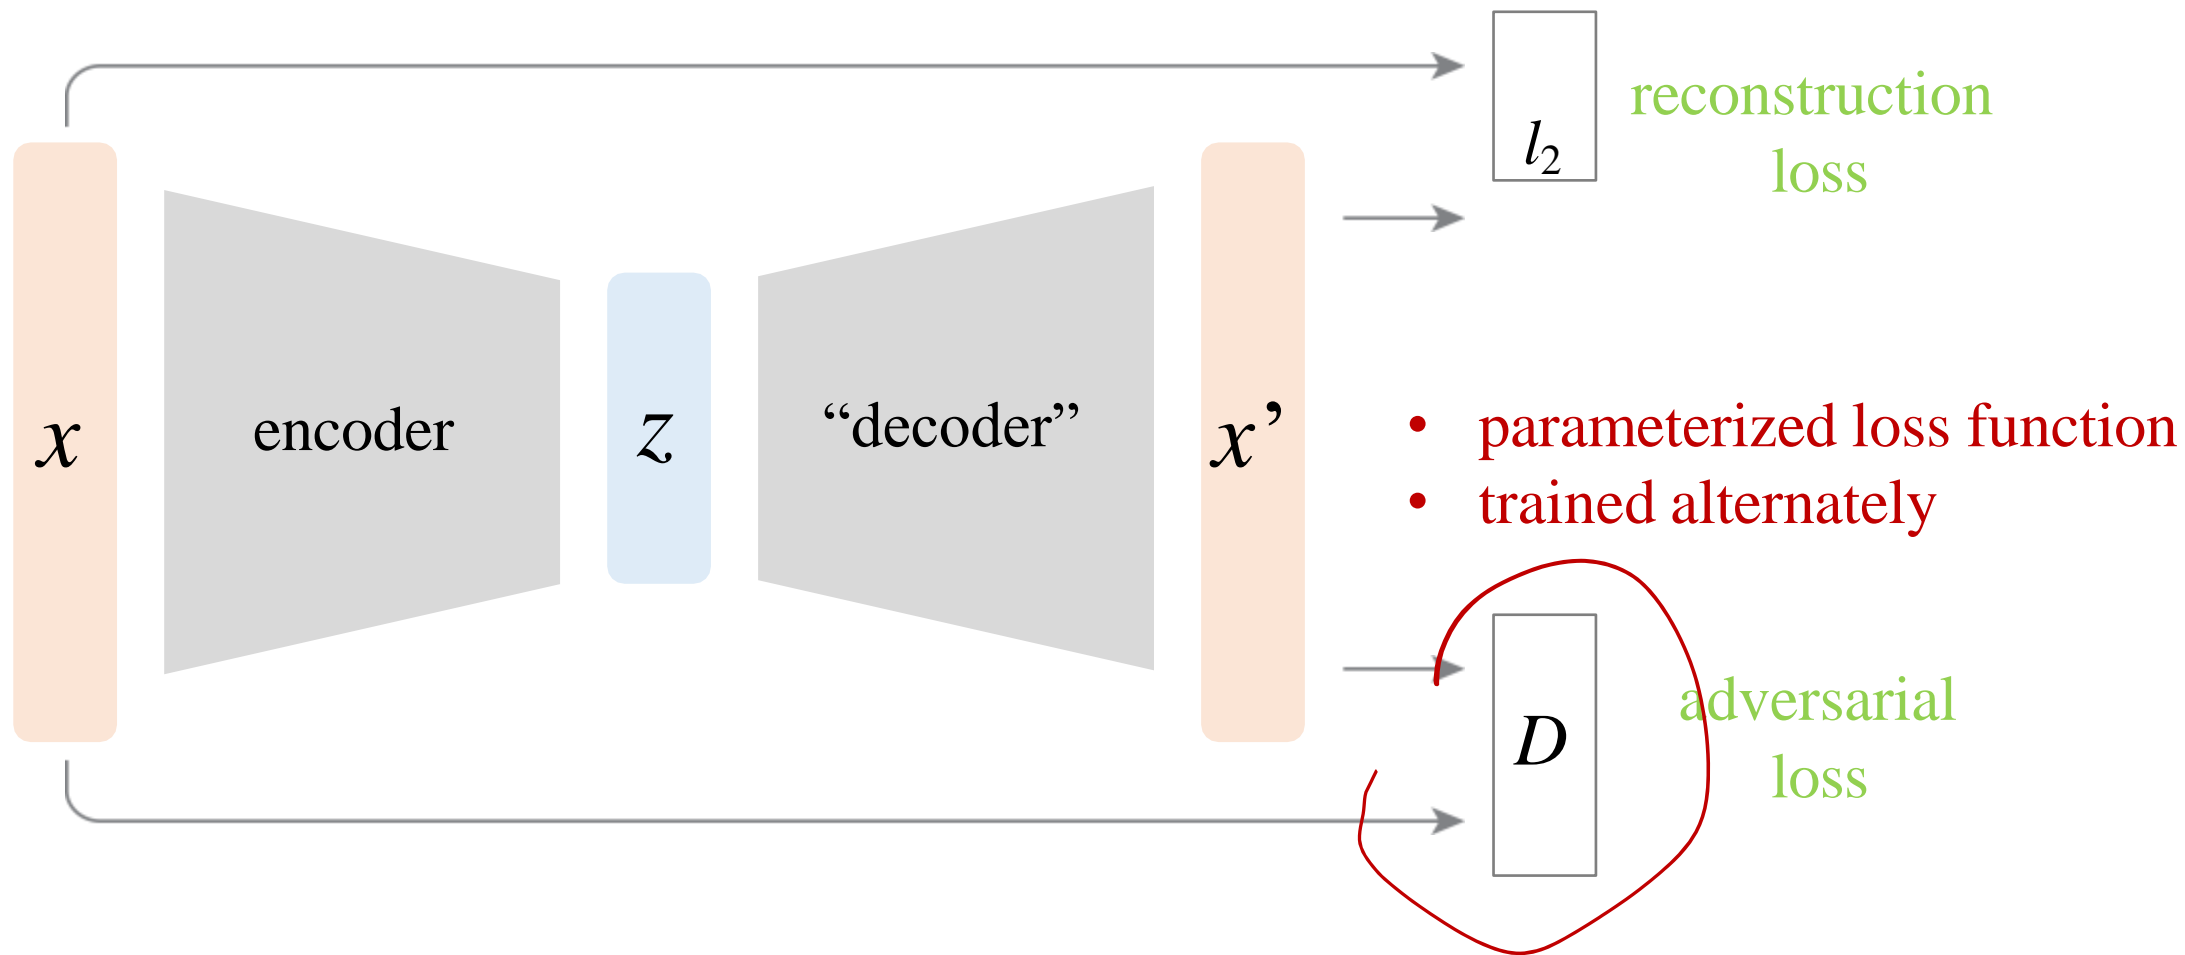
Adversary as a Loss Function

Input can be from another source



Adversary as a Loss Function

Input can be from another source



Example: Super-Resolution GAN

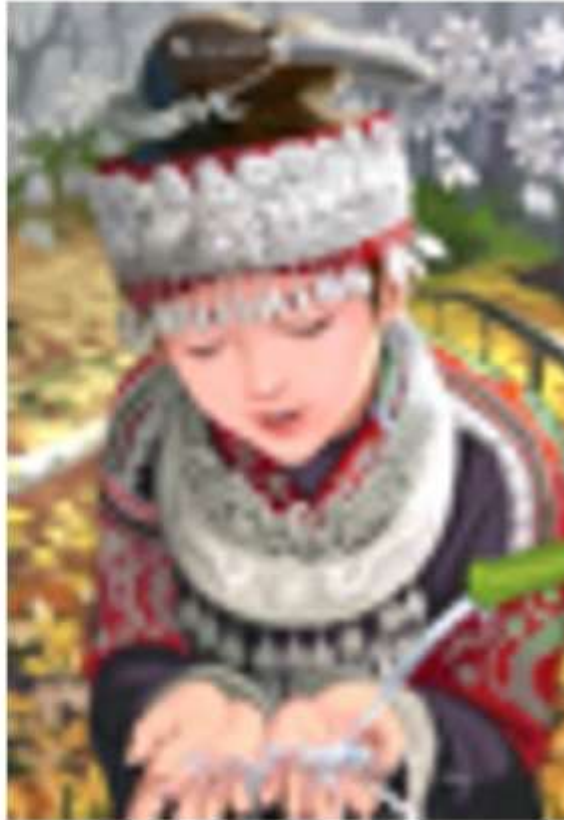
- better PSNR

- worse PSNR, but better visual quality

original



bicubic
(21.59dB/0.6423)



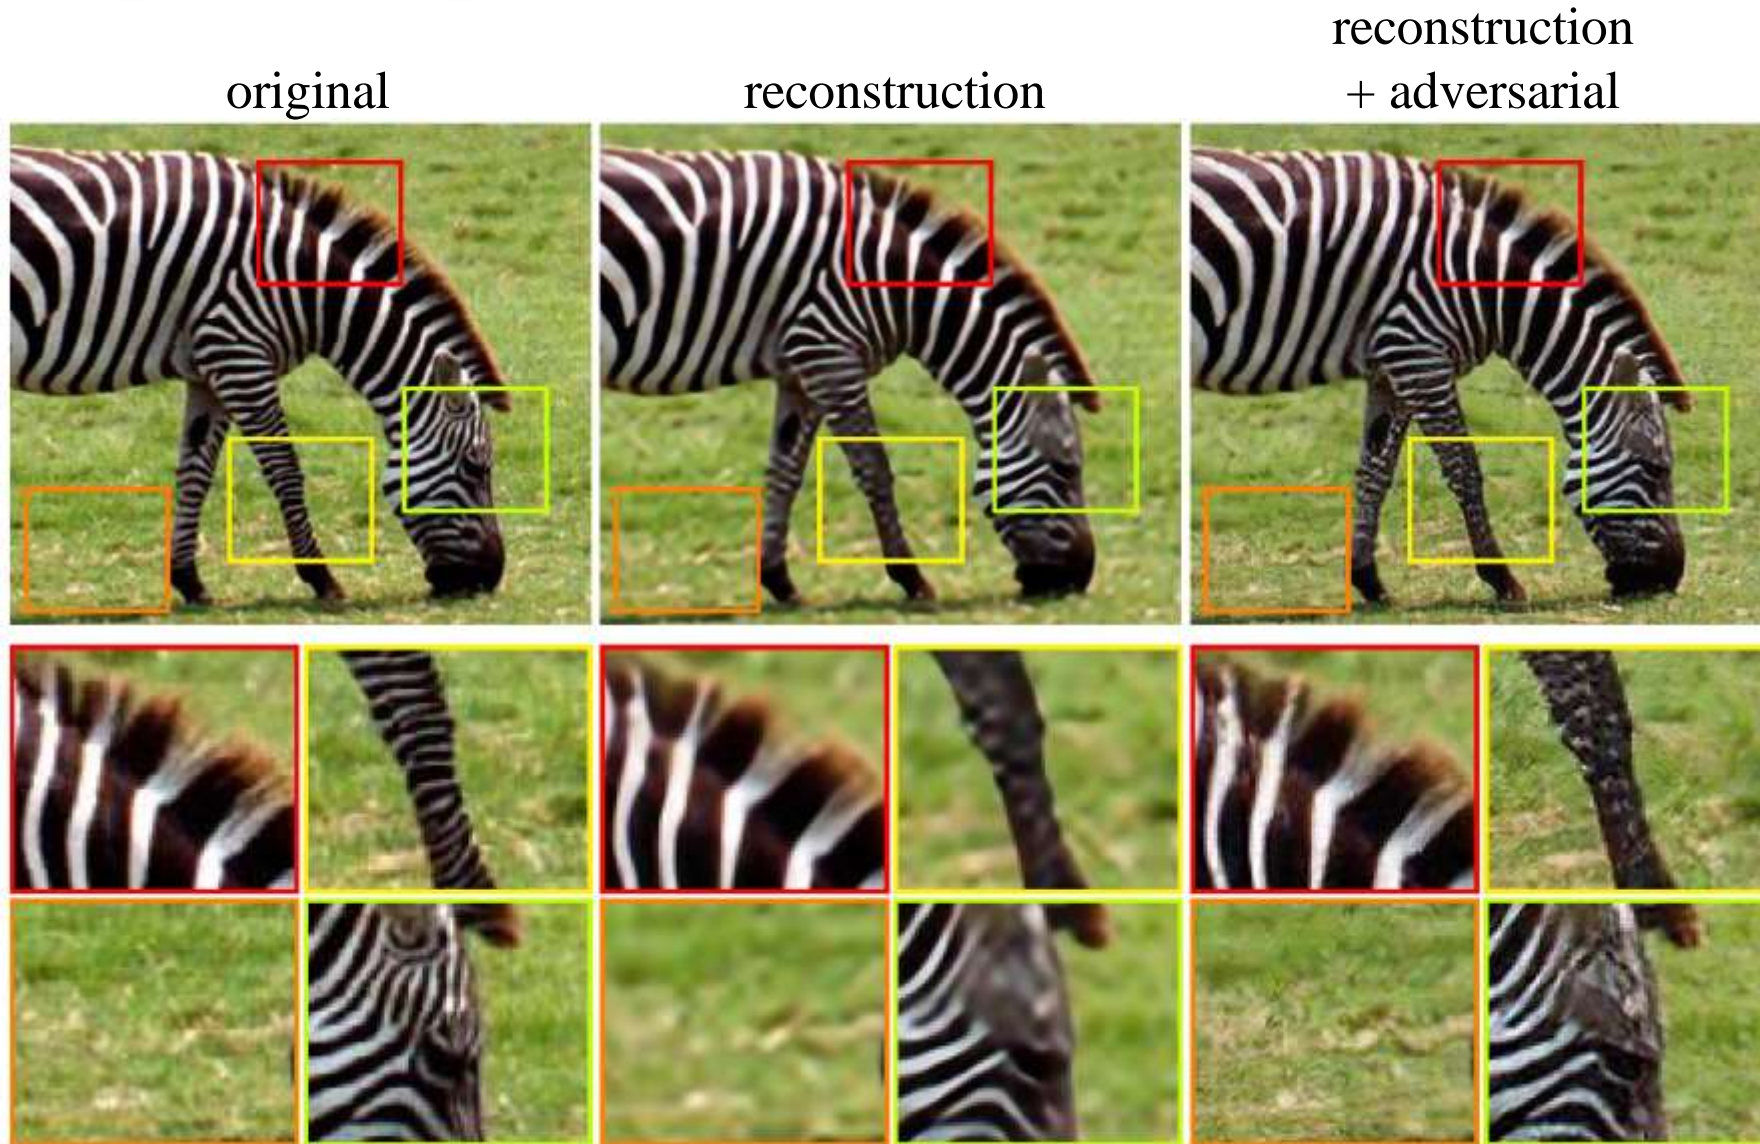
SRResNet
(23.44dB/0.7777)

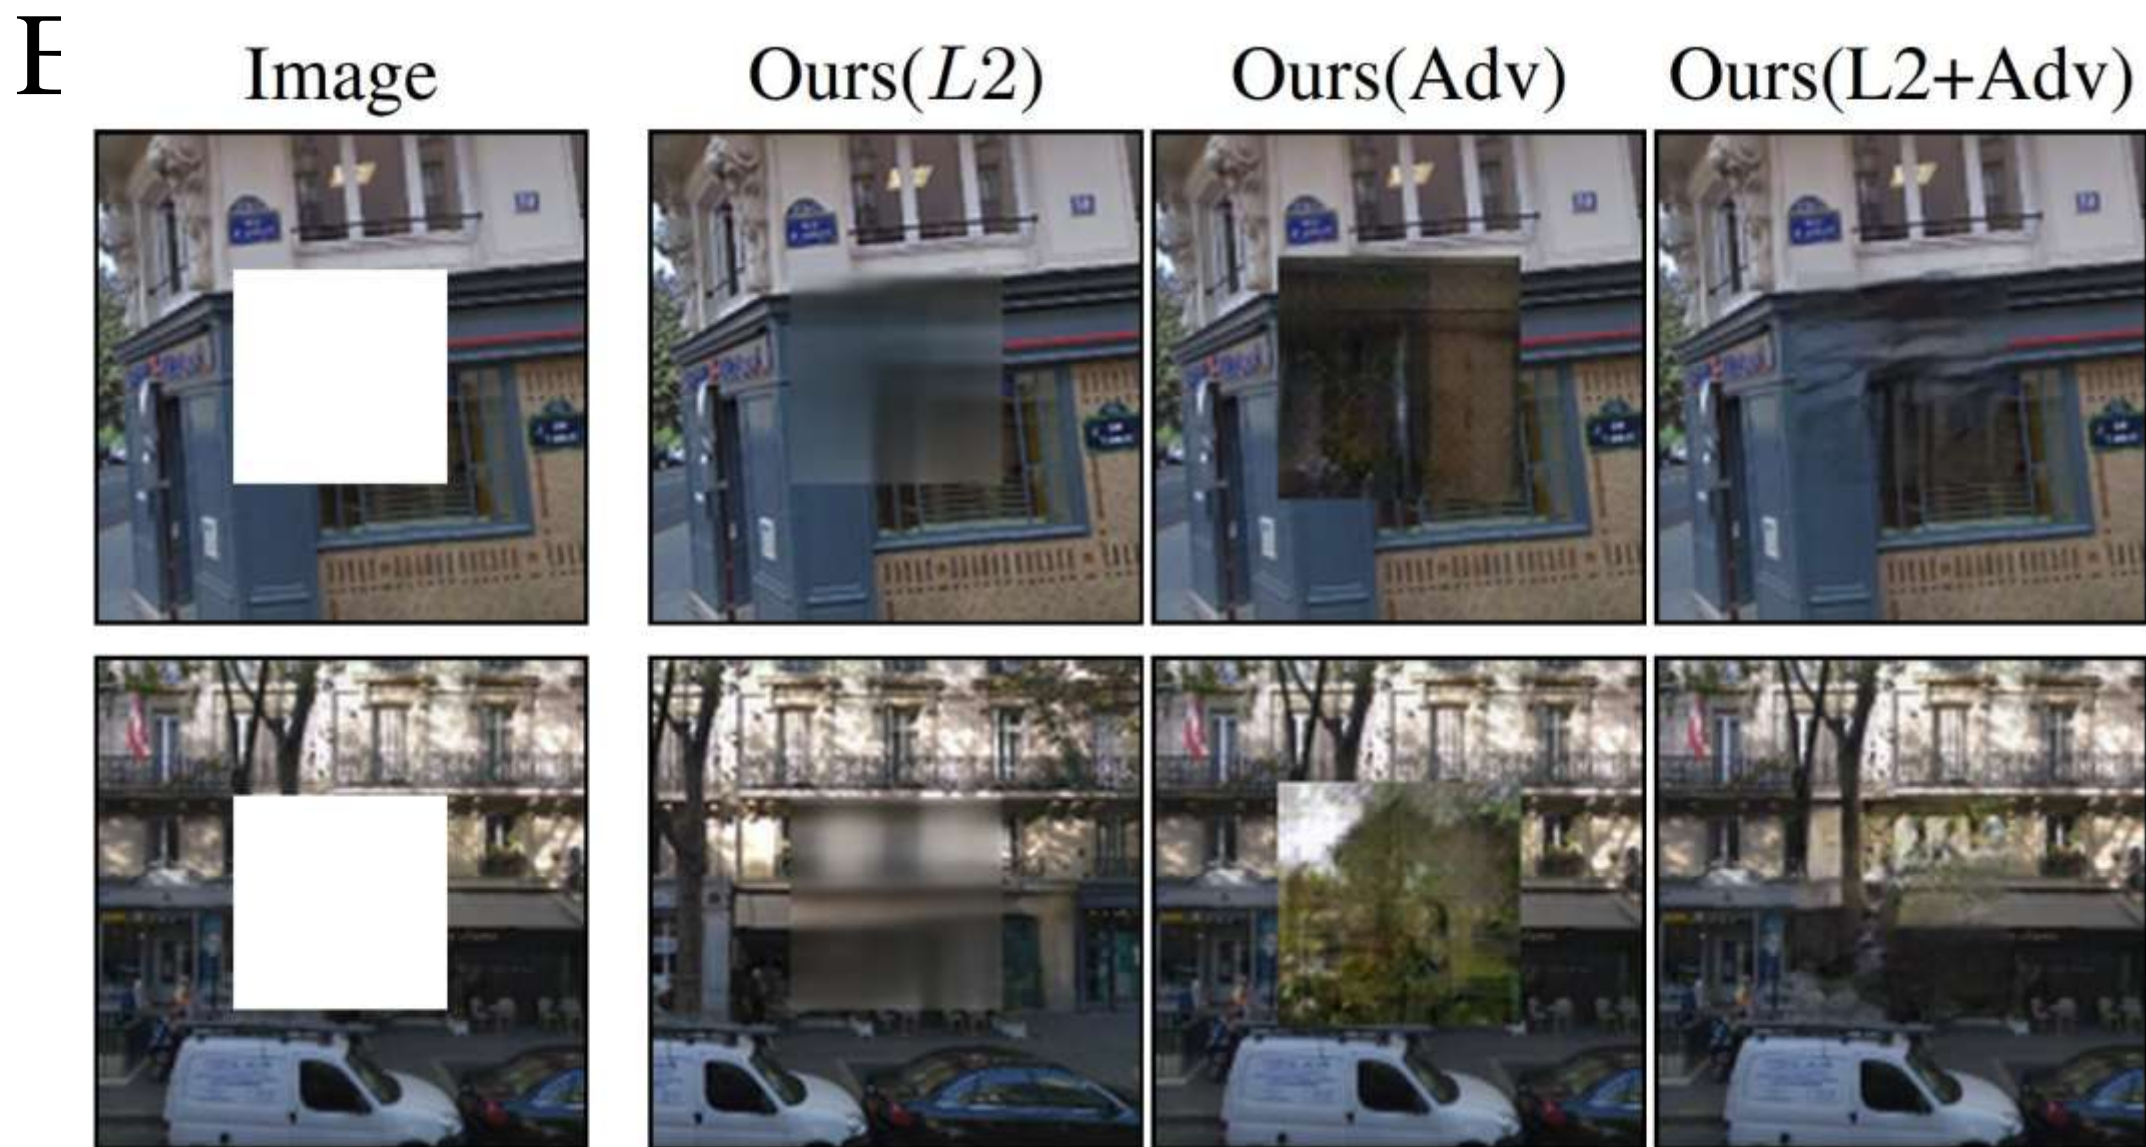


SRGAN
(20.34dB/0.6562)



Example: Super-Resolution GAN





Example: nix2nix



Example: CycleGAN




zebra → horse




horse → zebra

\mathbf{x}



G



Generator

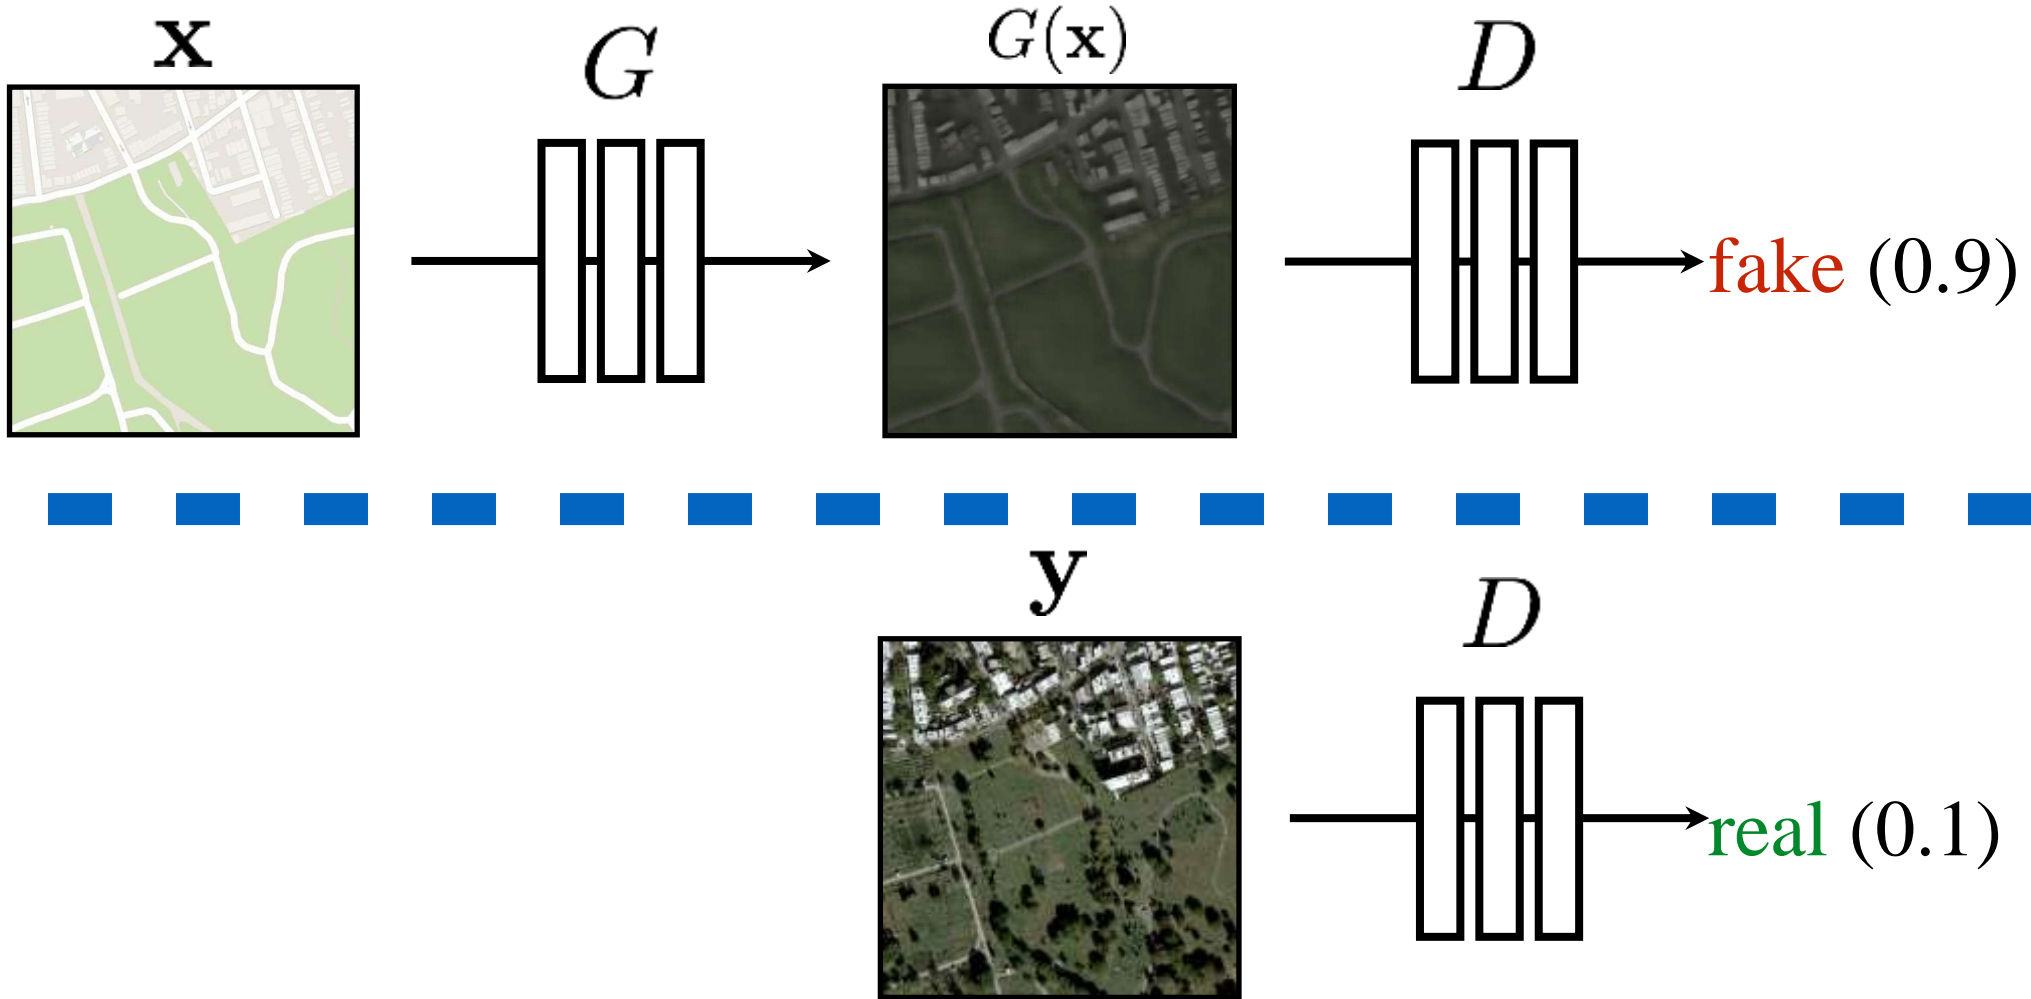
$G(\mathbf{x})$



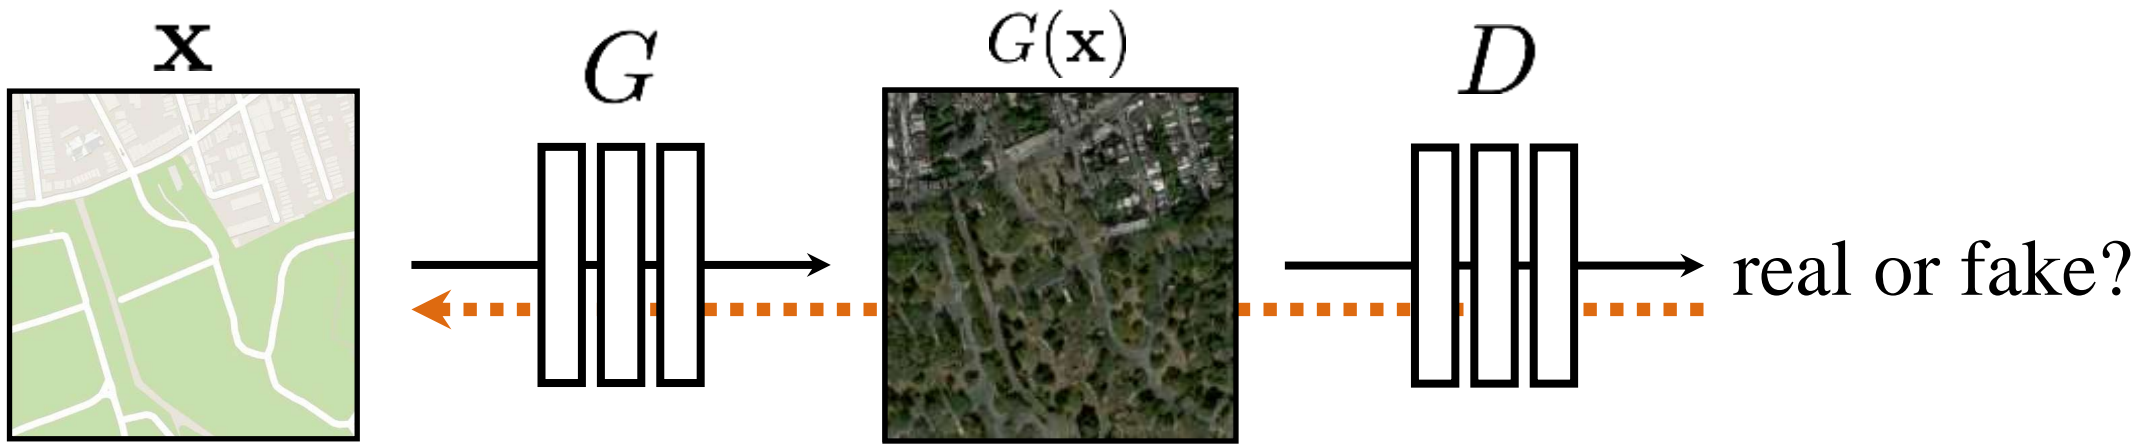


G tries to synthesize fake images that fool D

D tries to identify the fakes

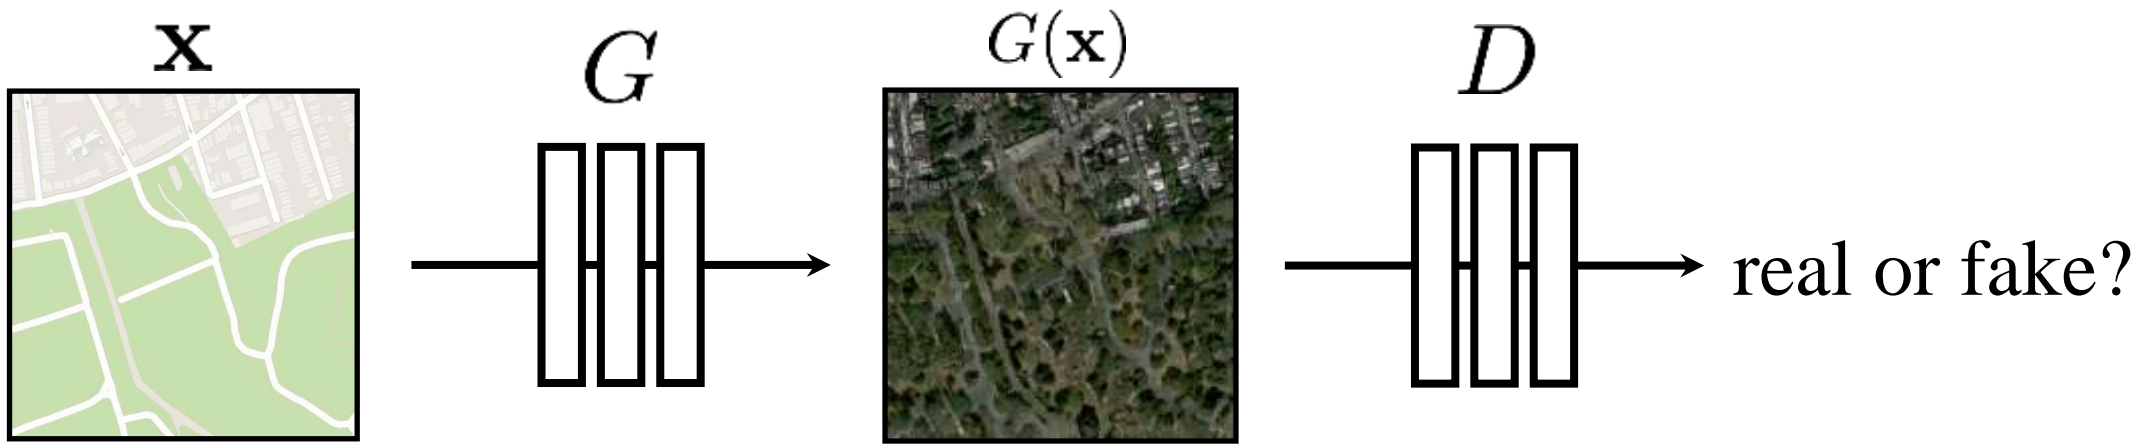


$$\arg \max_D \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\log D(G(\mathbf{x})) + \log(1 - D(\mathbf{y}))]$$



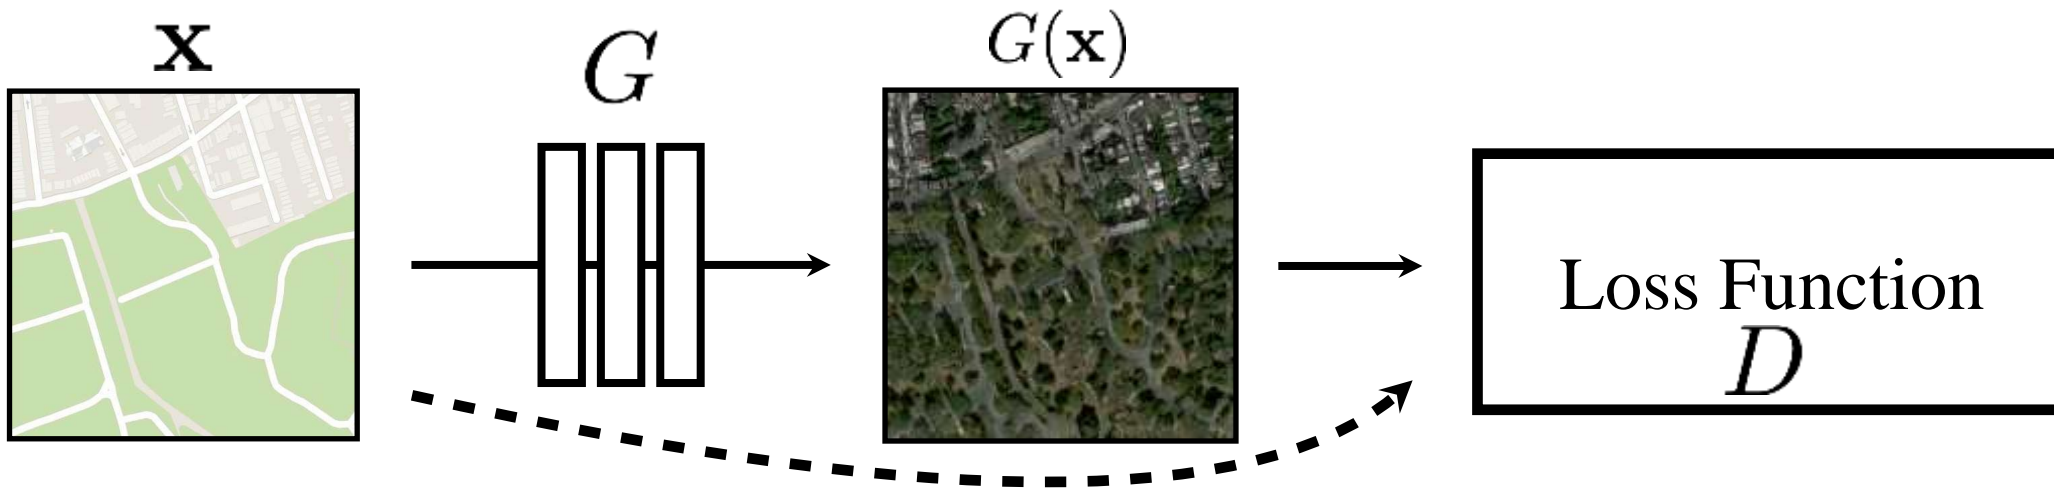
G tries to synthesize fake images that *fool* D :

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



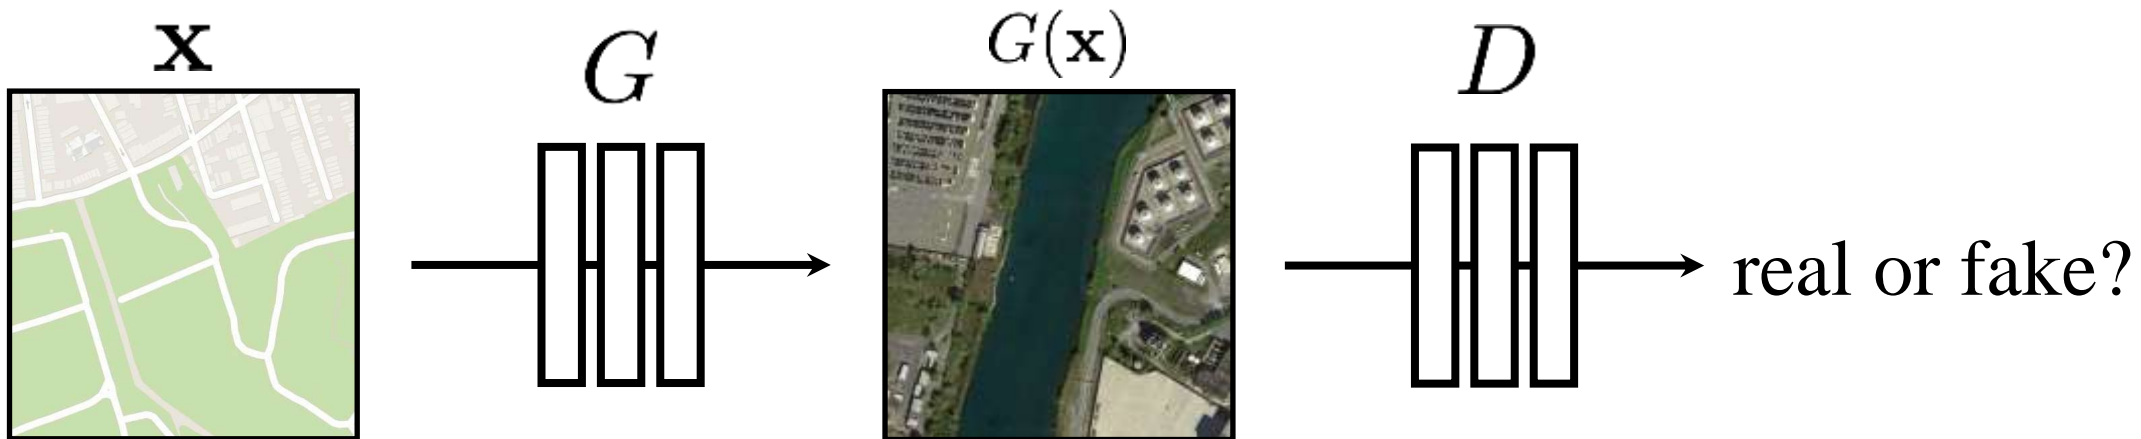
G tries to synthesize fake images that *fool* the *best* D:

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

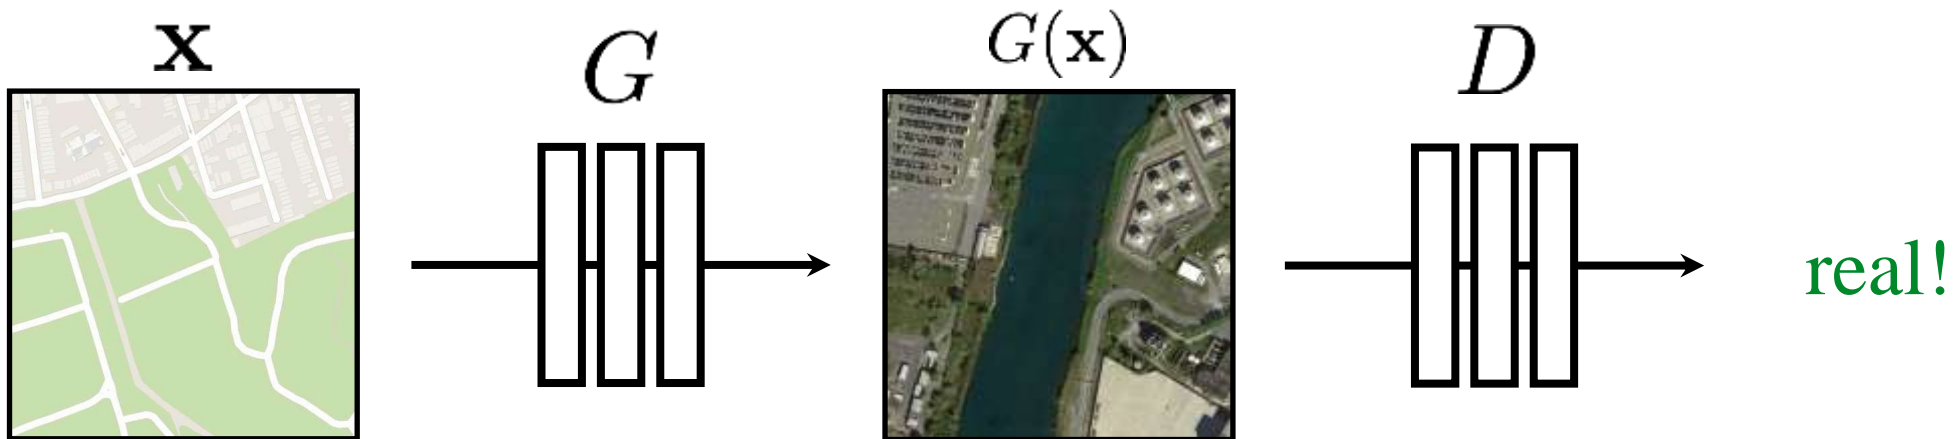


G's perspective: D is a loss function.

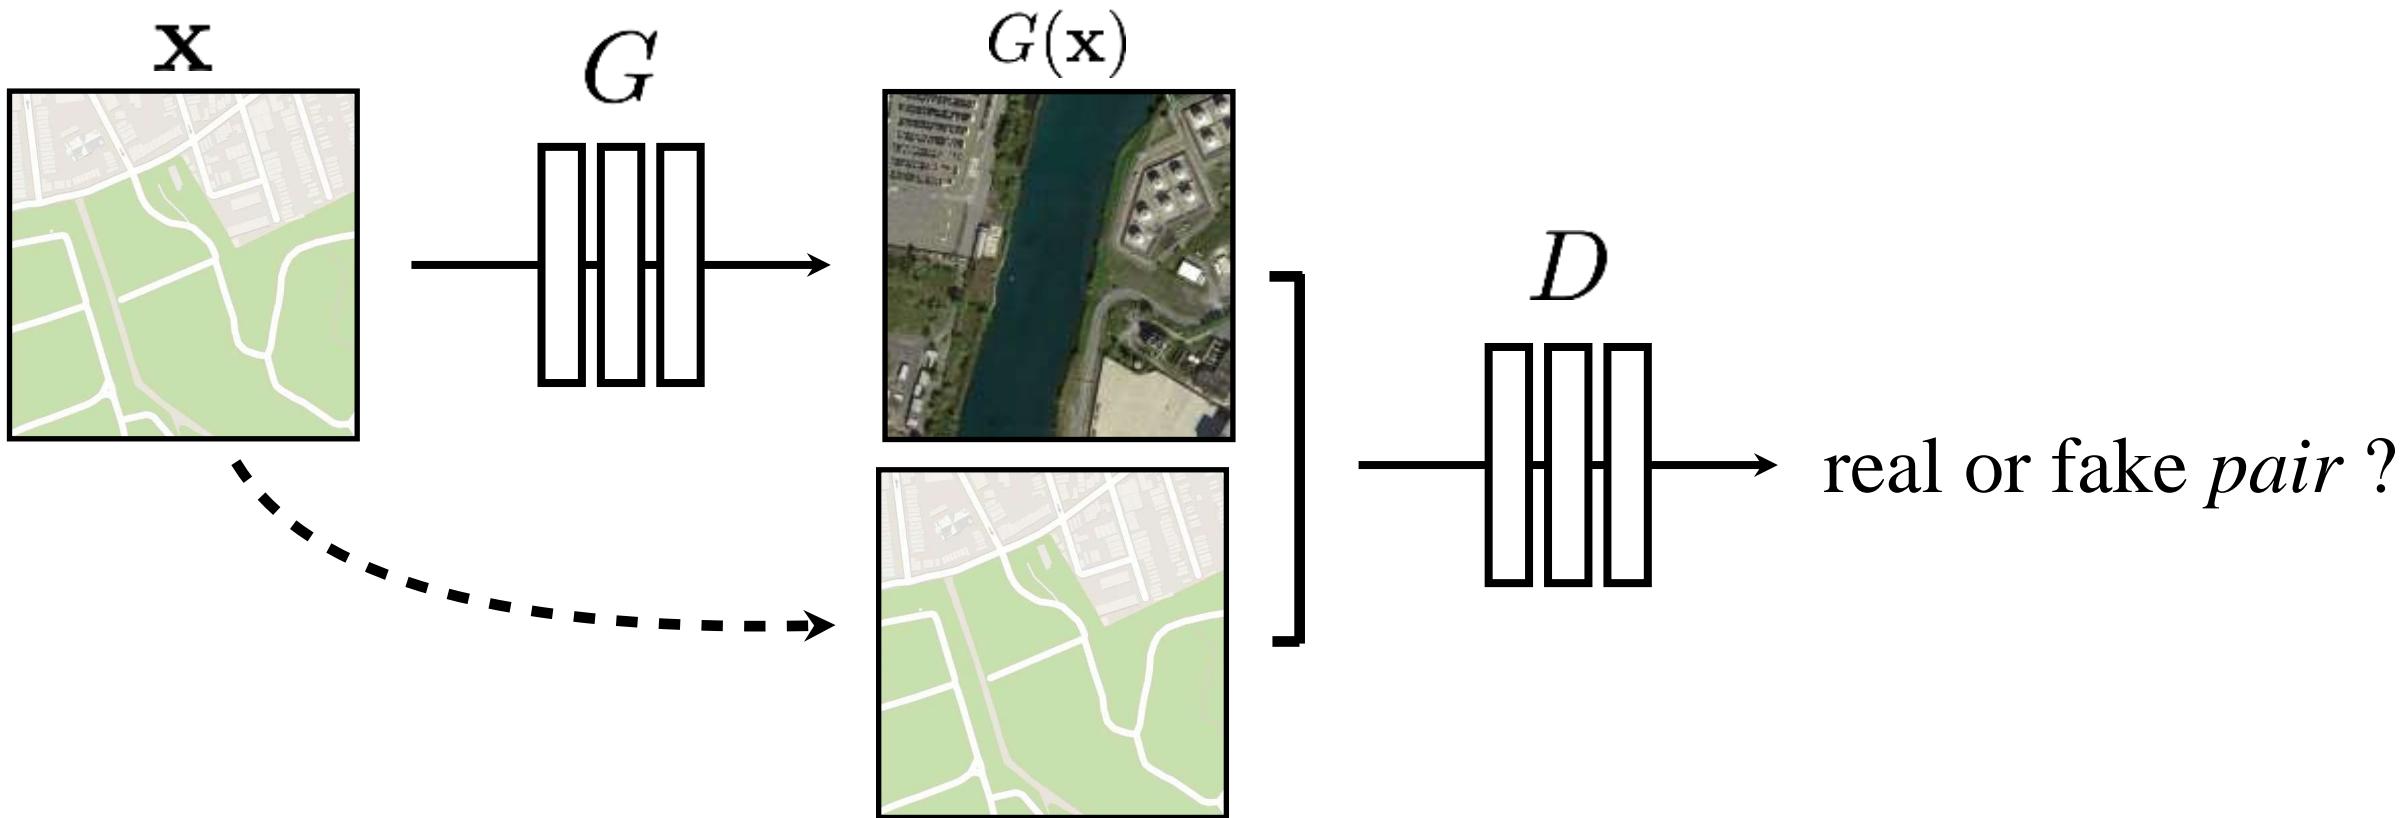
Rather than being hand-designed, it is *learned* and *highly structured*.



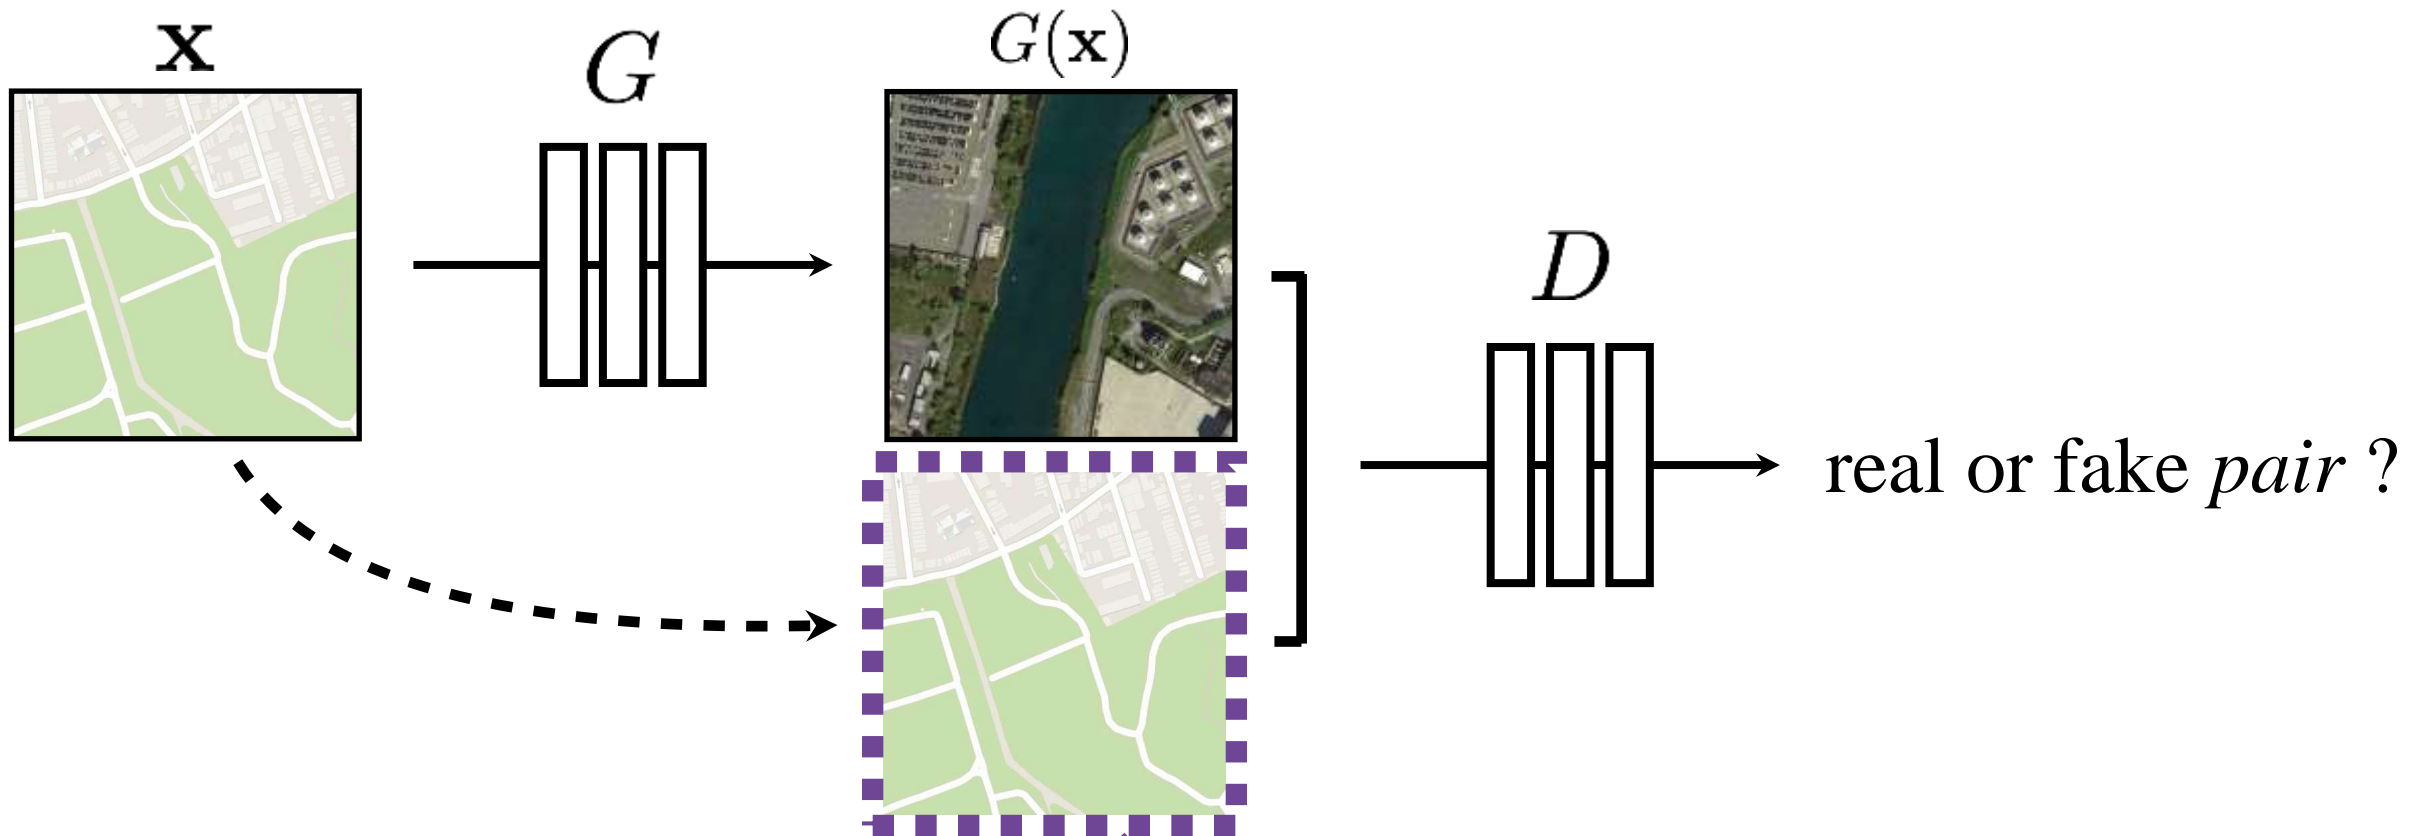
$$\arg \min_G \max_D \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\log D(G(\mathbf{x})) + \log(1 - D(\mathbf{y}))]$$



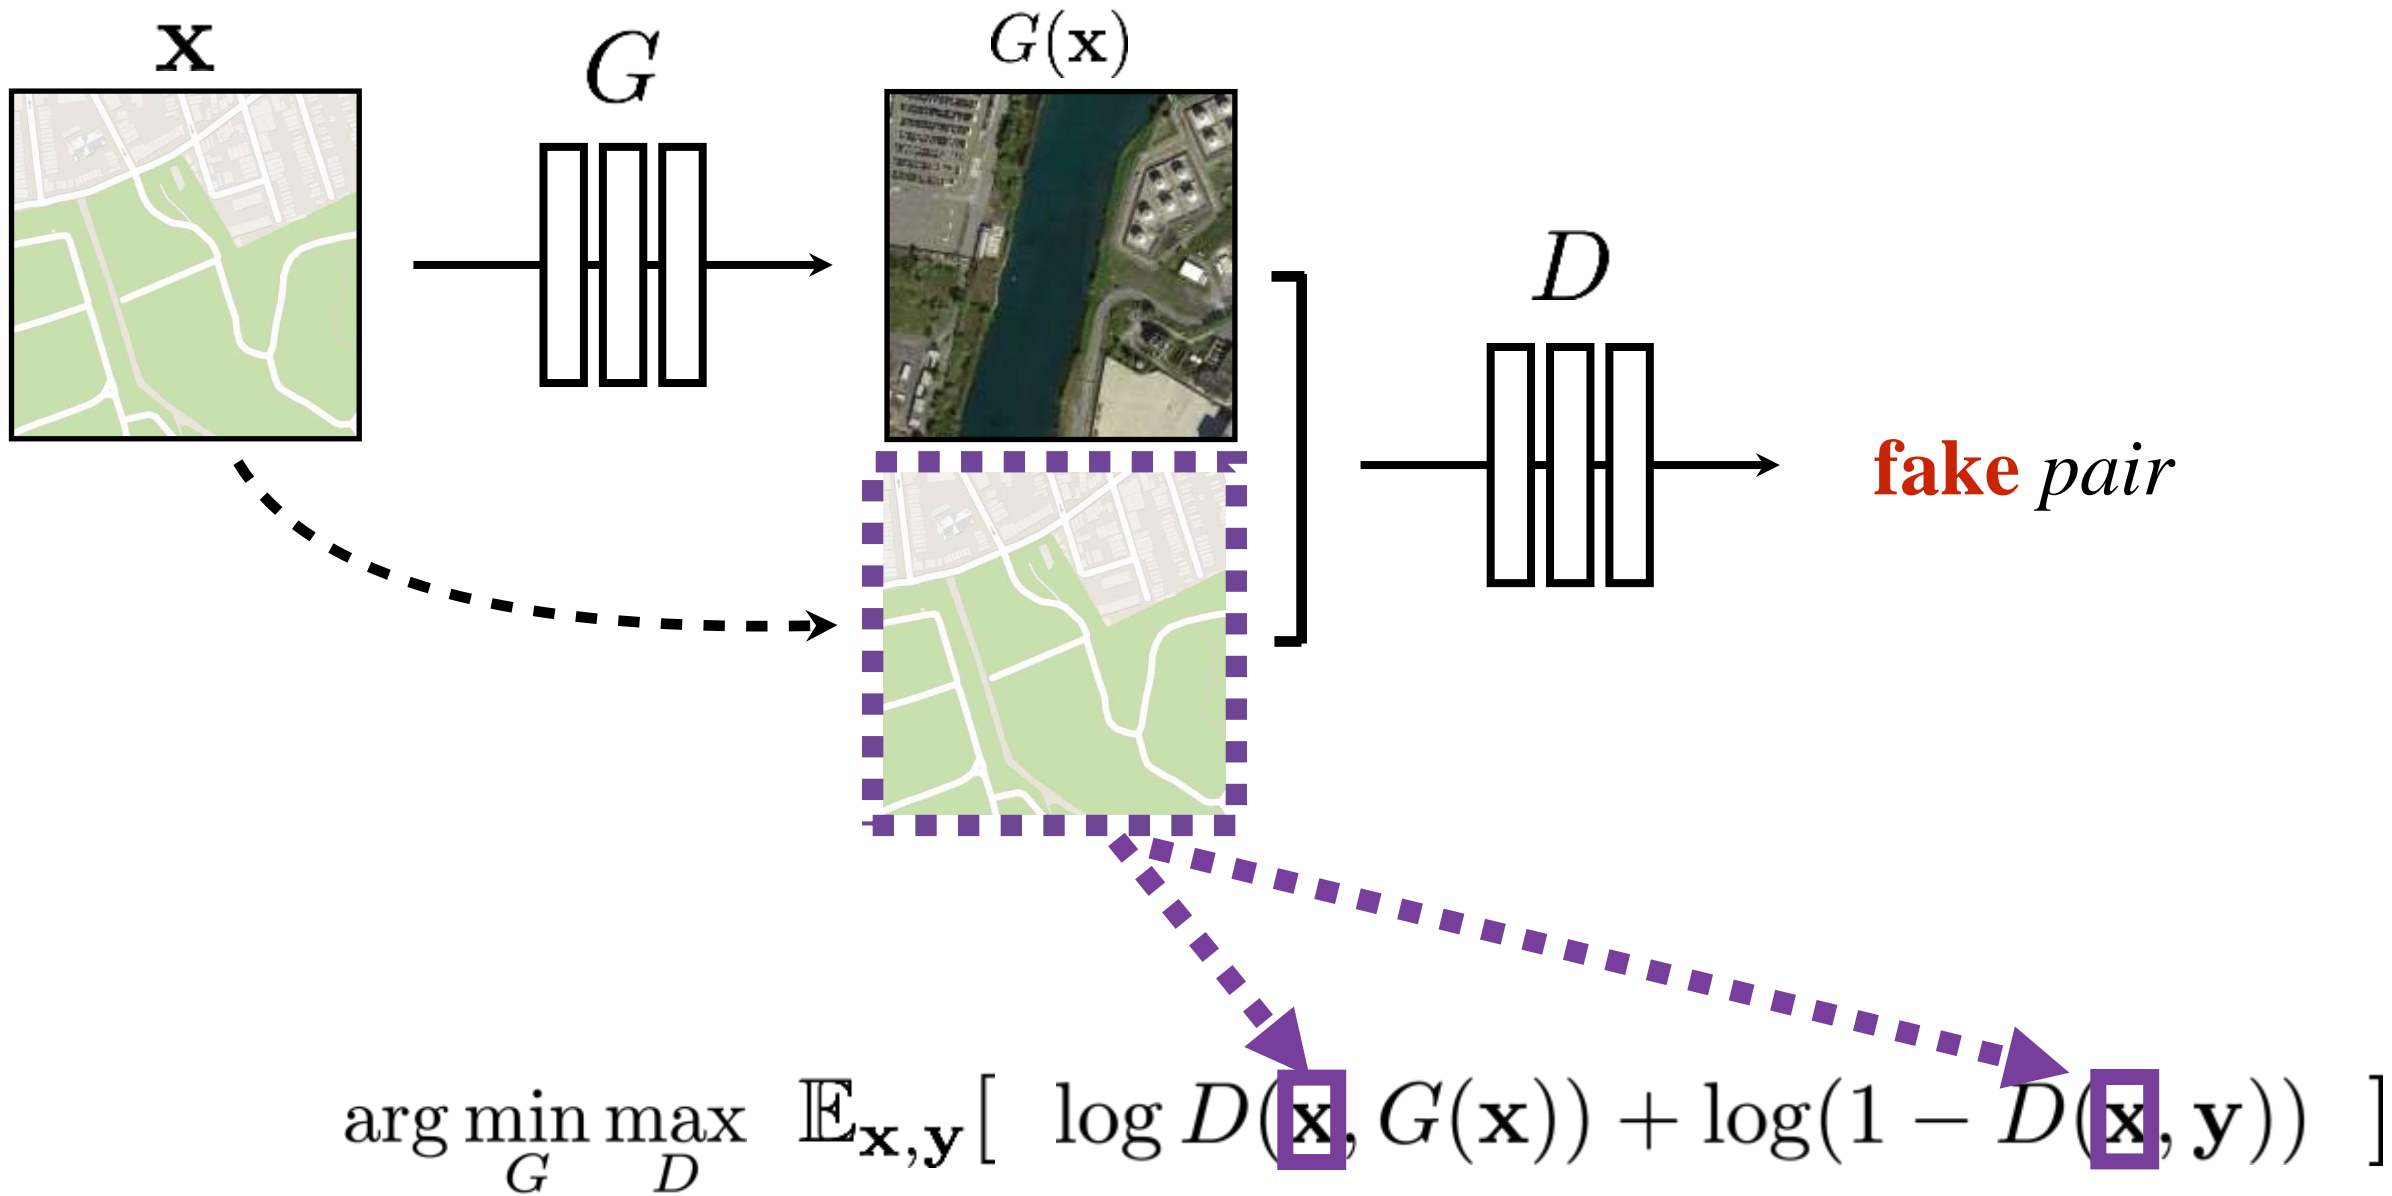
$$\arg \min_G \max_D \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\log D(G(\mathbf{x})) + \log(1 - D(\mathbf{y}))]$$

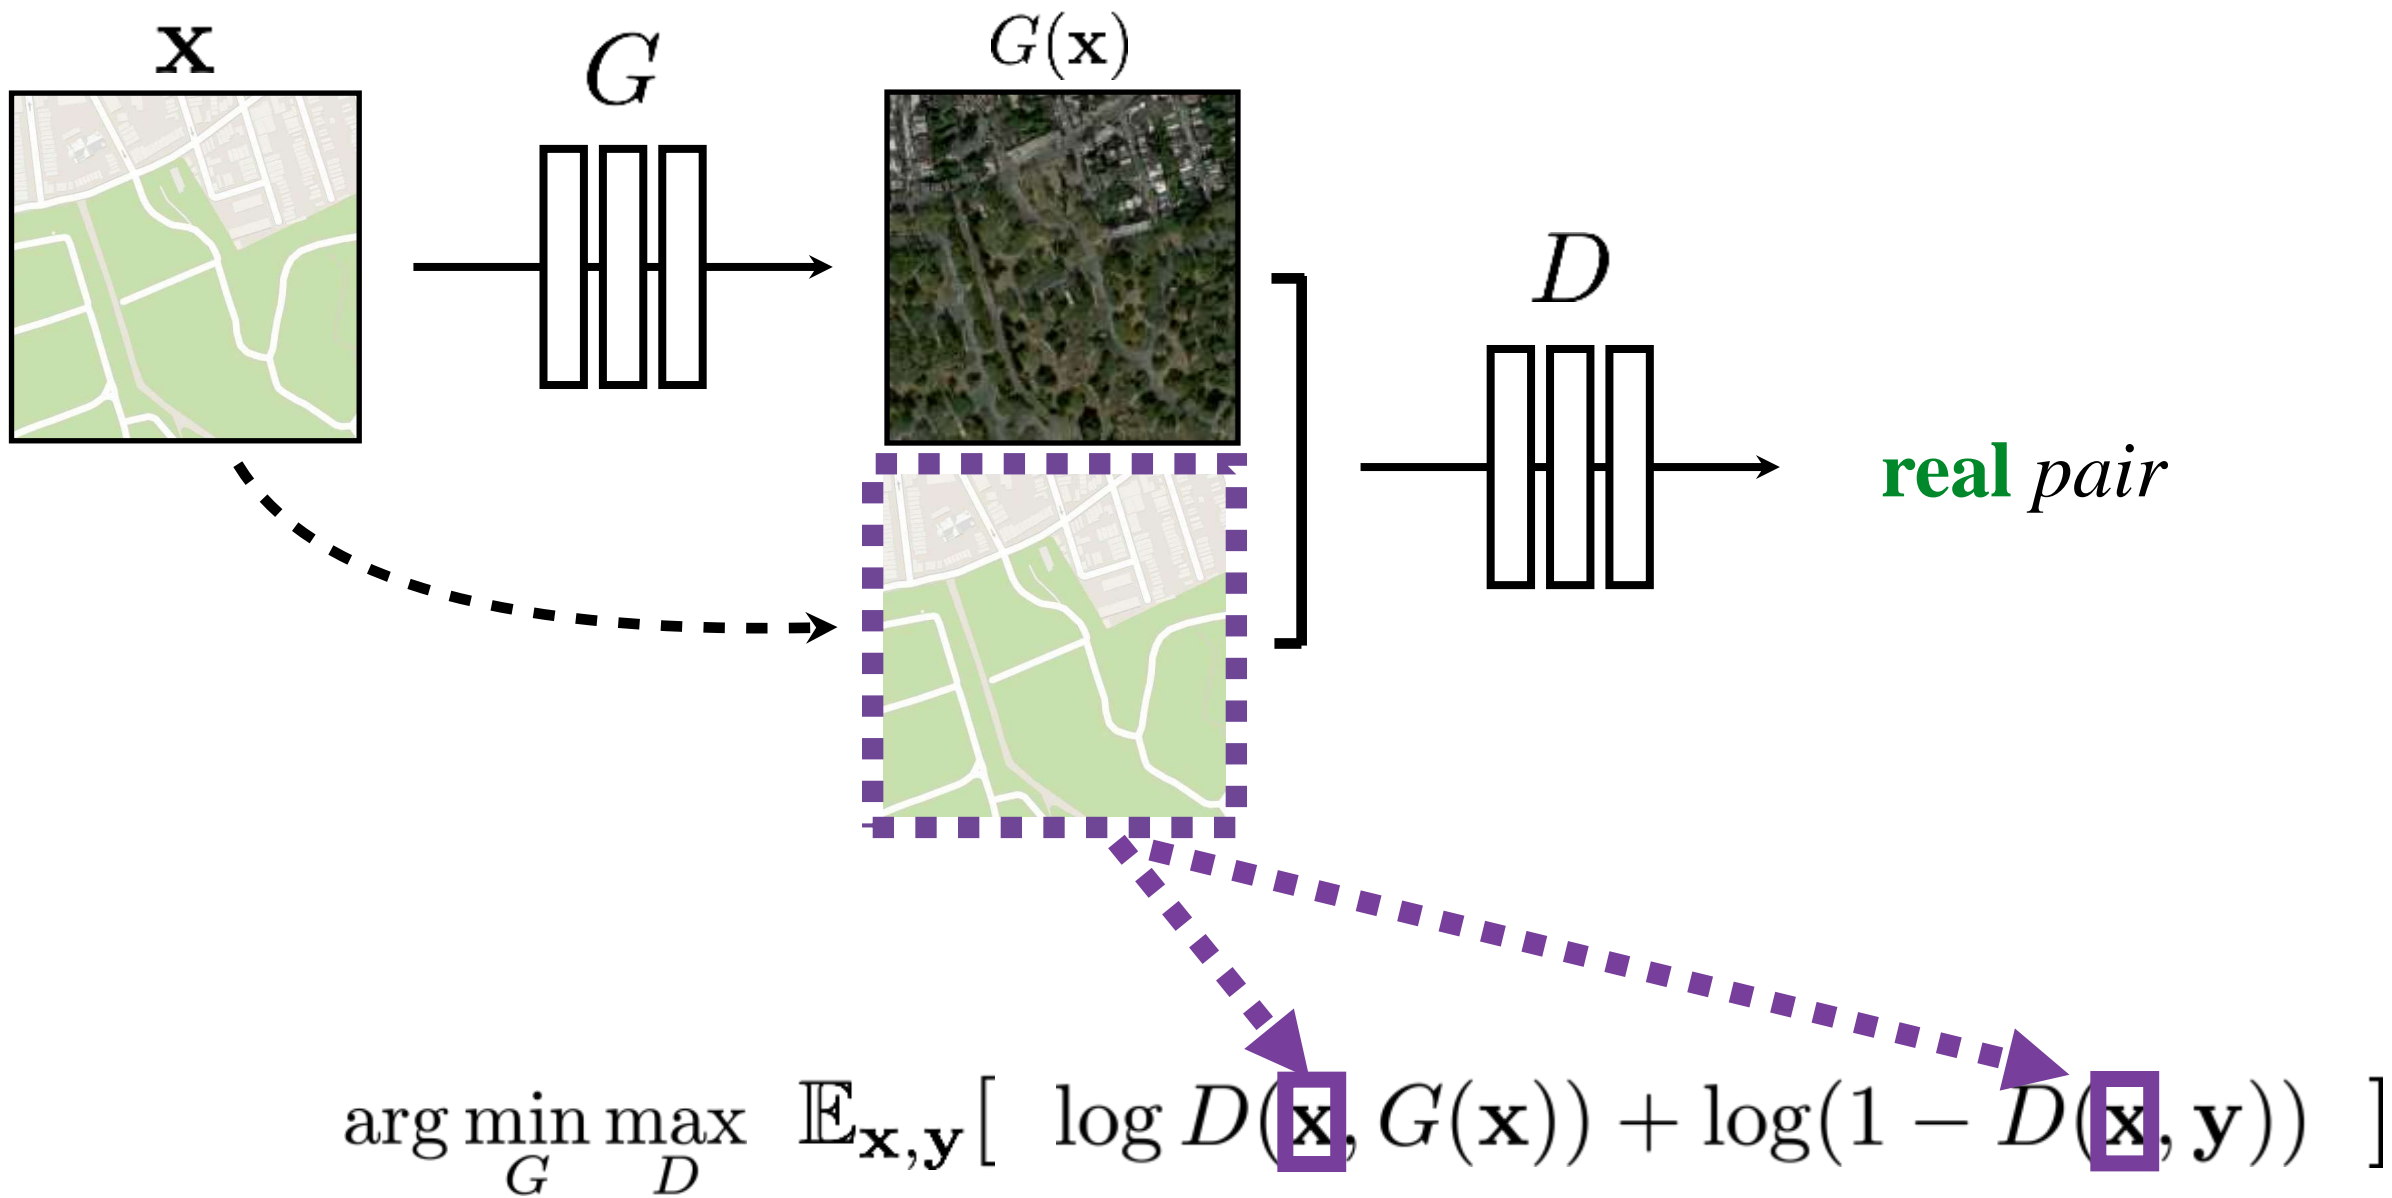


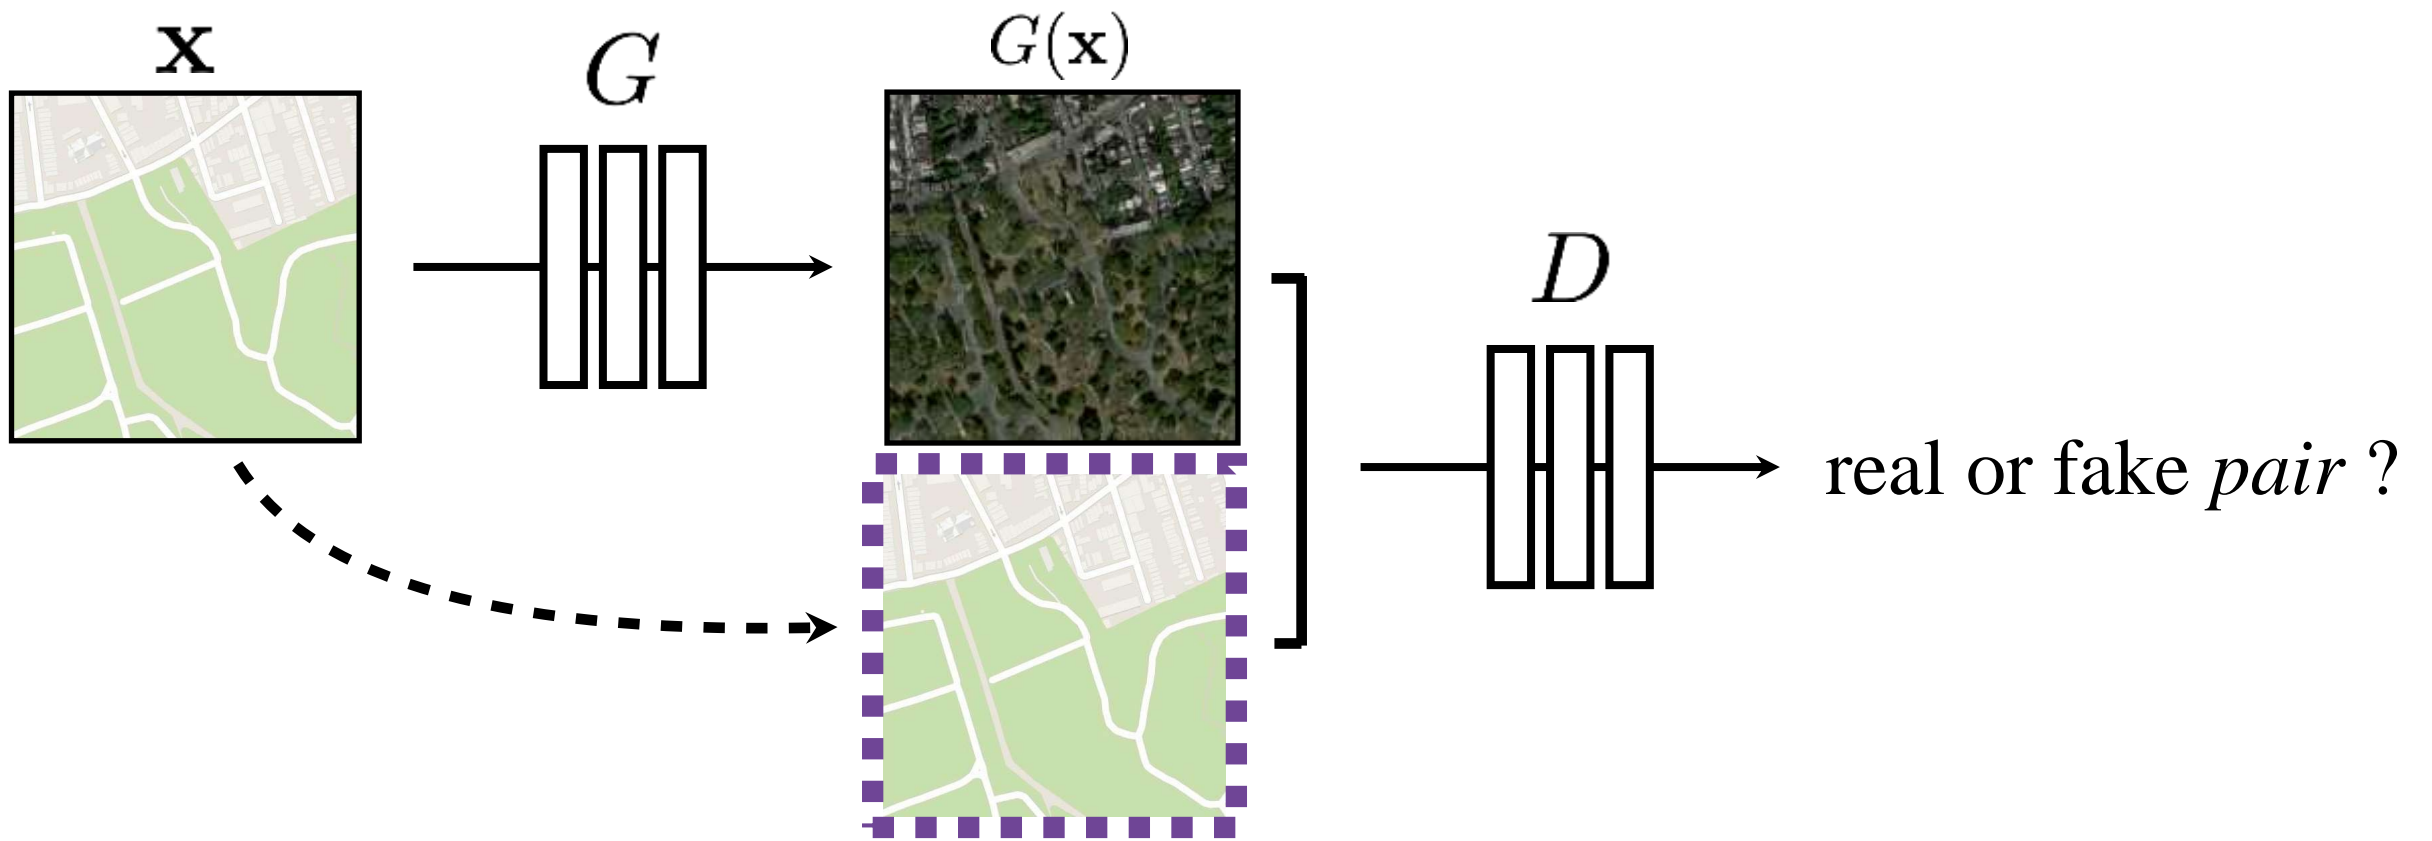
$$\arg \min_G \max_D \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\log D(G(\mathbf{x})) + \log(1 - D(\mathbf{y}))]$$



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








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

1. Image synthesis
2. Structured prediction
3. Domain mapping



Domain mapping

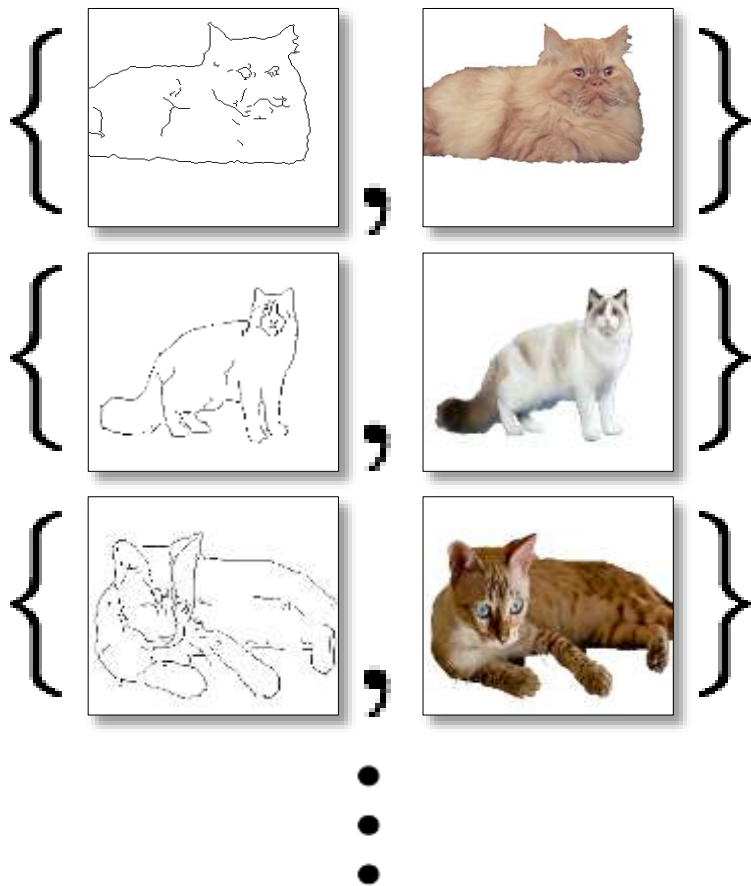
[Includes slides from Jun-Yan Zhu, Taesung Park]

[Cartoon: The Computer as a Communication Device, Licklider & Taylor 1968]

Paired data

X_i

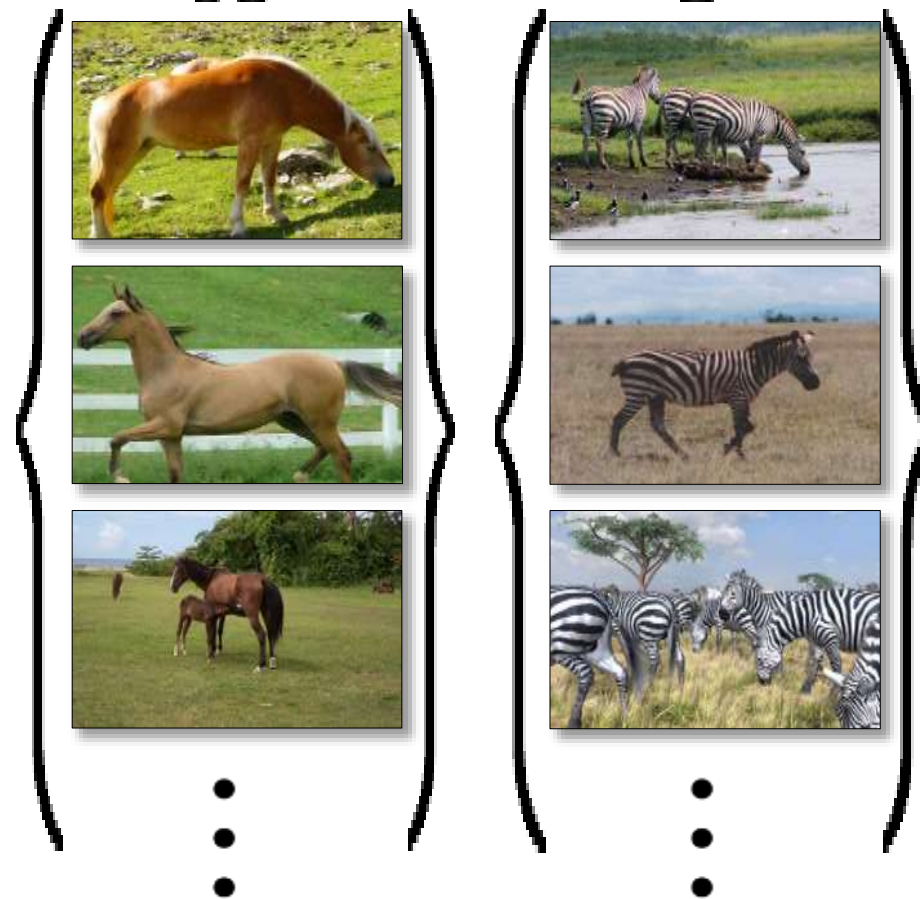
Y_i

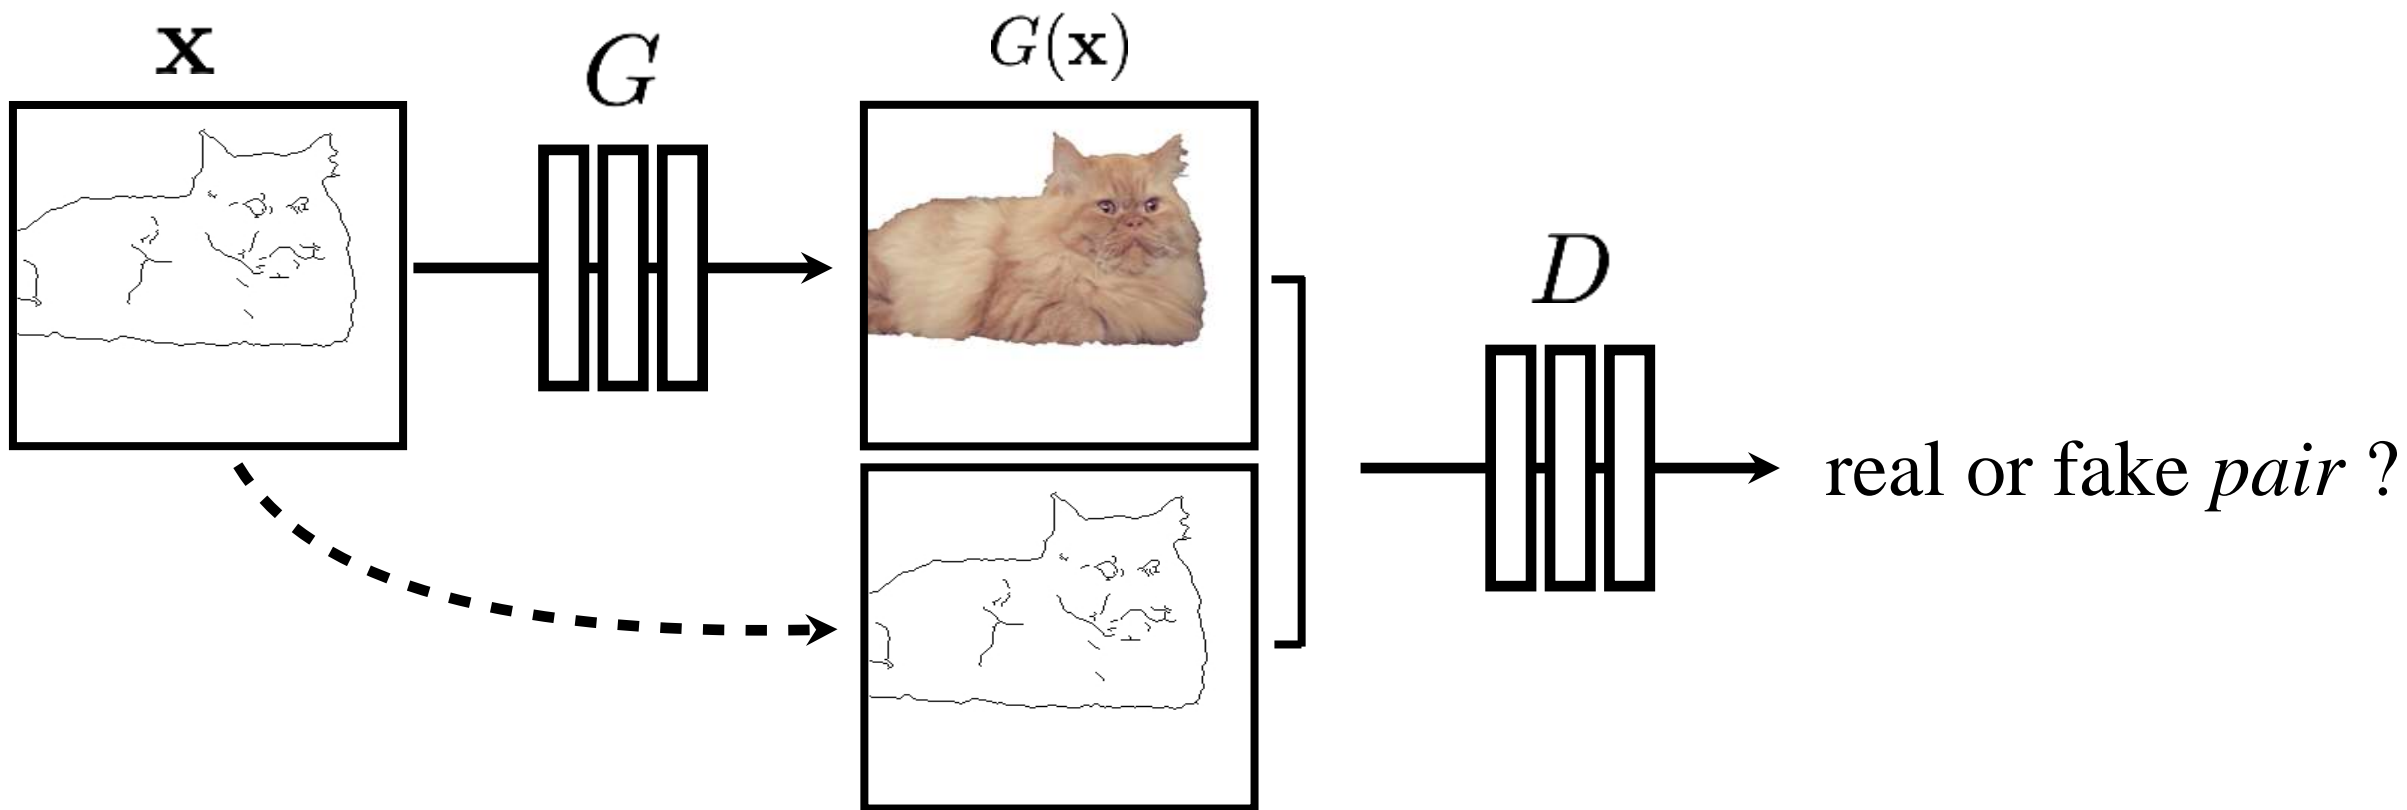


Unpaired data

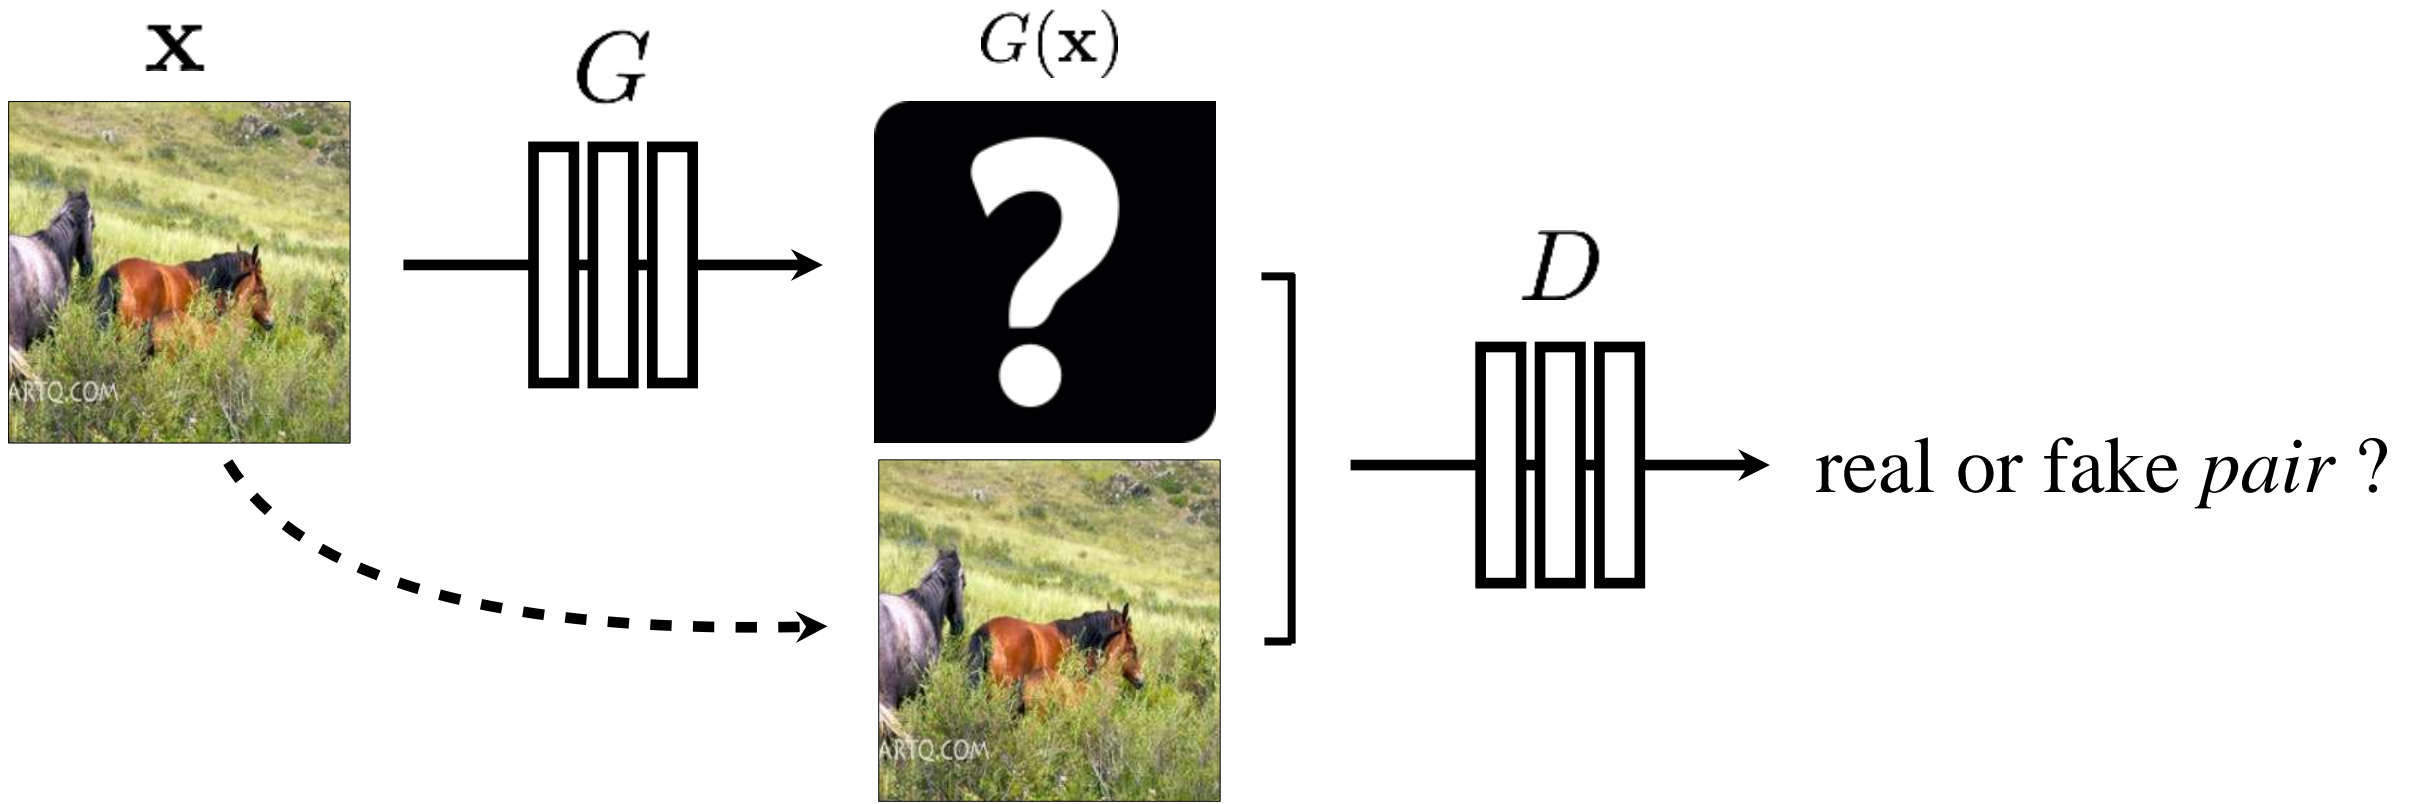
X

Y





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



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

No input-output pairs!



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

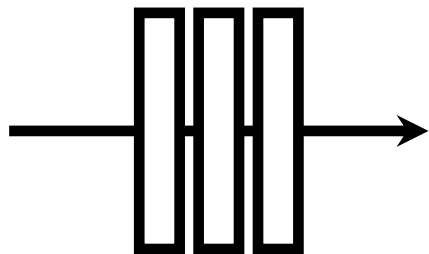
Usually loss functions check if output matches a target *instance*

GAN loss checks if output is part of an admissible *set*

\mathbf{x}




G



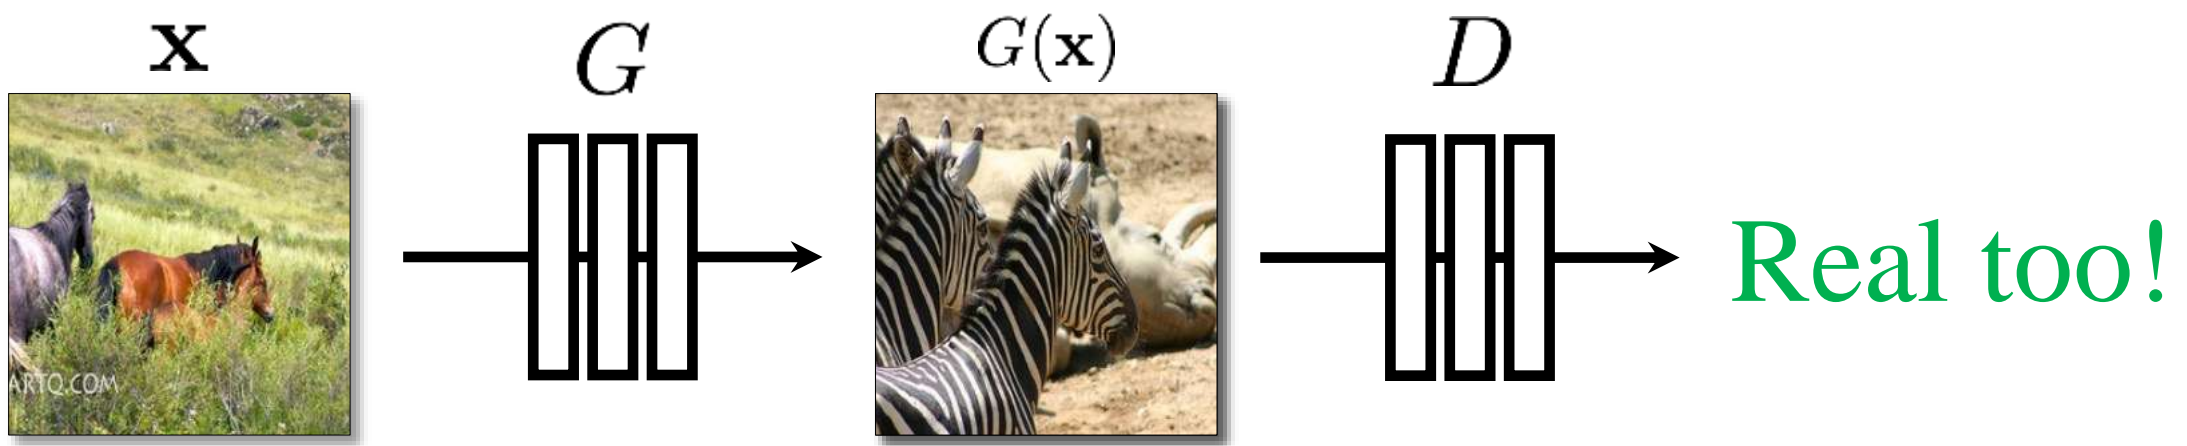
$G(\mathbf{x})$



D

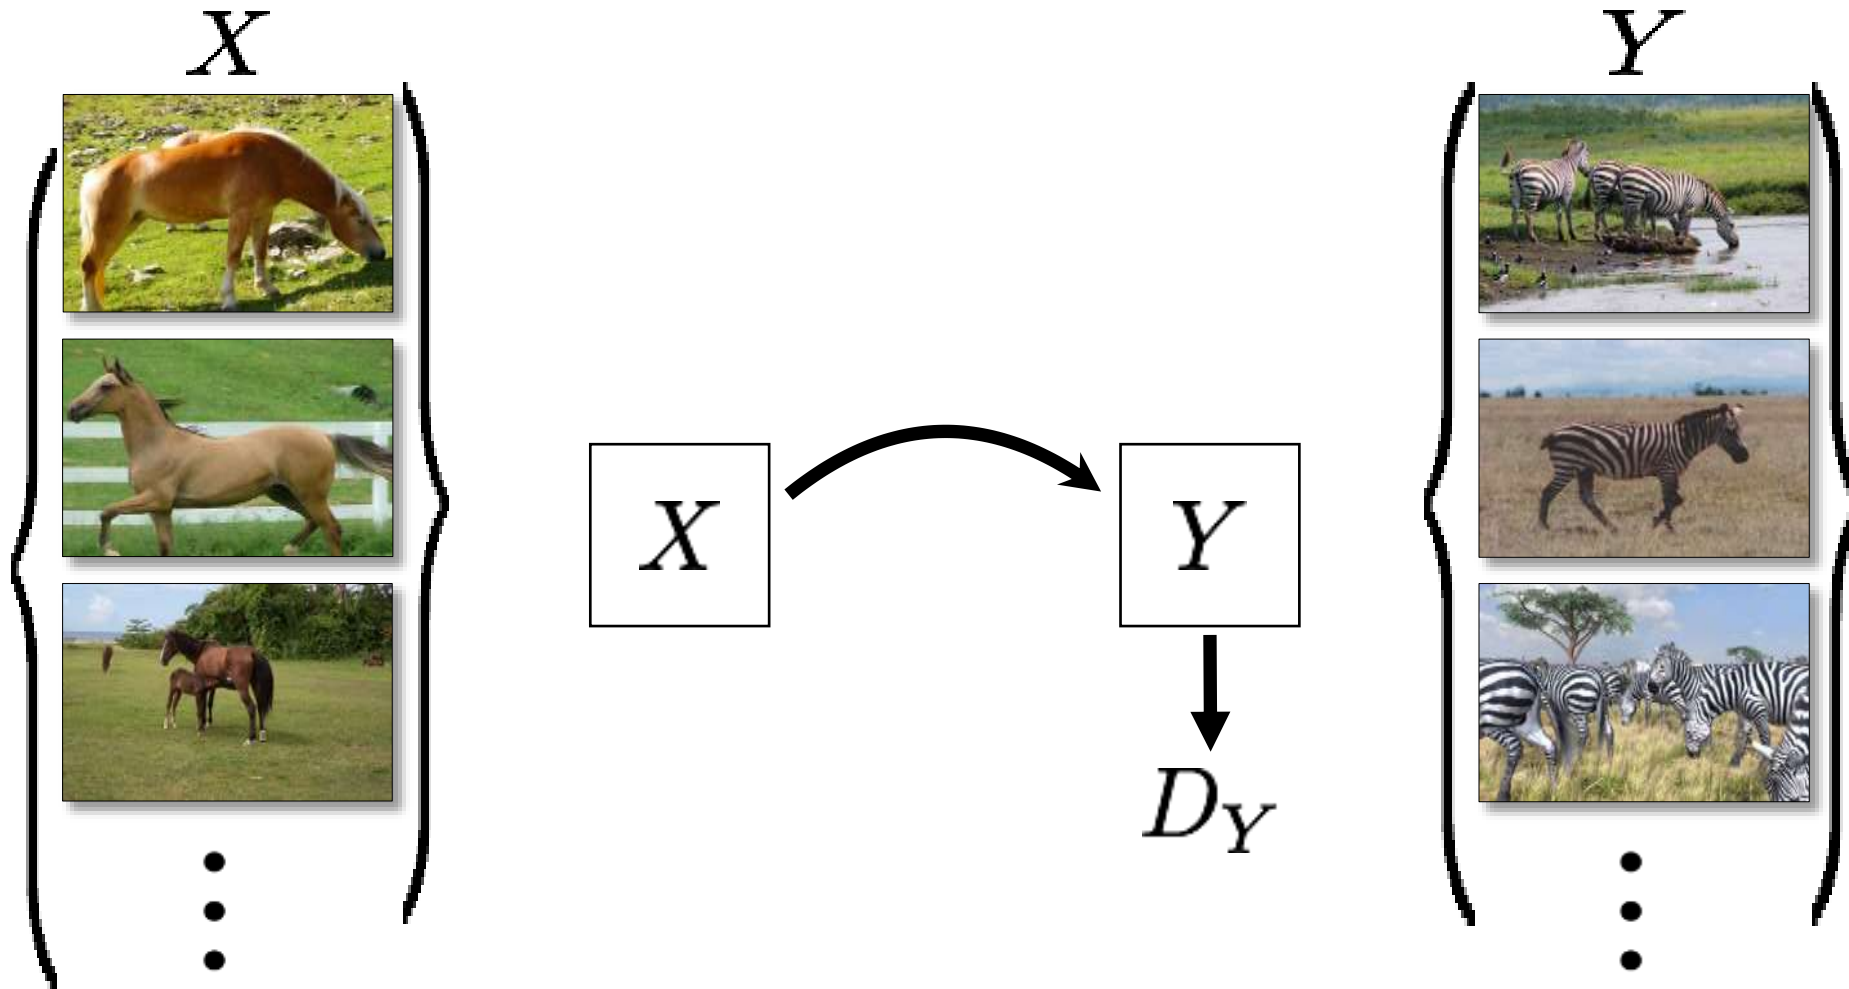


Real!



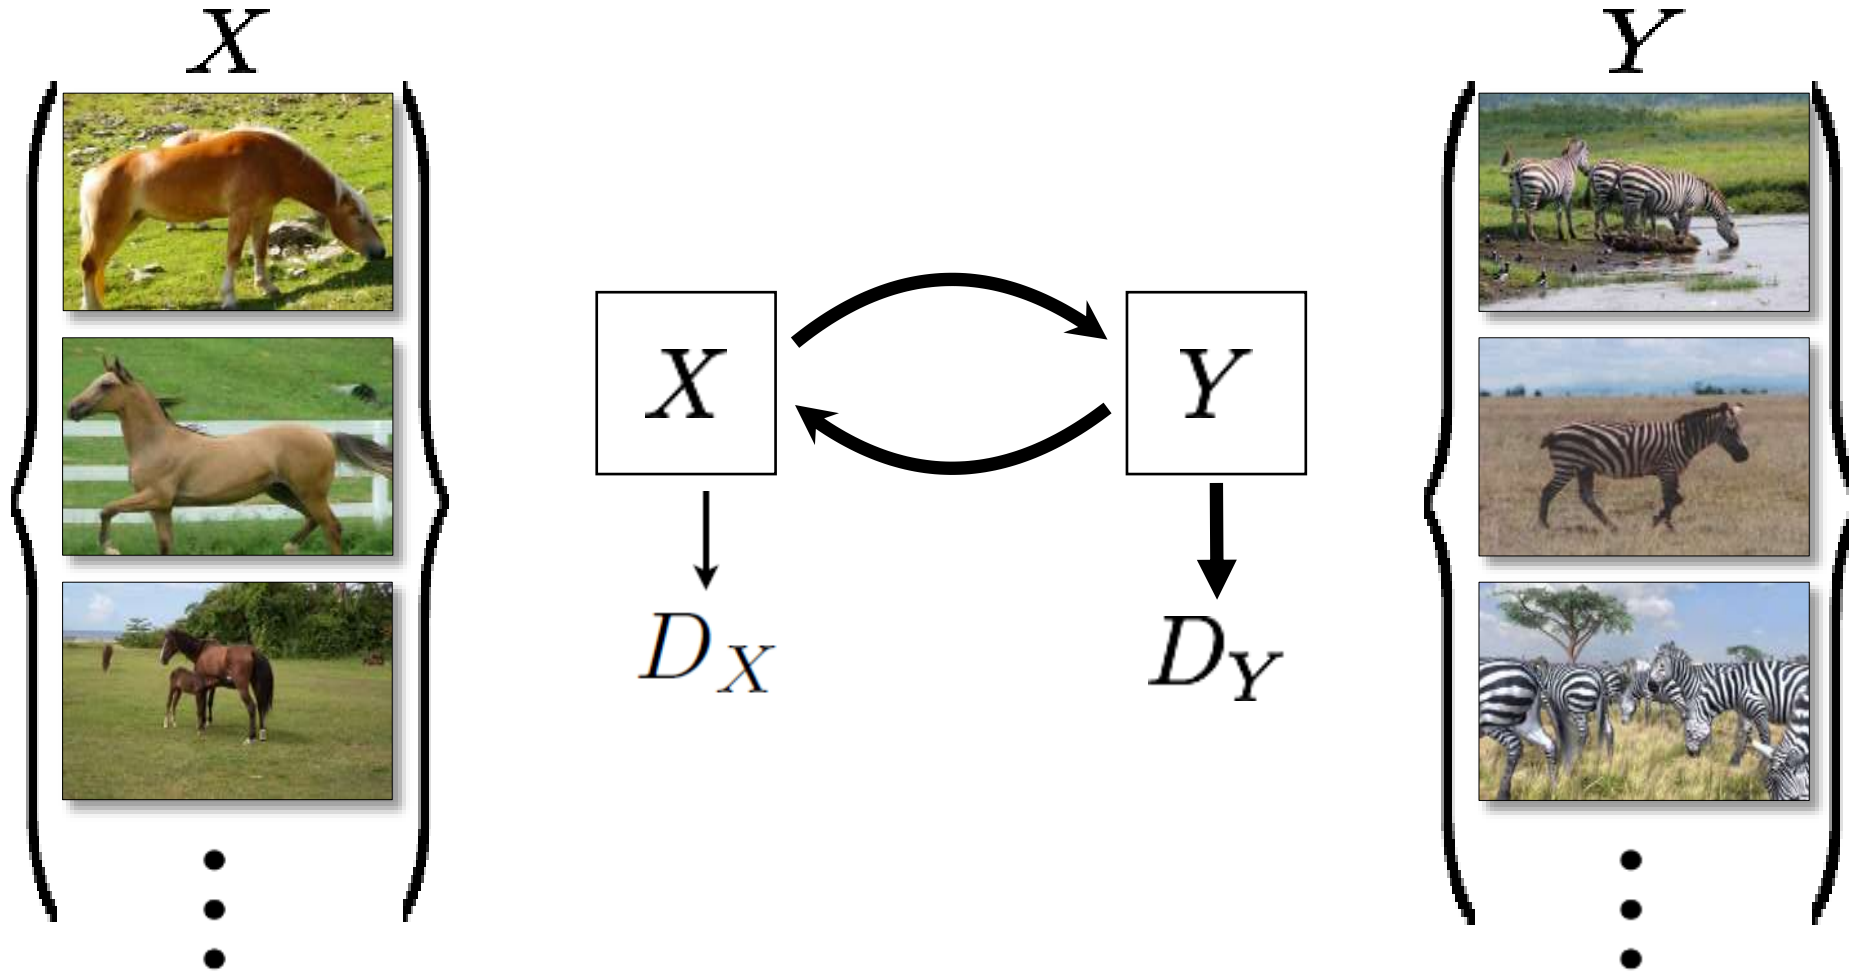
Nothing to force output to correspond to input

CycleGAN, or there and back aGAN

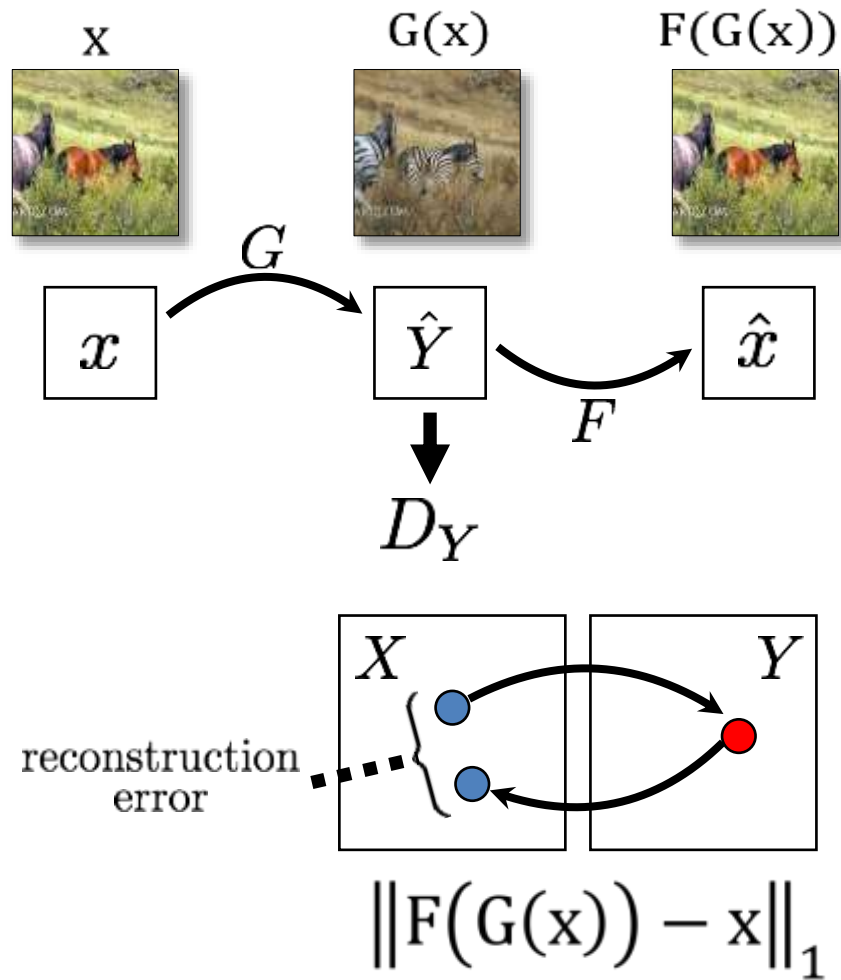


[Zhu*, Park* et al. 2017], [Yi et al. 2017], [Kim et al. 2017]

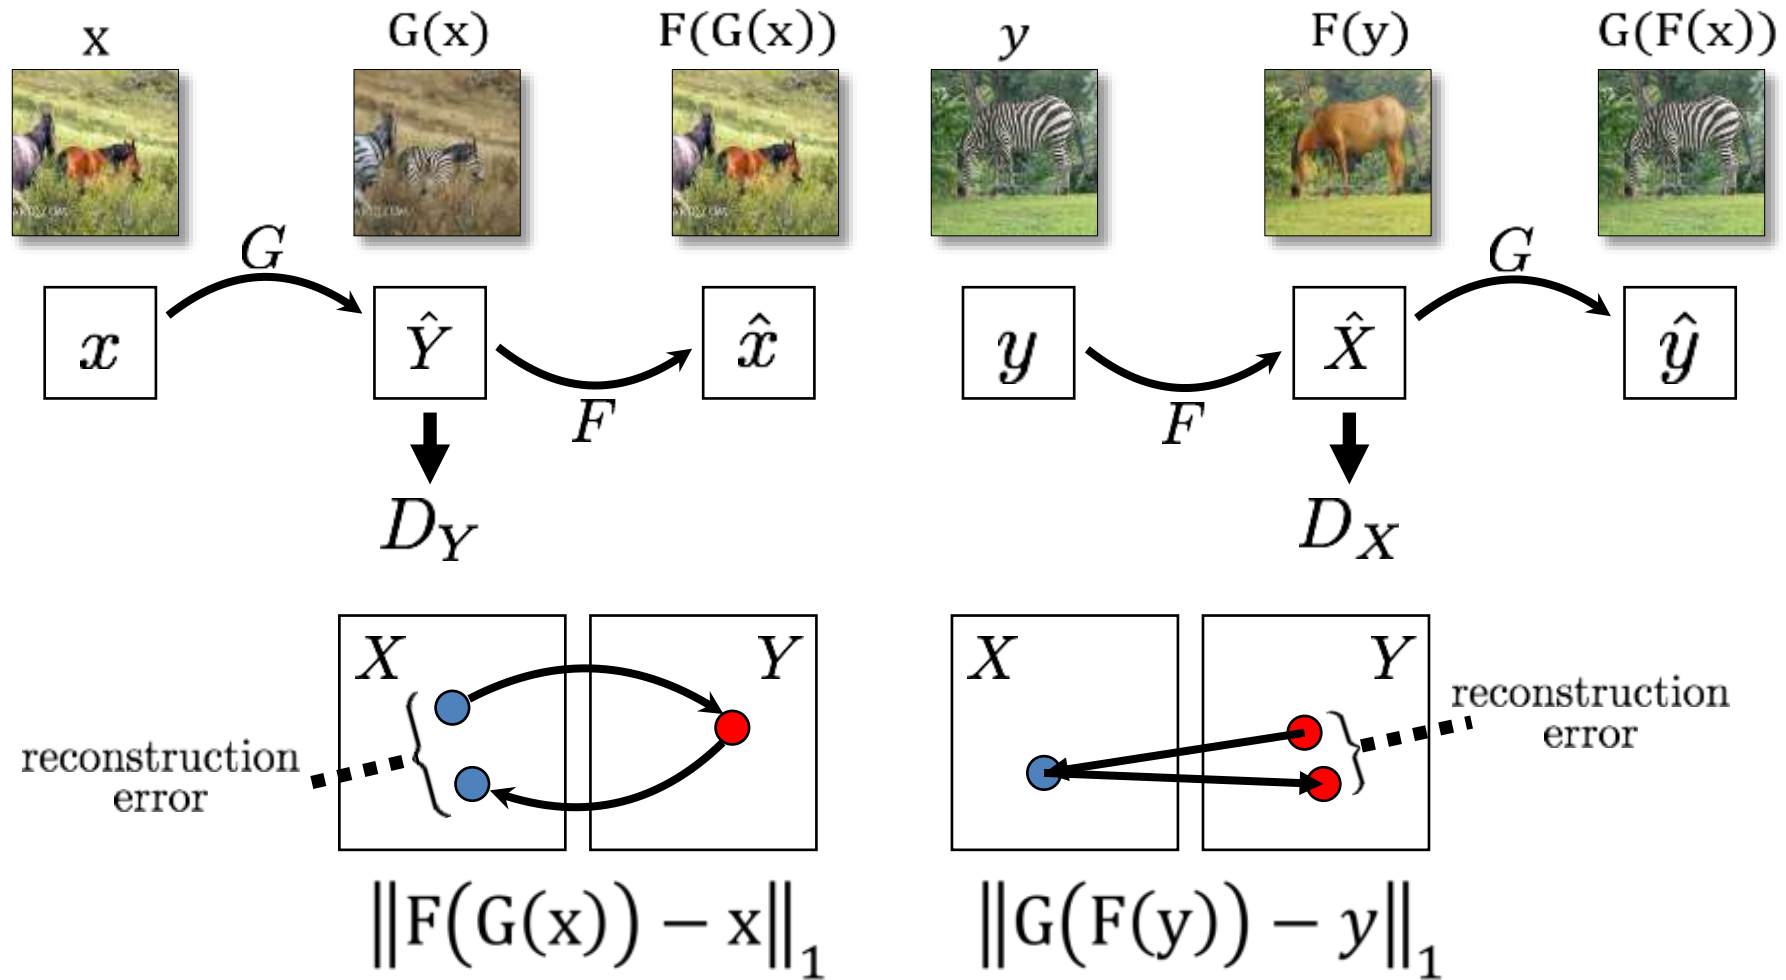
CycleGAN, or there and back aGAN



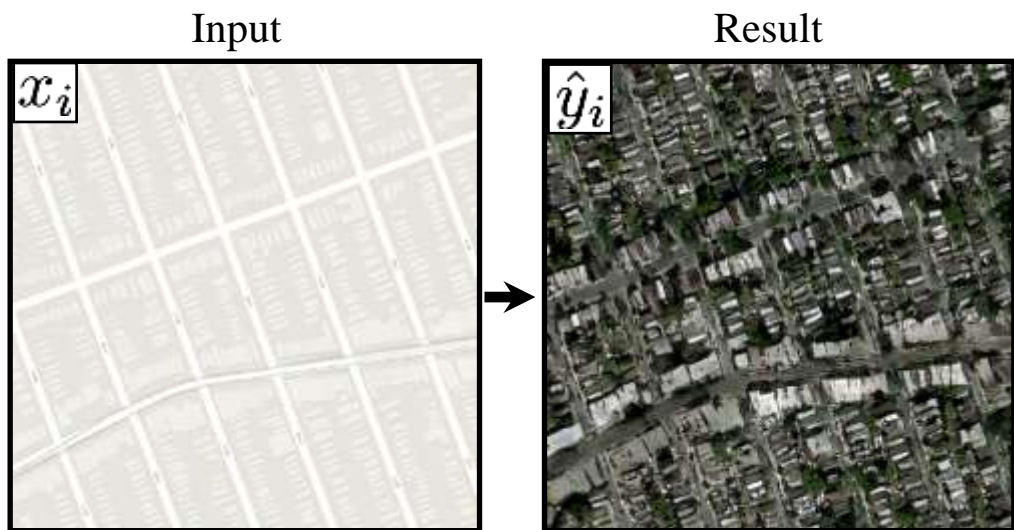
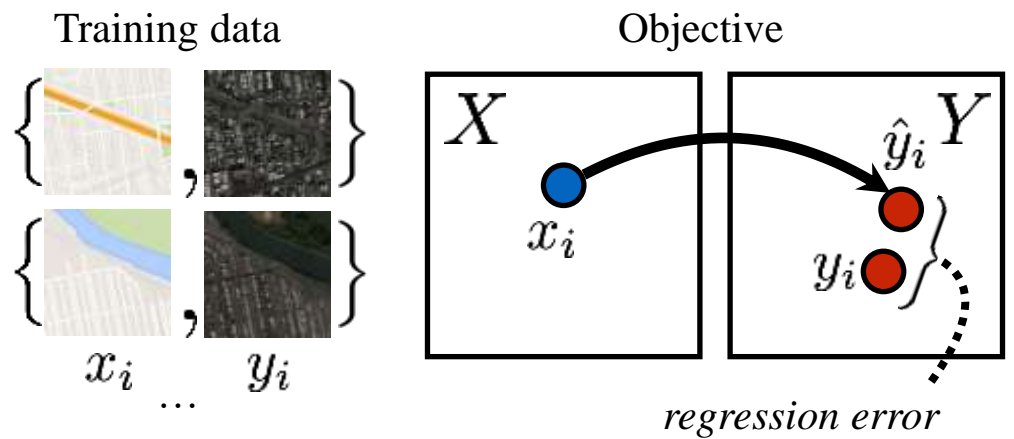
Cycle Consistency Loss



Cycle Consistency Loss




Paired translation



[“pix2pix”, Isola, Zhu, Zhou, Efros, 2017]

Unpaired translation



[“CycleGAN”, Zhu*, Park*, Isola, Efros*, 2017]





Input



Monet





Van Gogh



Cezanne



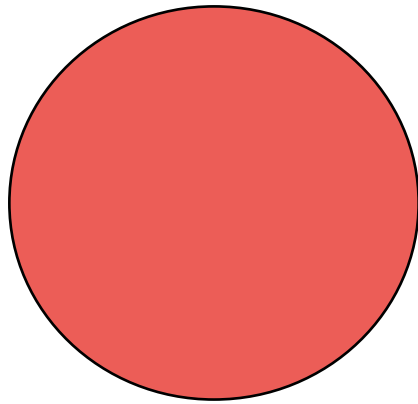
Ukiyo-e



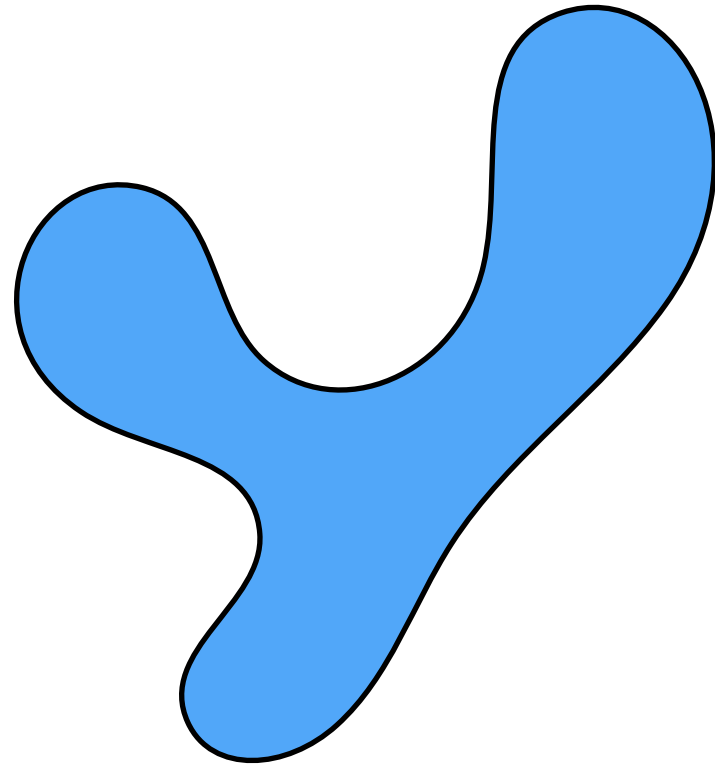
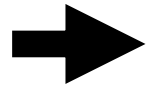
GANs

Gaussian

Target distribution



Z

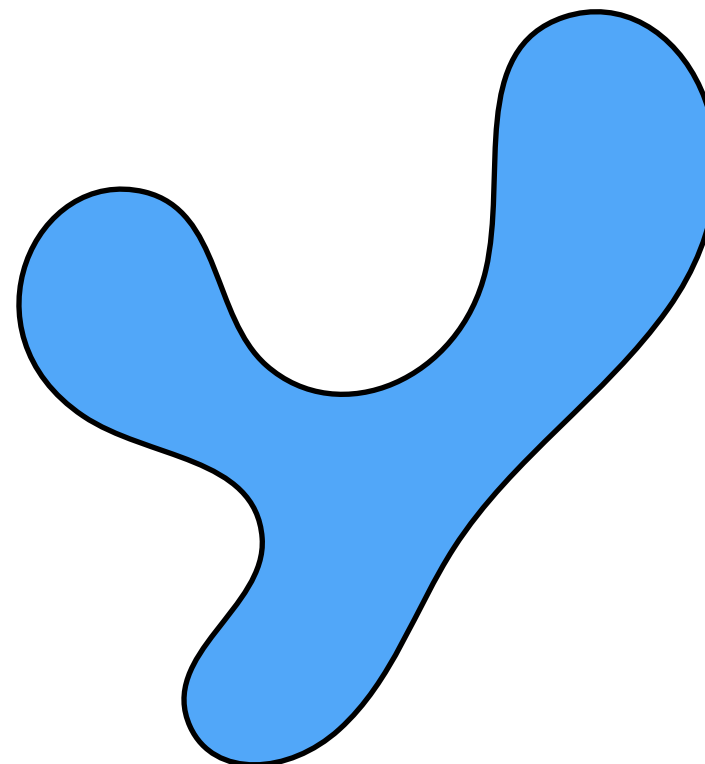
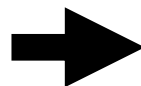
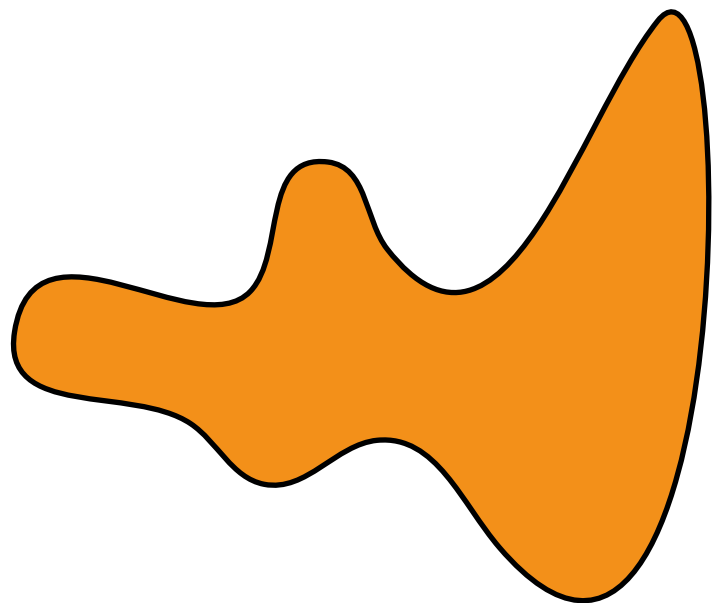


Y

CycleGAN

Horses

Zebras



X

Y