Plan

Recap

Logistics

GANs

Extensions

Conditional GANs

Recap

(s, a, r, s')

State Action Reward

$Q(s,a)$ = expected longterm discounted reward

$$Q(s,a) = r + \gamma \max_{a'} Q(s',a')$$

$$\mathcal{L} = \left(Q(s,a) - \left(r + \gamma \max_{a'} Q(s',a')\right)\right)^2$$

1. sample action with ϵ-greedy
2. update r, s'
3. $\theta \leftarrow \theta - \eta \frac{\partial \mathcal{L}}{\partial \theta}$

Logistics

Form: 23/25

- ↳ Contrastive learning
- ↳ Diffusion
- ↳ Implicit Regularization

Q Learning Problem

↳ more guidance now

Generative Tasks

- ☑ Supervised Learning
- ☑ Reinforcement Learning
- ◯ Unsupervised Learning
  - ↳ embedding

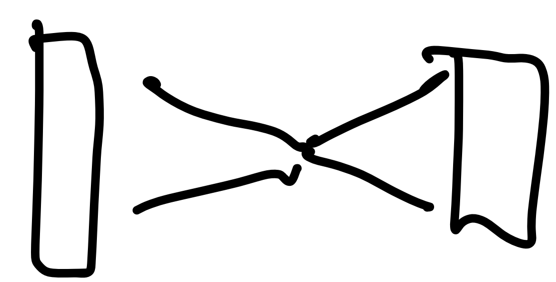

dense, LSTM, CNN, gradient descent

Our techniques still work

Attention, ℓ_2 norm, RNN, cross entropy

Build Realistic Images

Generator $G_\theta(z)$

z

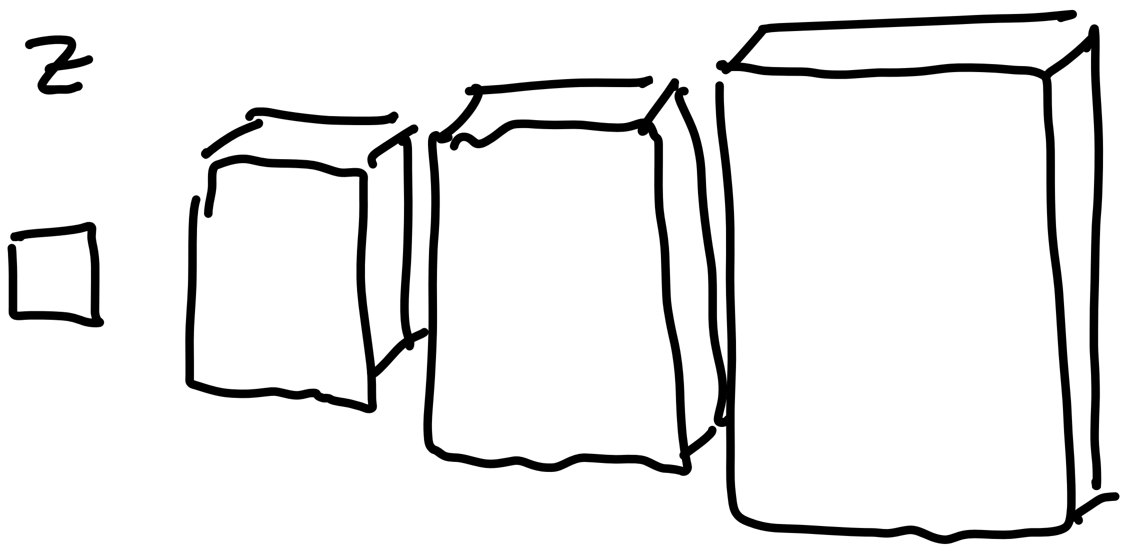

noise input

θ parameters of Generator

image output

☑ Architecture

☐ Loss?

Human would work but too slow.

Discriminator

image input

ψ "psi" parameters of discriminator

realness score output

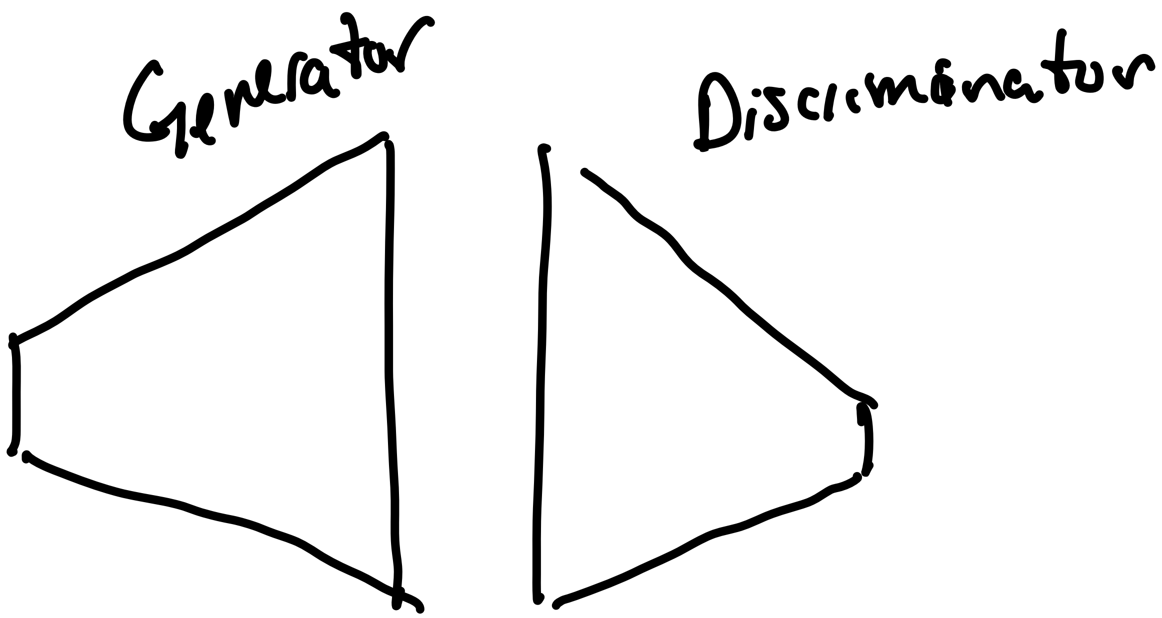

Both need to learn

↳ training for D

X , y
image , boolean real

$$D_\psi(x) = \hat{y}$$

$$\mathcal{L}(\psi) = \text{cross entropy}(y, \hat{y})$$

$$= \sum_{i=1}^{d} y_i \cdot -\log(\hat{y}_i)$$

$$y = \begin{bmatrix}0\\1\end{bmatrix} \leftarrow \text{fake} \qquad = \cancel{0 \cdot -\log(\hat{y}_1)} + 1 \cdot -\log(\hat{y}_2)$$

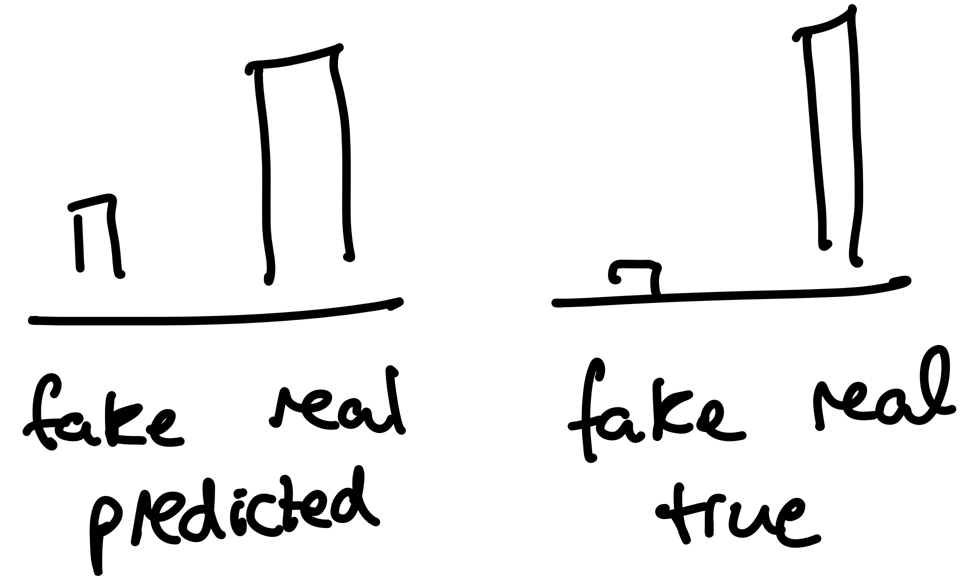

↳ Training for G

Fool D


Alexis
Liam
Violet
Jared
Hugh
Asif
Alex
Sujay
Julian
Jessica

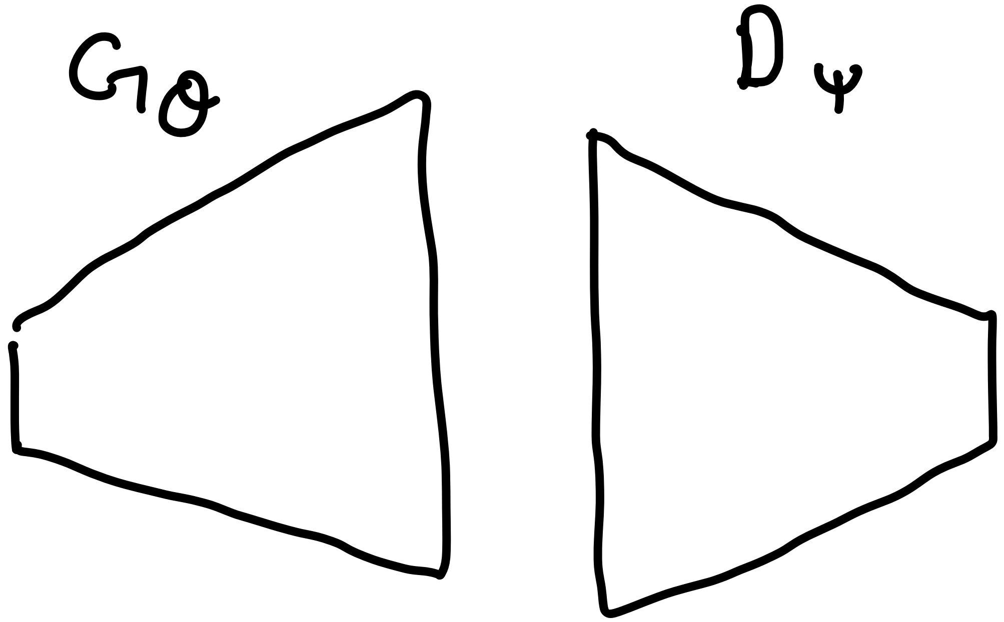

<u>cross entropy</u>

$y = 1$ if real 0 else

$X \in \mathbb{R}^{d \times d}$

$$\mathcal{L}(\psi) = -y \log(D_\psi(x)) - (1-y)\log(1 - D_\psi(x))$$

$$= -\log D_\psi(x)$$ ← real image

$$- \log(1 - D_\psi(\text{fake}))$$ ← fake image

$$\text{fake} = G_\theta(z)$$

$$\mathcal{L}(\psi, \theta) = \mathbb{E}_{x \sim \text{real}}[\log D_\psi(x)] + \mathbb{E}_{z \sim \text{random}}[\log(1 - D_\psi(G_\theta(z)))]$$

Generator wants to ... minimize

Discriminator wants to ... maximize

$$\theta \leftarrow \theta - \eta \frac{\partial \mathcal{L}}{\partial \theta}$$

$$\psi \leftarrow \psi + \eta \frac{\partial \mathcal{L}}{\partial \psi}$$

$$\nabla_\theta \mathcal{L}(\theta, \psi) = \nabla_\theta \mathbb{E}_{z \sim \text{random}}\left[\log(1 - D_\psi(G_\theta(z)))\right]$$

<u>Hack</u>

Train generator more

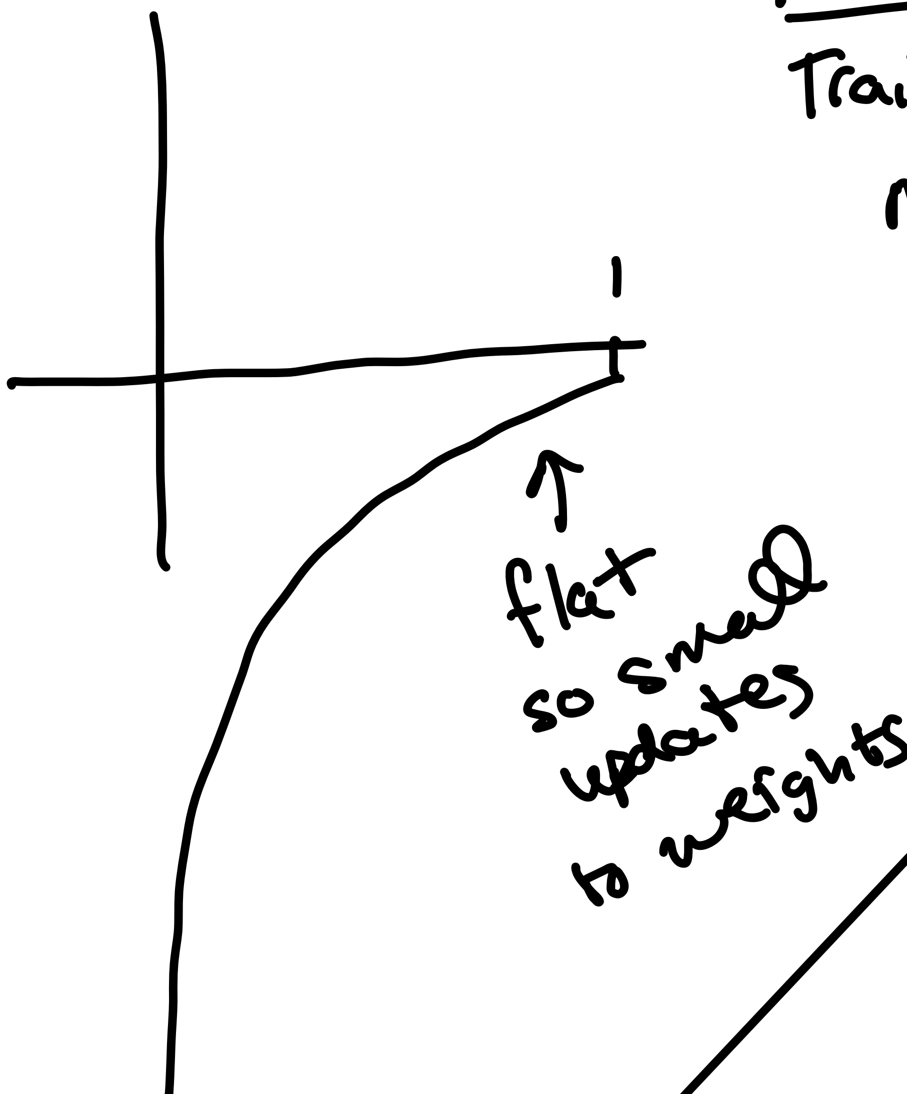

$$\mathcal{L}(\theta, \psi) = \mathbb{E}_{x \sim \text{real}}\left[D_\psi(x)\right] + \mathbb{E}_{z \sim \text{random}}\left[-D_\psi(G_\theta(z))\right]$$

Wasserstein
aka earth mover

WGAN

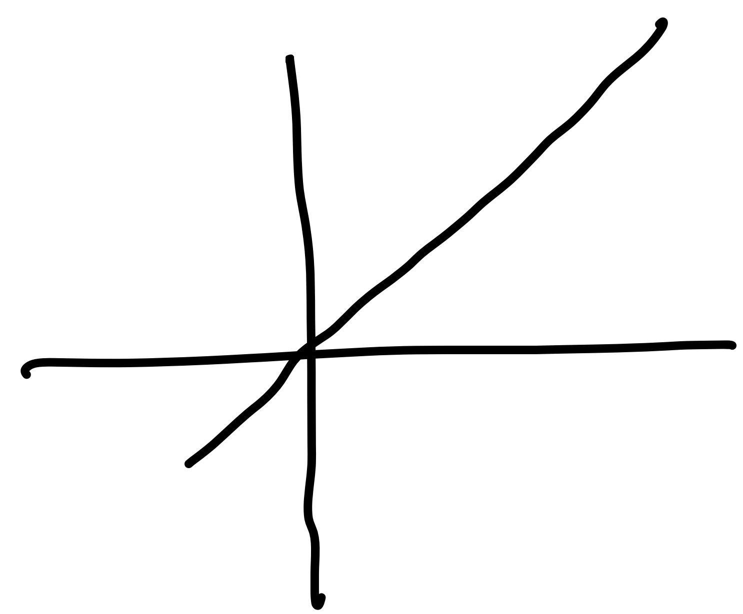

Generative Adversarial Network

G can cheat to win

Mode collapse

↳ early stopping

↳ diversity of real images

↳ dropout

Conditional GAN

↳ category dependent generation

$$\mathcal{L}(\theta, \varphi) = \mathbb{E}_{x \sim real}\left[D_\varphi(x, c)\right] - \mathbb{E}_{z \sim random}\left[D_\varphi(G_\theta(x, c))\right]$$

c one-hot encoded class

Discriminator can discern if either fake image OR wrong label

↳ harder for generator

Style Transfer

G_1 : Style 1 to Style 2

Train discriminates with real examples of Style 2

G_2 : Style 2 to Style 1

ℓ_2 loss so $x \approx G_2(G_1(x))$

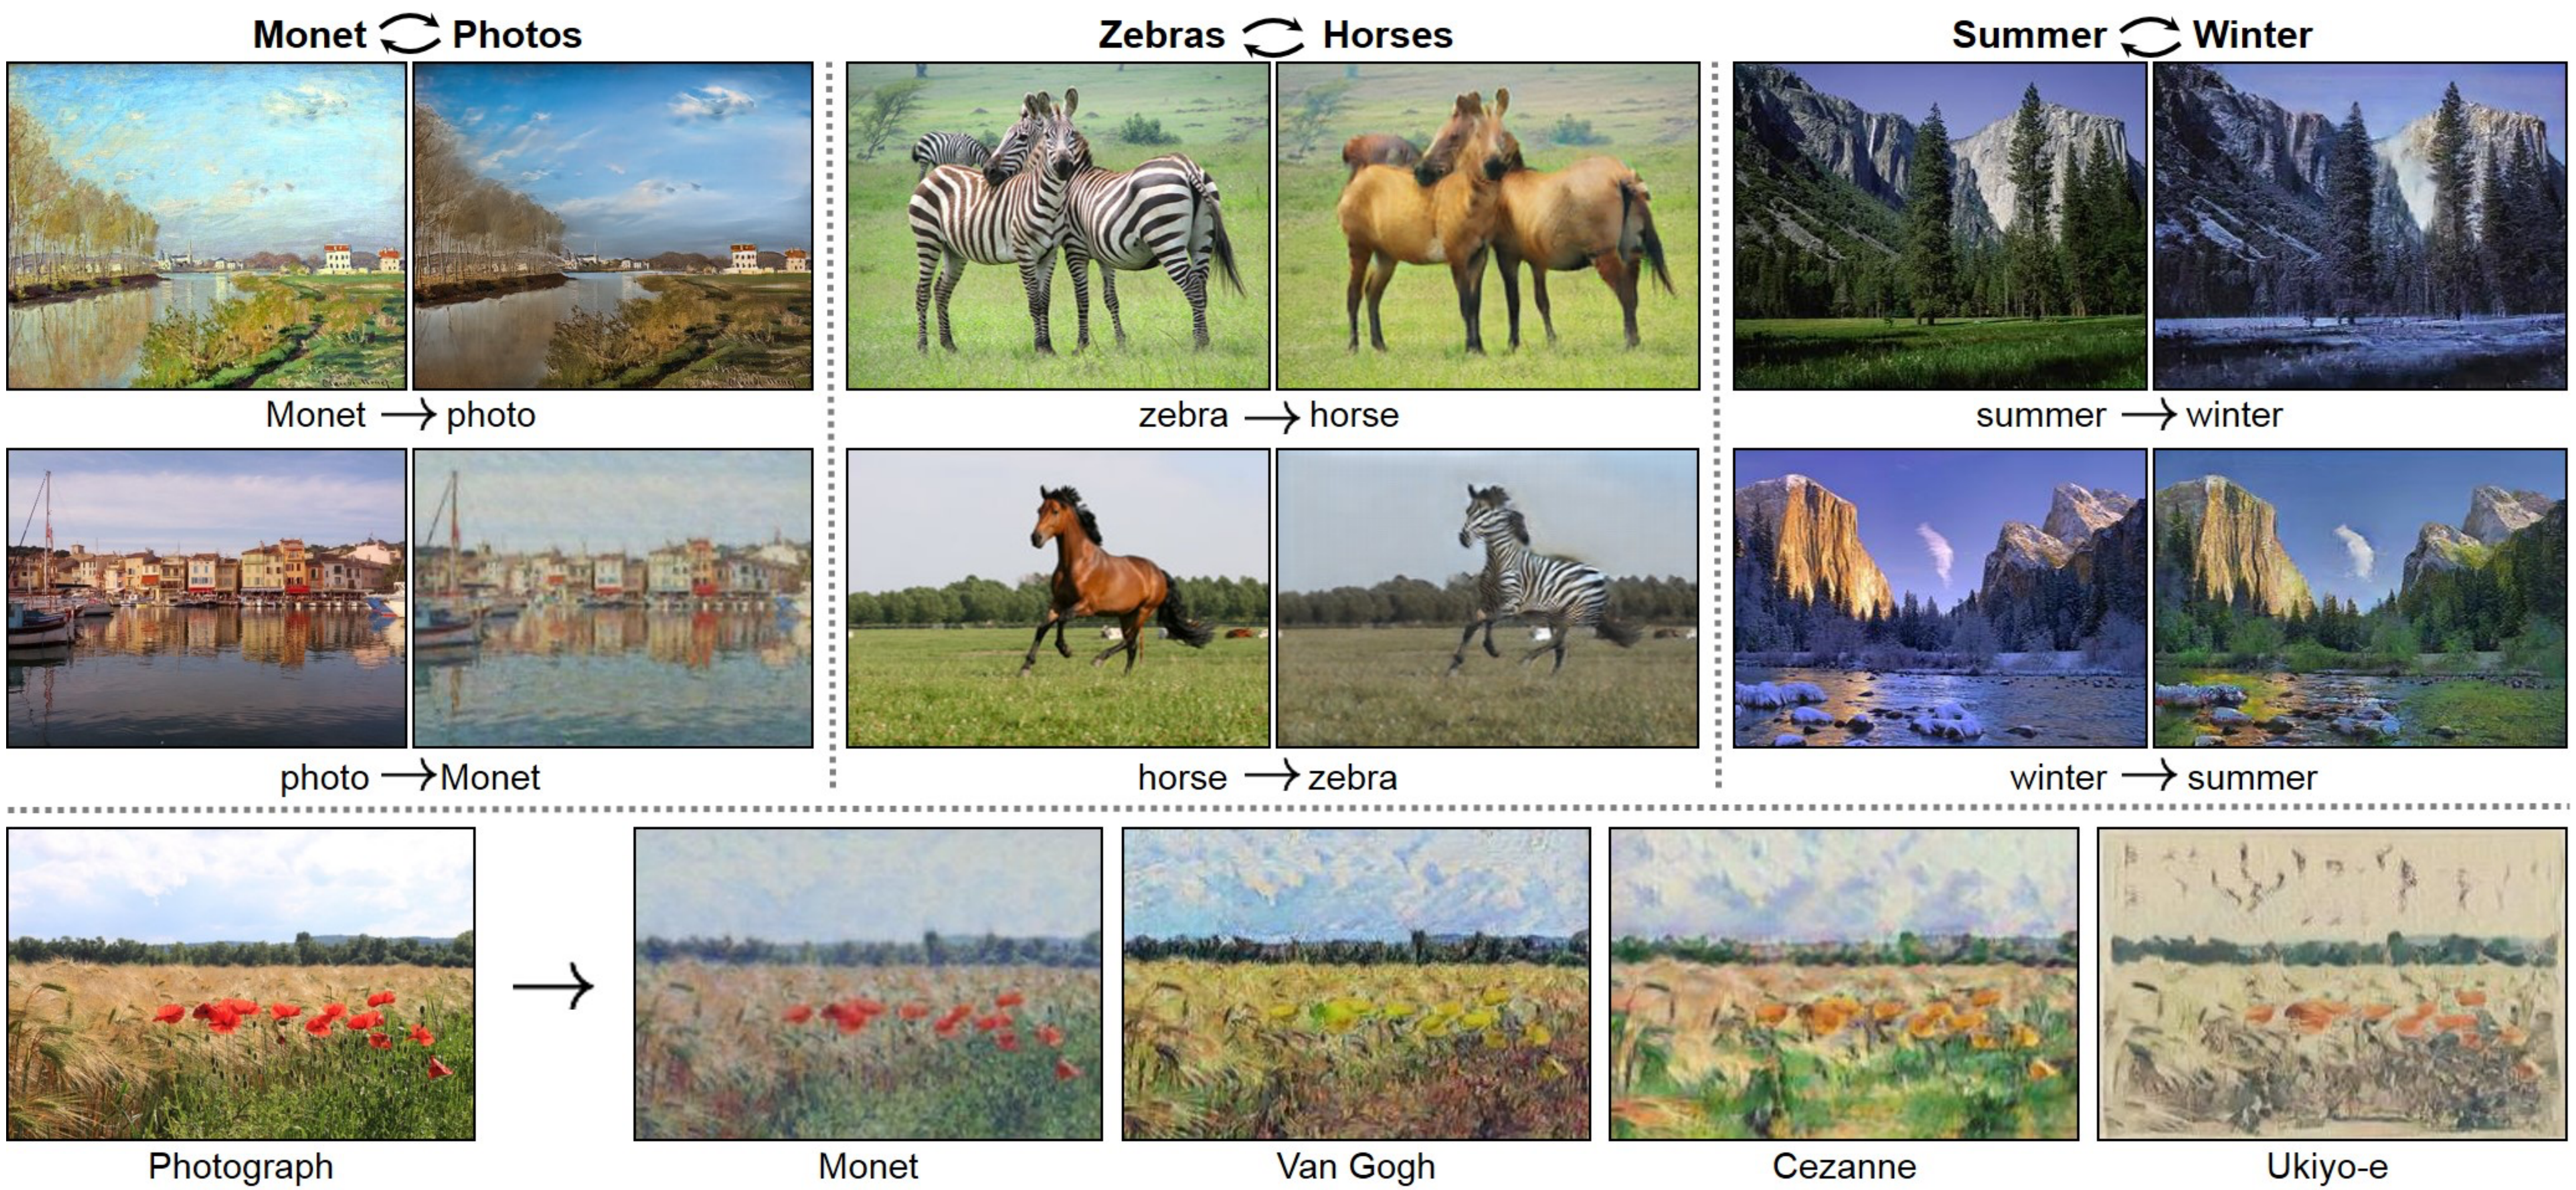

Monet ⇄ Photos
Zebras ⇄ Horses
Summer ⇄ Winter
Monet → photo
zebra → horse
summer → winter
photo → Monet
horse → zebra
winter → summer
Photograph
→
Monet
Van Gogh
Cezanne
Ukiyo-e