STOR566: Introduction to Deep Learning

Lecture 13: Generative Models II

Yao Li
UNC Chapel Hill

Oct 06, 2022

Materials are from *Deep Learning (UCLA)*

Generated Adversarial Network

- Discriminative models:
  - Given an image x, predict a label y
  - (by learning $P(y \mid x)$)
- Generative models:
  - Generate new images
  - Learn $P(x)$ (or $P(x, y)$, $P(x \mid y)$)

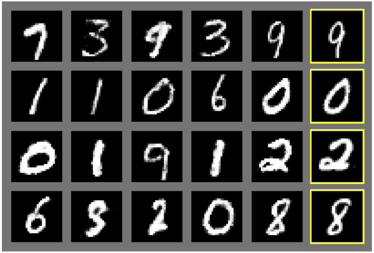

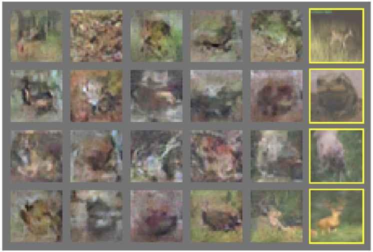

(Goodfellow et al., 2014)

How to represent a distribution

- Define the distribution implicitly
- Start from a random vector z: a simple distribution (e.g., sphere Gaussian)
- Define (the sampling process of) the distribution as a function G:

$$z \to G(z) = x$$

- Our goal is to learn this generator function G

How to represent a distribution

- Define the distribution implicitly
- Start from a random vector z: a simple distribution (e.g., sphere Gaussian)
- Define (the sampling process of) the distribution as a function G:

$$z \to G(z) = x$$

- Our goal is to learn this generator function G

Example:

- Gaussian with covariance matrix $N(0, \Sigma)$

$$z \sim N(0, I) \qquad \to \qquad \underbrace{\Sigma^{1/2} z}_{G(z)} \sim N(0, \Sigma)$$

Neural network as a generator

- Now we assume G is a neural network parameterized by θ
- Goal: learn θ to make generated distribution similar to the data distribution

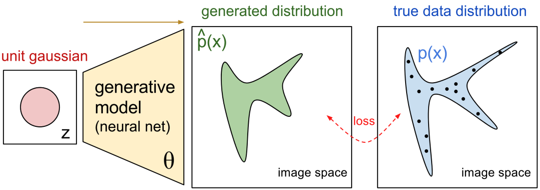

(figure from https://openai.com/blog/generative-models/)

- But how to evaluate the quality of generated distribution?

Generative Adversarial Network (GAN)

- A good measurement: whether there exists a discriminator (classifier) to distinguish real/fake images
- Generative Adversarial Network (GAN): Train two networks jointly
  - The generator network tries to produce realistic-looking images
  - The discriminator network tries to classify real vs fake images

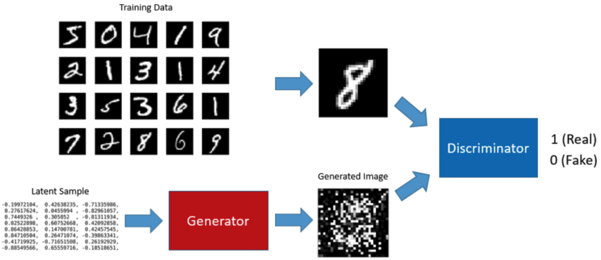

(figure fromhttps://naokishibuya.medium.com/understanding-generative-adversarial-networks)

Training objective

- The discriminator's goal: classify real/fake images

$$L_D = E_{x \sim \text{real data}}\big[-\log D(x)\big] + E_z\big[-\log(1 - D(G(z)))\big]$$

- Generator's goal: fool the discriminator
- A simple cost function for generator: the opposite of the discriminator's
- The minmax training objective:

$$\max_G \min_D L_D(G, D)$$

- GAN training: alternatively update G and D

Gradient vanishing problem

$$\max_G \min_D E_{x\sim \text{real data}}\big[-\log D(x)\big] + E_z\big[-\log(1-D(G(z)))\big]$$

- The discriminator is usually much better than generators ($D(G(z)) \to 0$), this implies the gradient of generator will vanish
- A modified generator loss:

$$L_G = E_z\big[-\log D(G(z))\big]$$

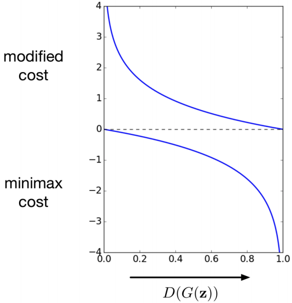

CNN for both generator and discriminator (DC-GAN)

- Discriminator: a regular classification network
- Generator: CNN with transposed convolution structure

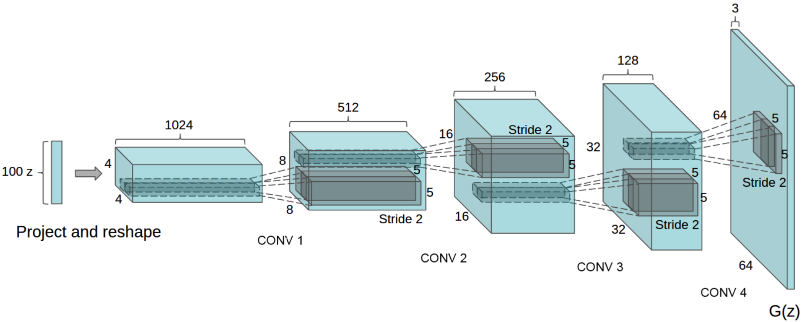

(Radford et al., 2015)

DC-GAN results

(Figure from Raford et al., 2015)

Many improvements have been made

- c-GAN (Mirza and Osindero, 2014): add class label into the generator
- AC-GAN (Odena et al., 2016): discriminator classifies both real/fake and class label
- WGAN (Arjovsky et al., 2017): use Wasserstein distance
- SN-GAN (Miyato et al., 2018): spectral regularization
- Big-GAN (Brock et al., 2018): large batch (2048), bigger model
- Fast-GAN (Liu and Hsieh, 2018), (Zhong et al., 2020): small batch (64) can also work with adversarial training
- Style-GAN1,2,3 (Karras et al., 2018; Karras et al., 2019; Karras et al., 2021): latent code transformation, progressive growing GAN

Big-GAN results

(Figure from Brock et al., 2018)

Image-to-image translation

Cycle GAN: Zhu et al., 2017

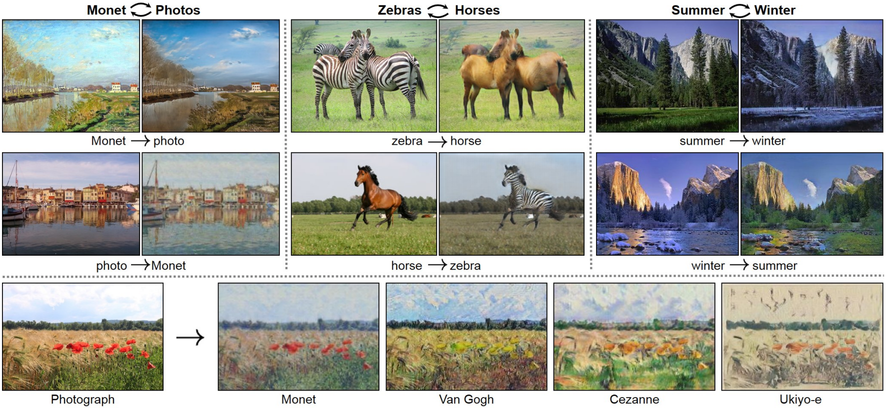

Many applications in bioinformatics

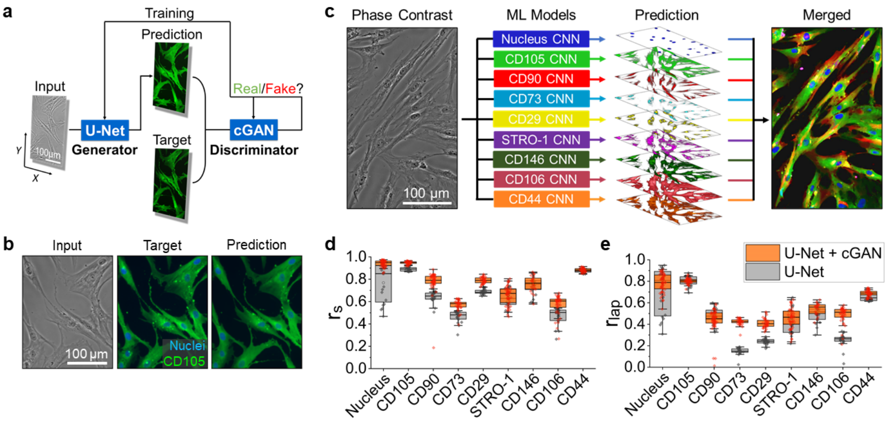

Image-to-image translation

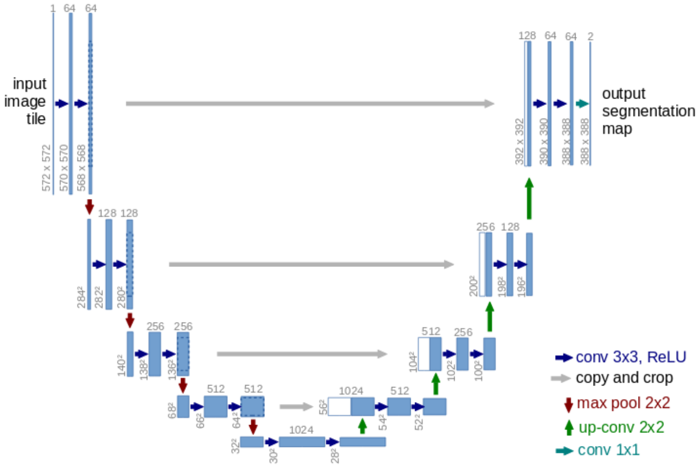

(The unet architecture)

Conclusions

- Generative Adversarial Network
- Variants of GAN

Questions?