ENGR 4350: Applied Deep Learning

Generative Adversarial Networks

11/30/2022

Outline

- Overview
- Architecture
- Training Process
- Case Study: StyleGAN

GAN Overview

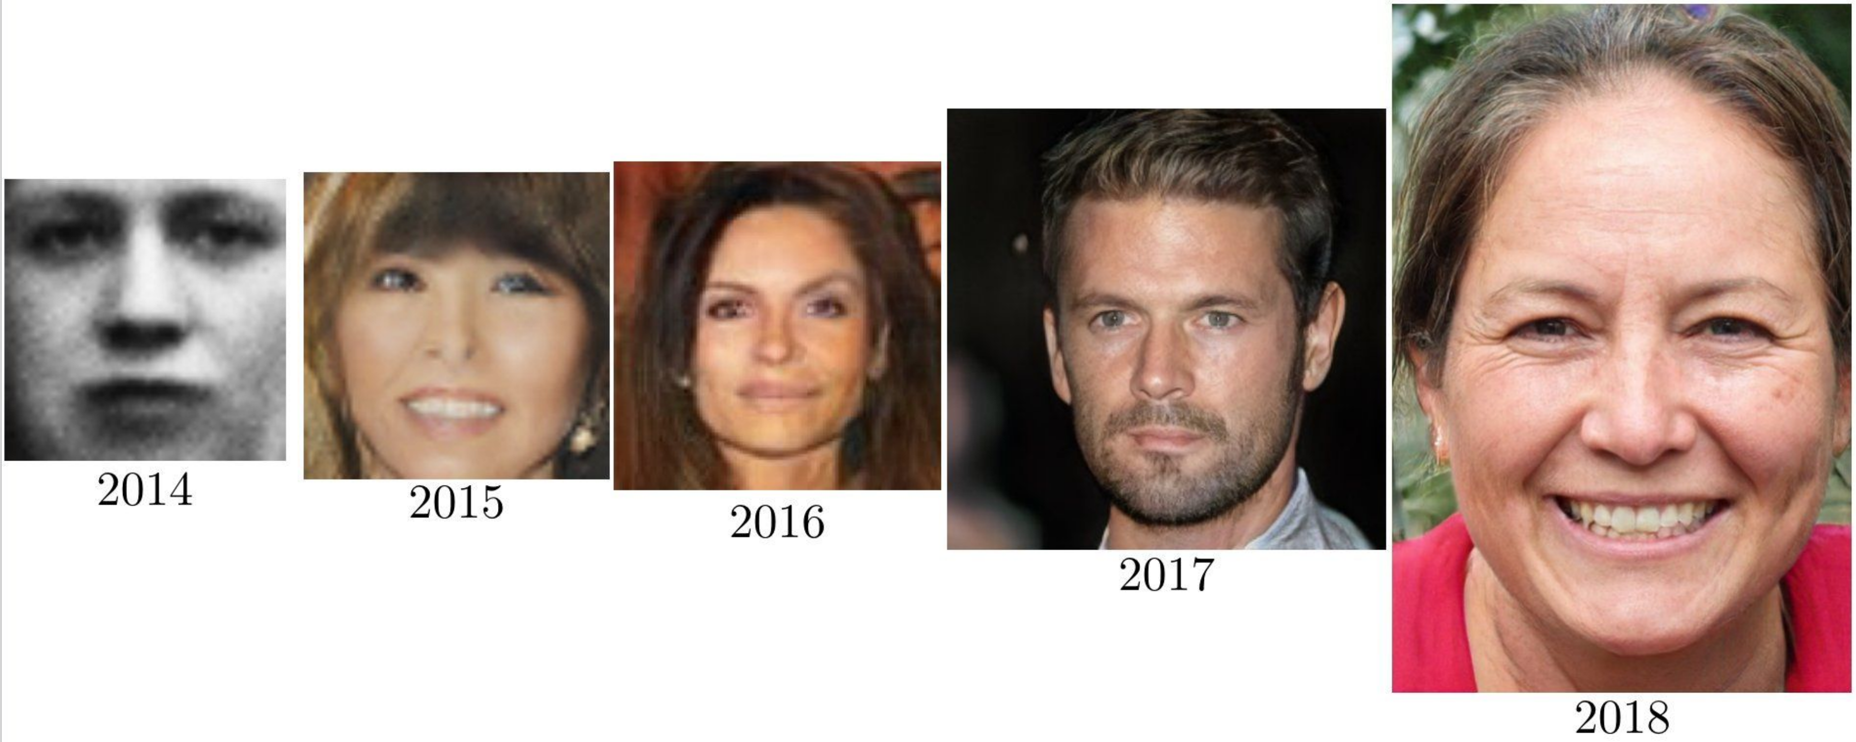

GAN Example: Style Transfer

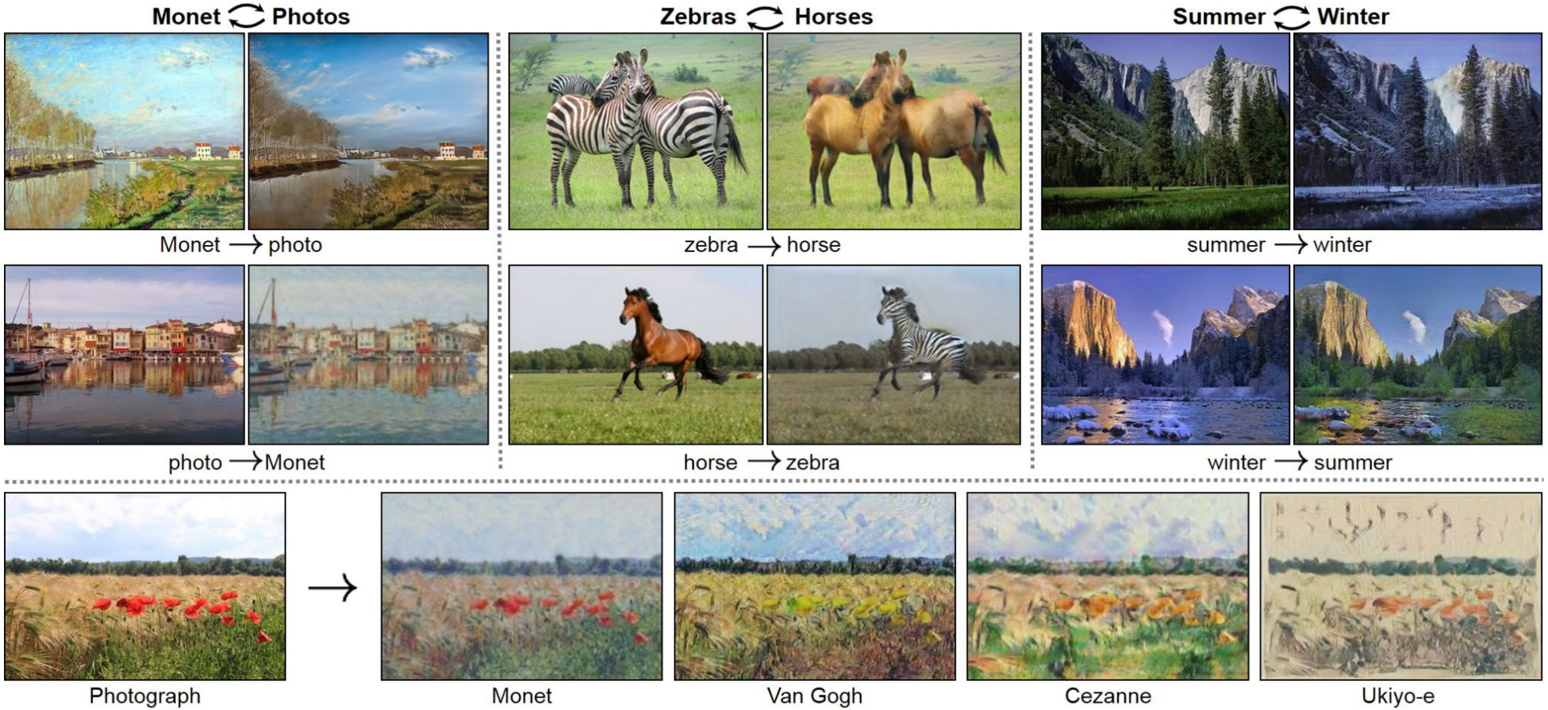

Zhu JY, Park T, Isola P, Efros AA. Unpaired image-to-image translation using cycle-consistent adversarial networks. InProceedings of the IEEE international conference on computer vision 2017 (pp. 2223-2232).

GAN Overview

Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J, Aila T. Analyzing and improving the image quality of stylegan. InProceedings of the IEEE/CVF conference on computer vision and pattern recognition 2020 (pp. 8110-8119).

GAN Overview

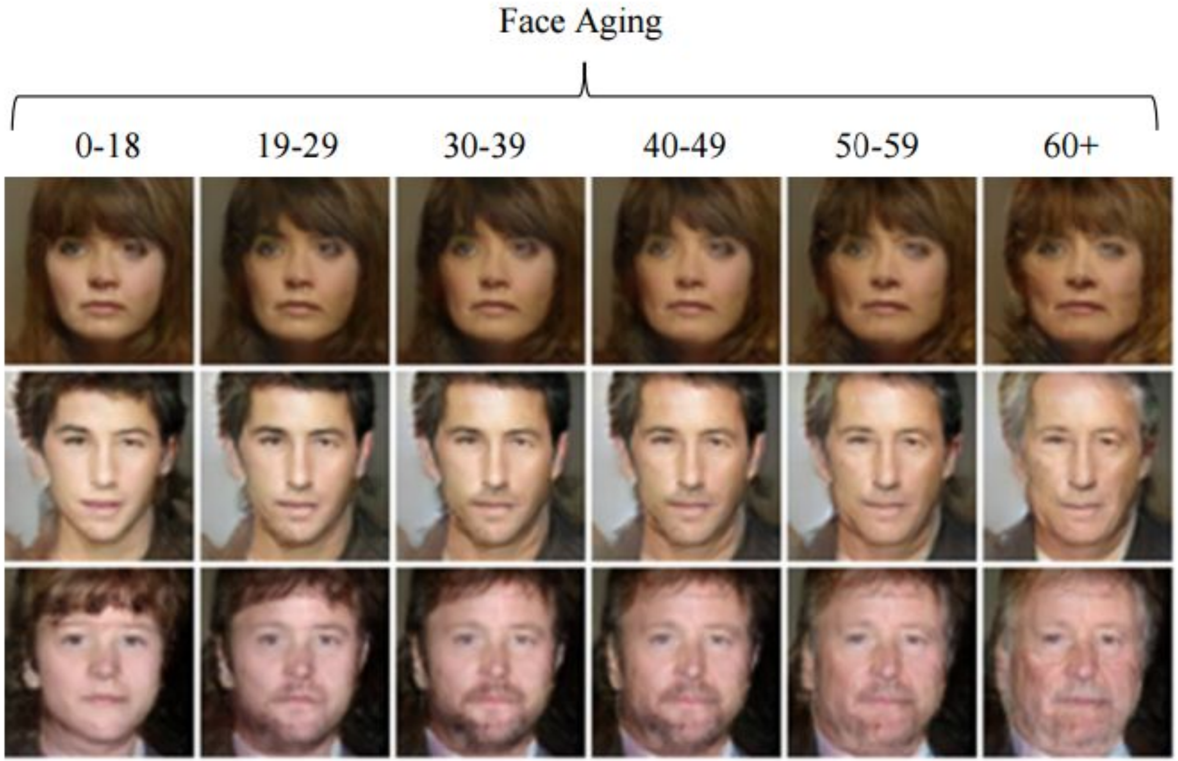

Antipov G, Baccouche M, Dugelay JL. Face aging with conditional generative adversarial networks. In2017 IEEE international conference on image processing (ICIP) 2017 Sep 17 (pp. 2089-2093). IEEE.

GAN Overview

Wang X, Li Y, Zhang H, Shan Y. Towards real-world blind face restoration with generative facial prior. InProceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition 2021 (pp. 9168-9178).

GAN Overview

https://thispersondoesnotexist.com/

GAN Architecture

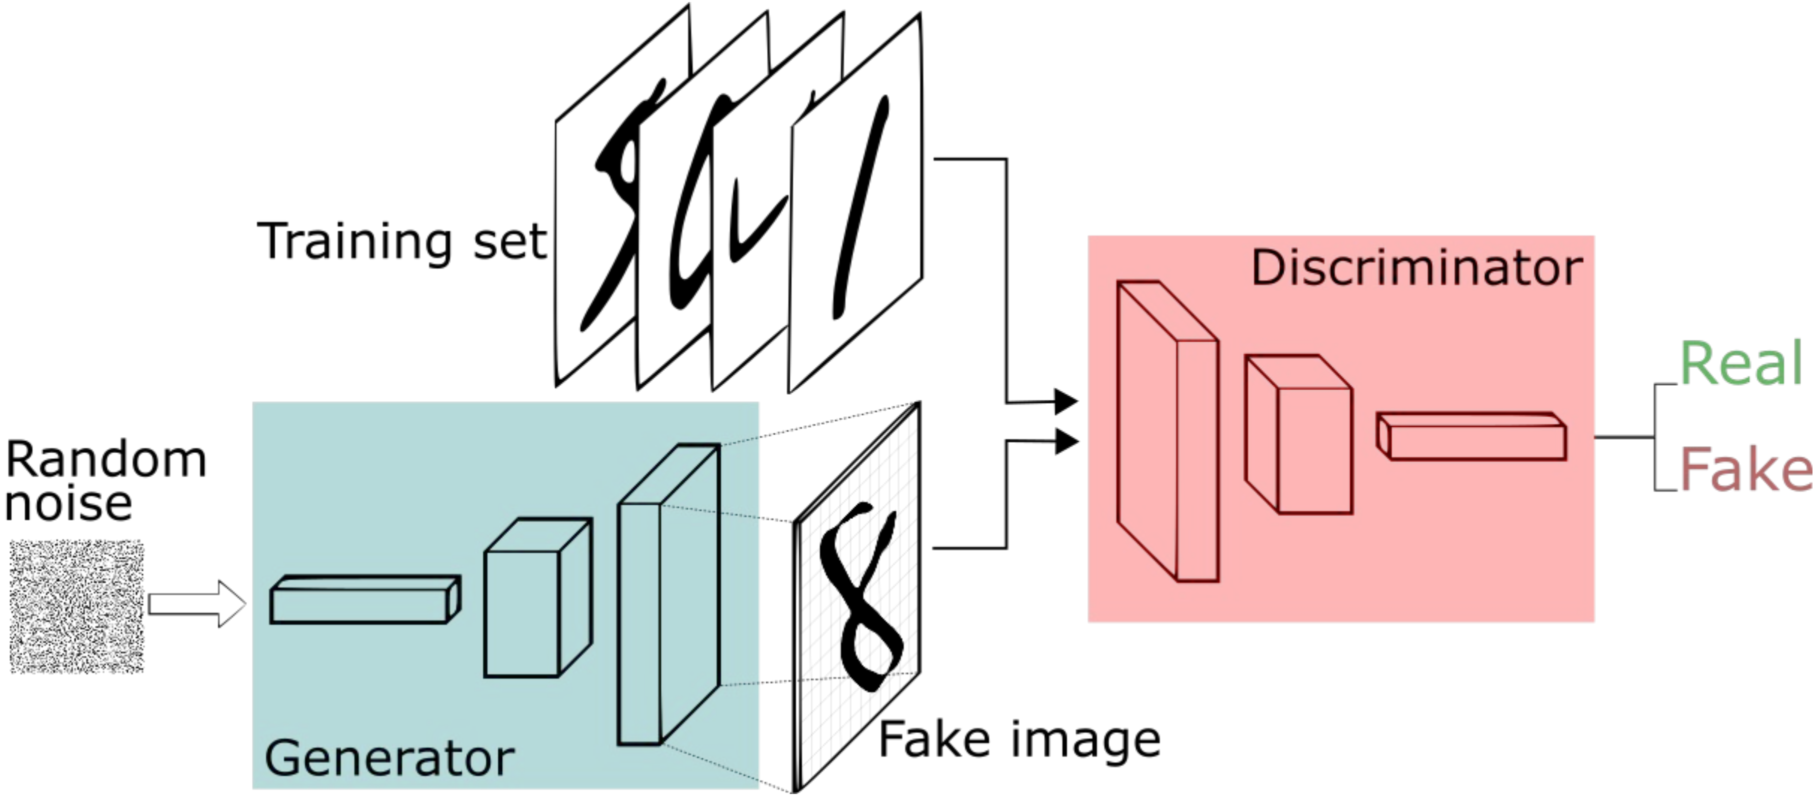

Training Process

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

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

  end for

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

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

end for

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

GAN Objective Functions

GAN Type	Key Take-Away
GAN	The original (JSD divergence)
WGAN	EM distance objective
Improved WGAN	No weight clipping on WGAN
LSGAN	L2 loss objective
RWGAN	Relaxed WGAN framework
McGAN	Mean/covariance minimization objective
GMMN	Maximum mean discrepancy objective
MMD GAN	Adversarial kernel to GMMN
Cramer GAN	Cramer distance
Fisher GAN	Chi-square objective
EBGAN	Autoencoder instead of discriminator
BEGAN	WGAN and EBGAN merged objectives
MAGAN	Dynamic margin on hinge loss from EBGAN

GAN Objective Functions: GANs and Their Variations

Case Study: StyleGAN

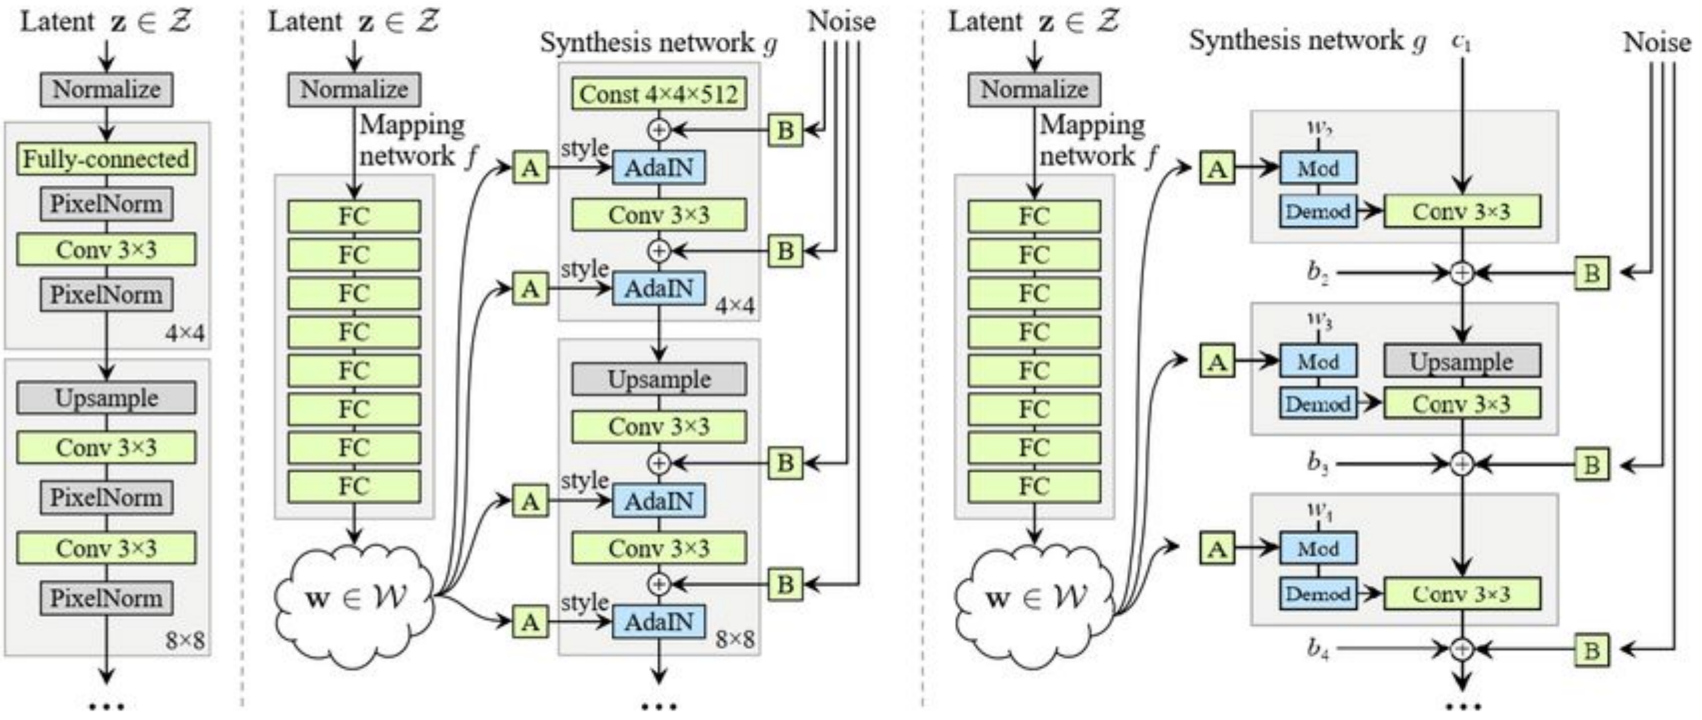

(a) Traditional (b) StyleGAN generator (c) StyleGAN2 generator

Case Study: StyleGAN

Karras T, Laine S, Aila T. A style-based generator architecture for generative adversarial networks. InProceedings of the IEEE/CVF conference on computer vision and pattern recognition 2019 (pp. 4401-4410).

Case Study: StyleGAN

StyleGAN 2 fixed subtle defects generated by StyleGAN.