Generative Adversarial Networks

Deep Learning — Unit 9

Dr. Jon Krohn

jon@untapt.com

December 7th,, 2019

Outline

Outline

Outline

Outline

GANs

Goodfellow et al. (2014)

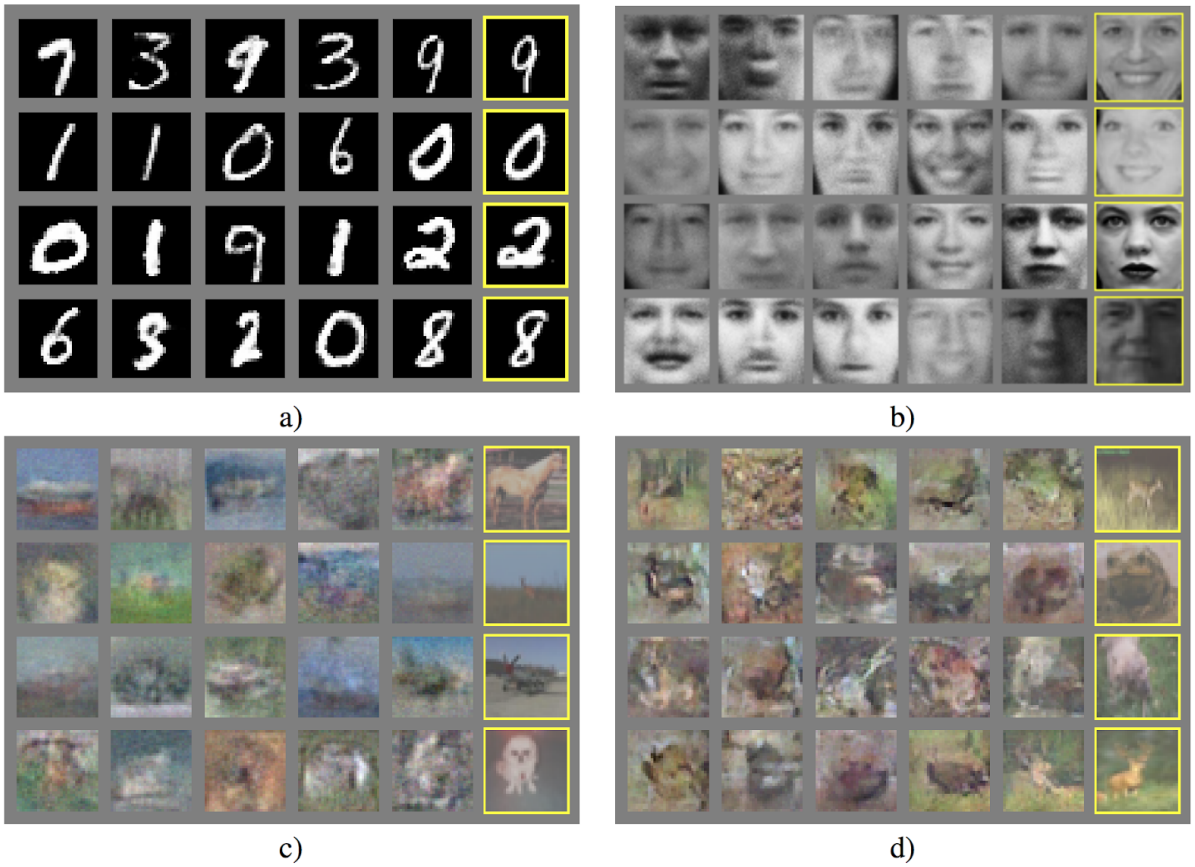

DCGANs

Radford et al. (2016)

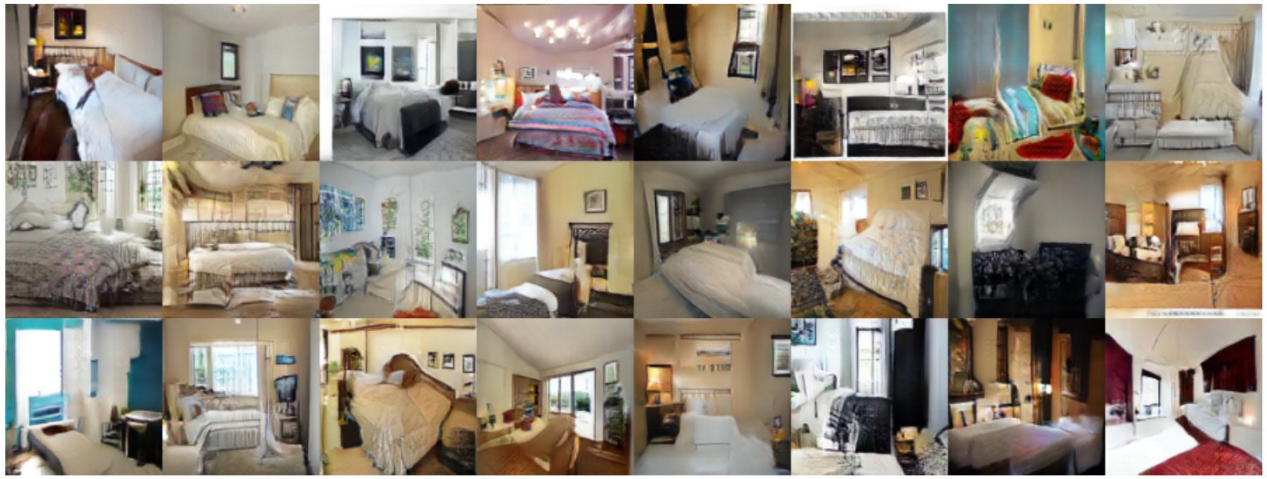

(a) Generated by LSGANs.

(b) Generated by DCGANs (Reported in [13]).

Figure 5: Generated images on LSUN-bedroom.

DCGANs

Radford et al. (2016)

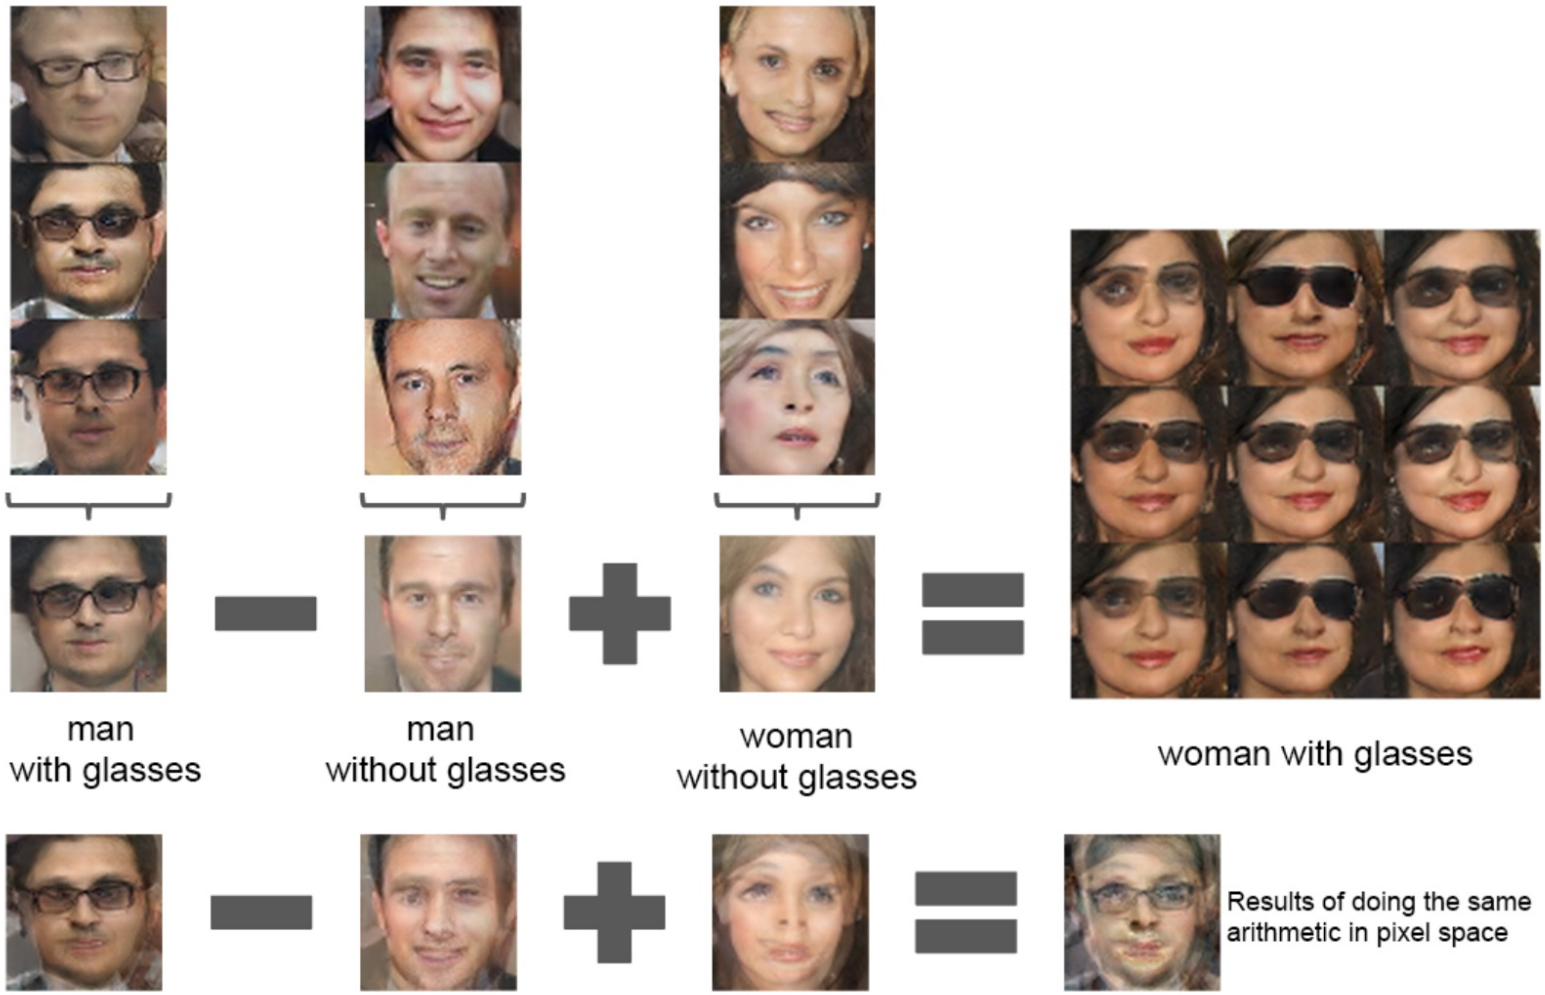

CycleGANs

Zhu et al. (2017)

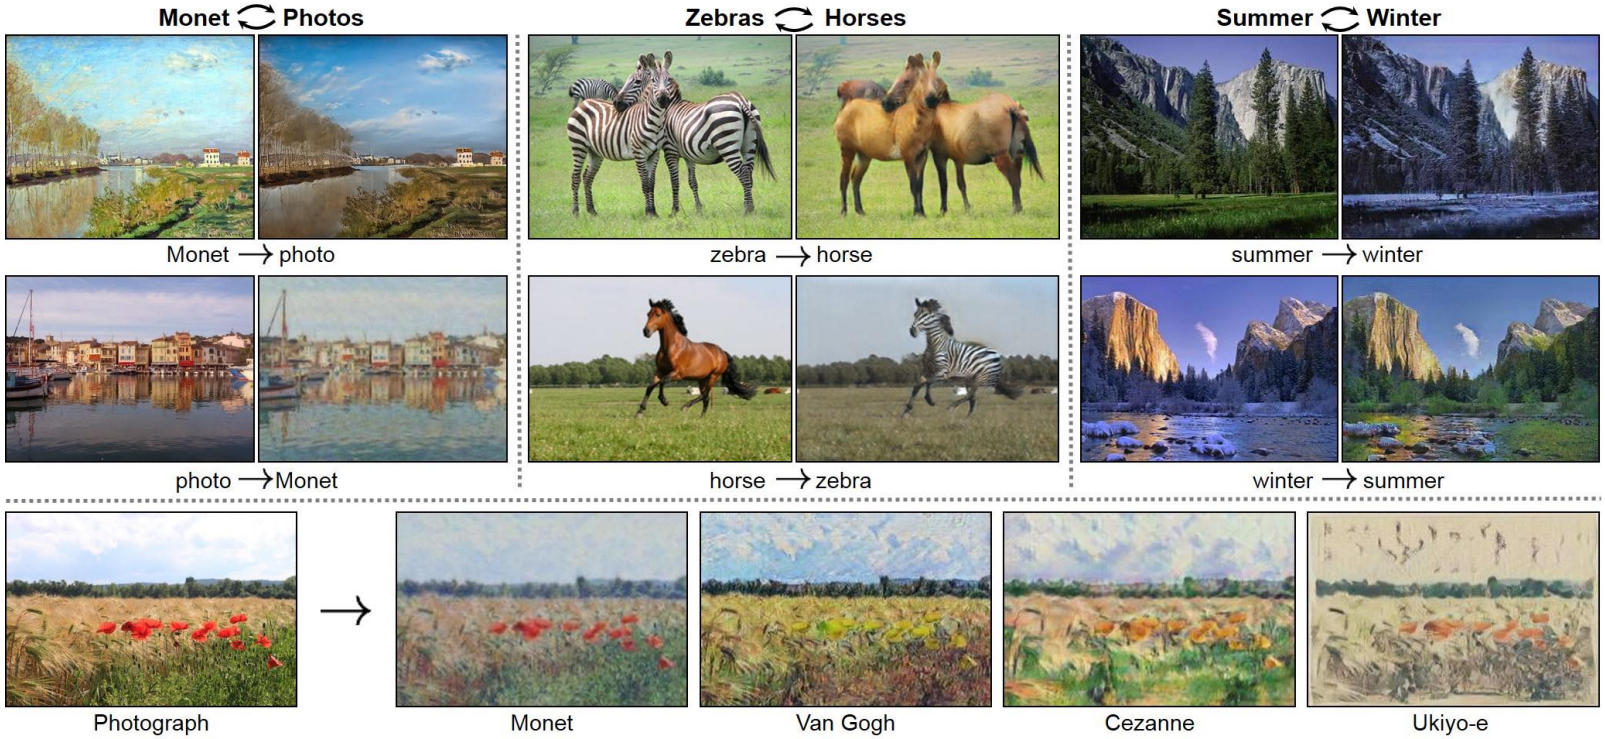

https://junyanz.github.io/CycleGAN

pix2pix

Isola et al. (2017)

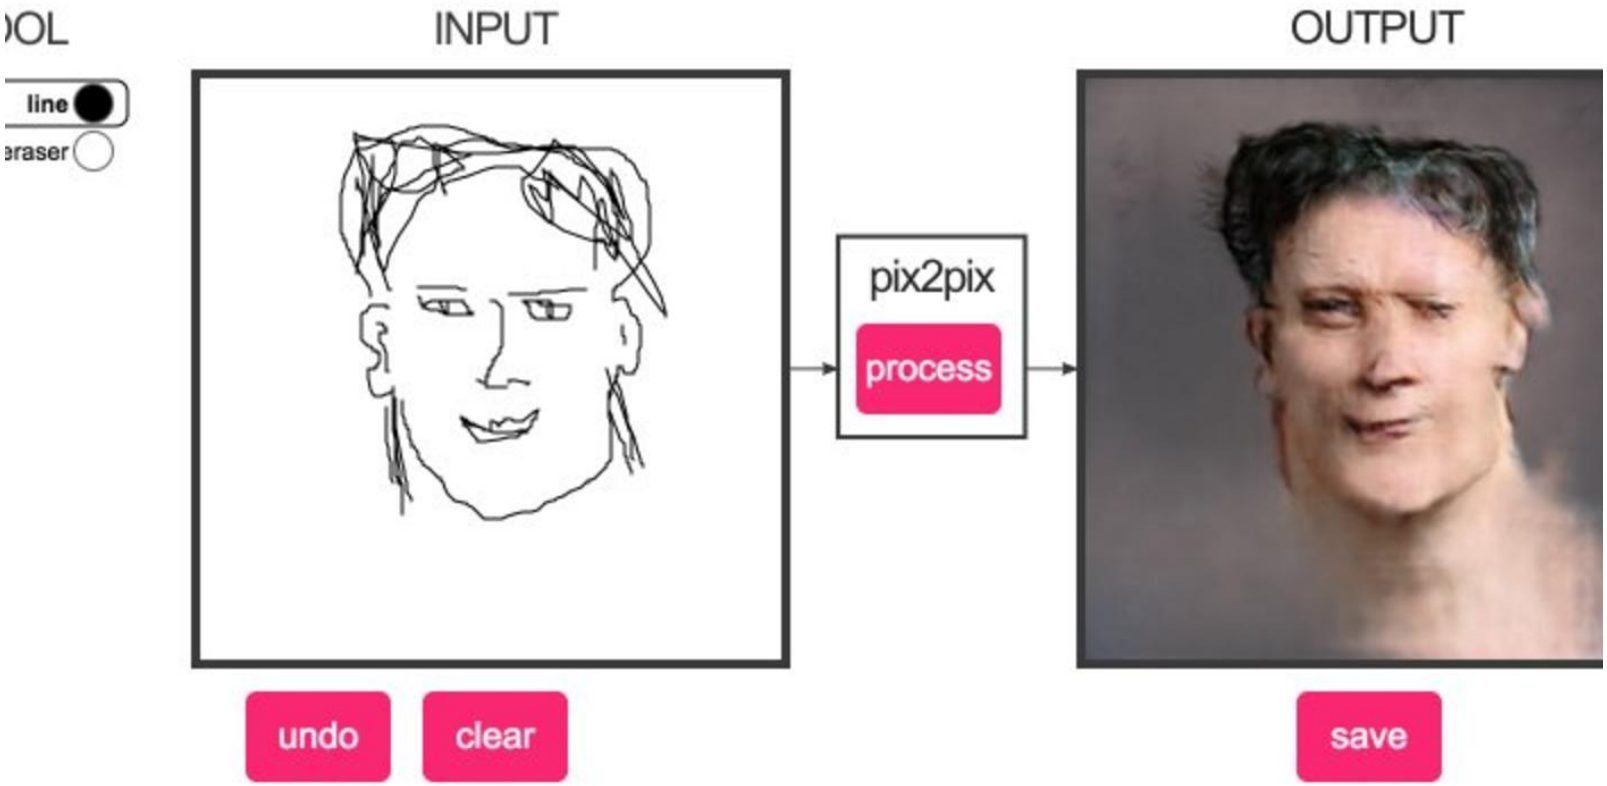

https:

StackGAN

Zhang et al. (2017)

Figure 3. Example results by our proposed StackGAN, GAWWN [20], and GAN-INT-CLS [22] conditioned on text descriptions from CUB test set. GAWWN and GAN-INT-CLS generate 16 images for each text description, respectively. We select the best one for each of them to compare with our StackGAN.

[Which Face is Real?]

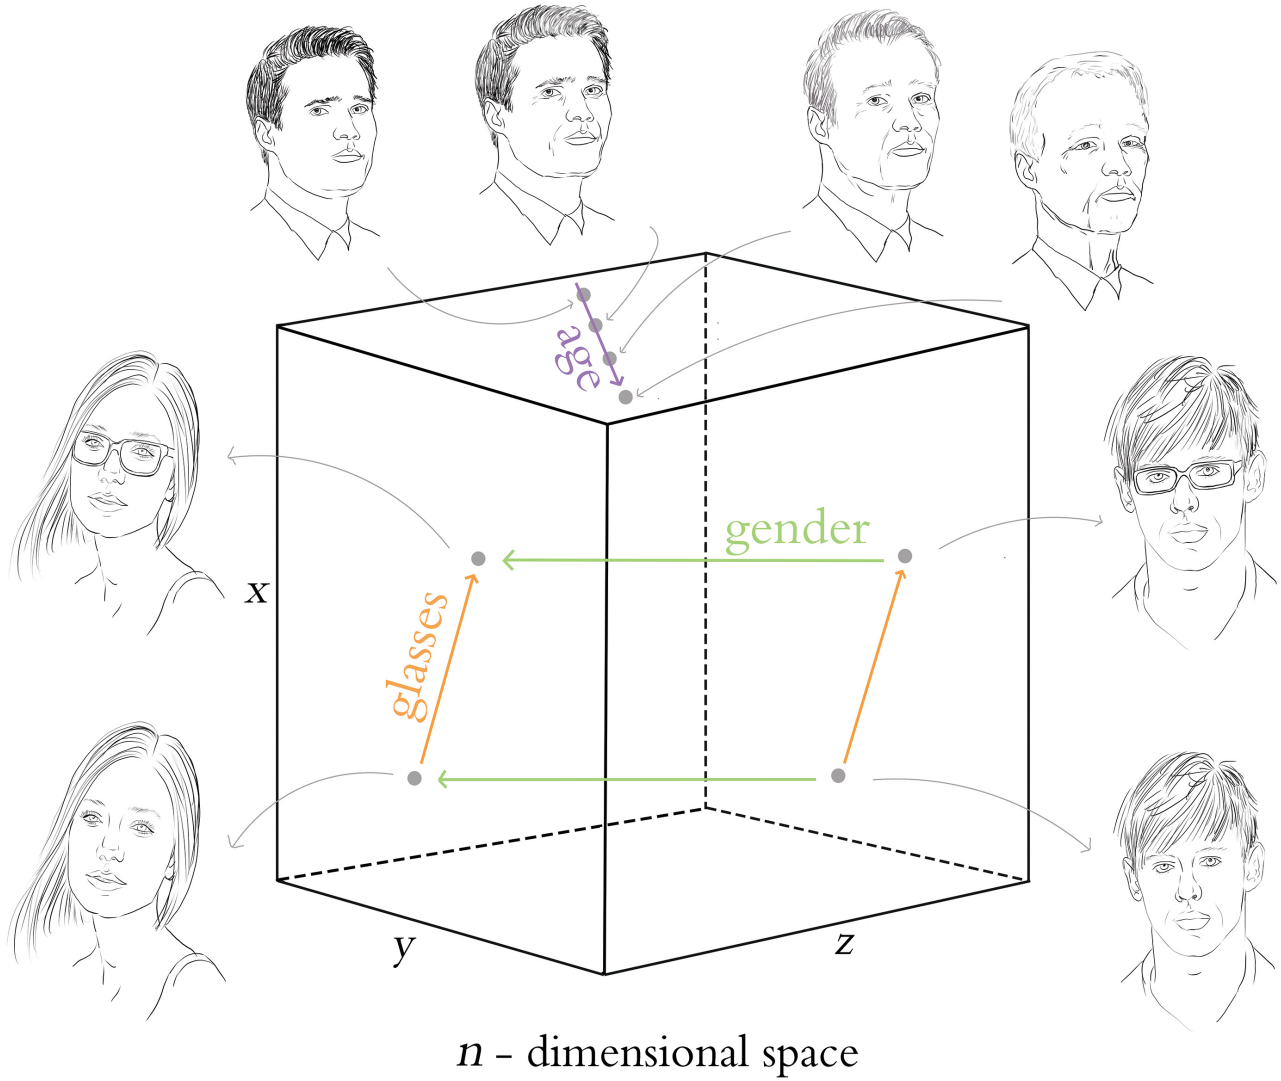

["celebrity" latent-space interpolation]

[Mona Lisa frown]

Latent-Space Interpolation

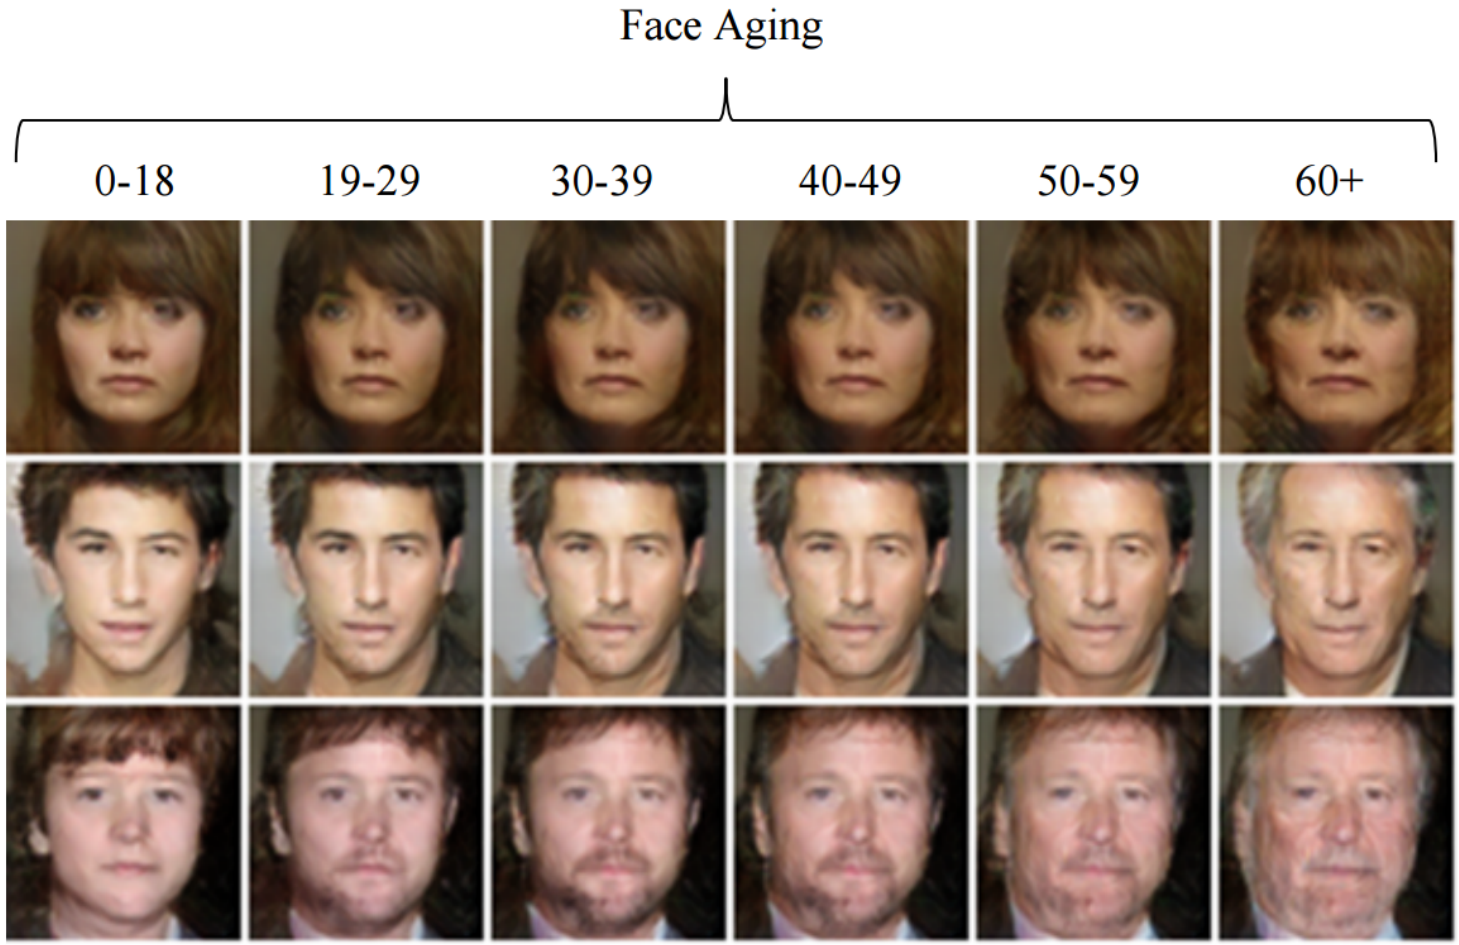

[Ganvatar interactive demo]

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Use Cases

- [make $ selling art :)]
- increase the resolution of an image
- simulate data, e.g., for training autonomous vehicles
- predict next frames of video
- speed fashion/architectural design (sketches to photorealism)
- edit images with realistic, nuanced changes
- [artificial intelligence augmentation (AIA)]
- also can generate time series like text, prices, audio

Outline

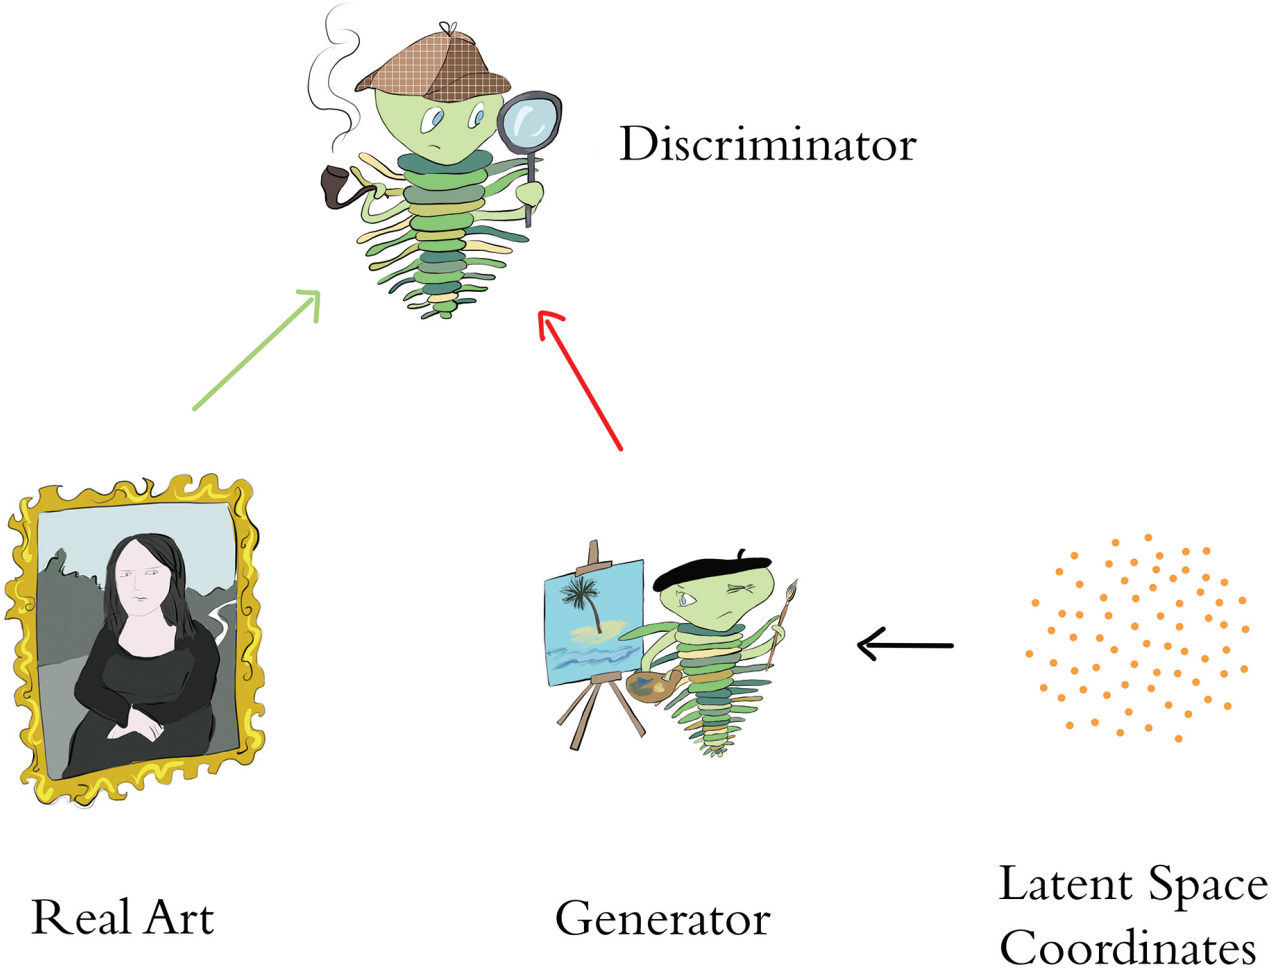
Discriminator
Real Art
Generator
Latent Space Coordinates

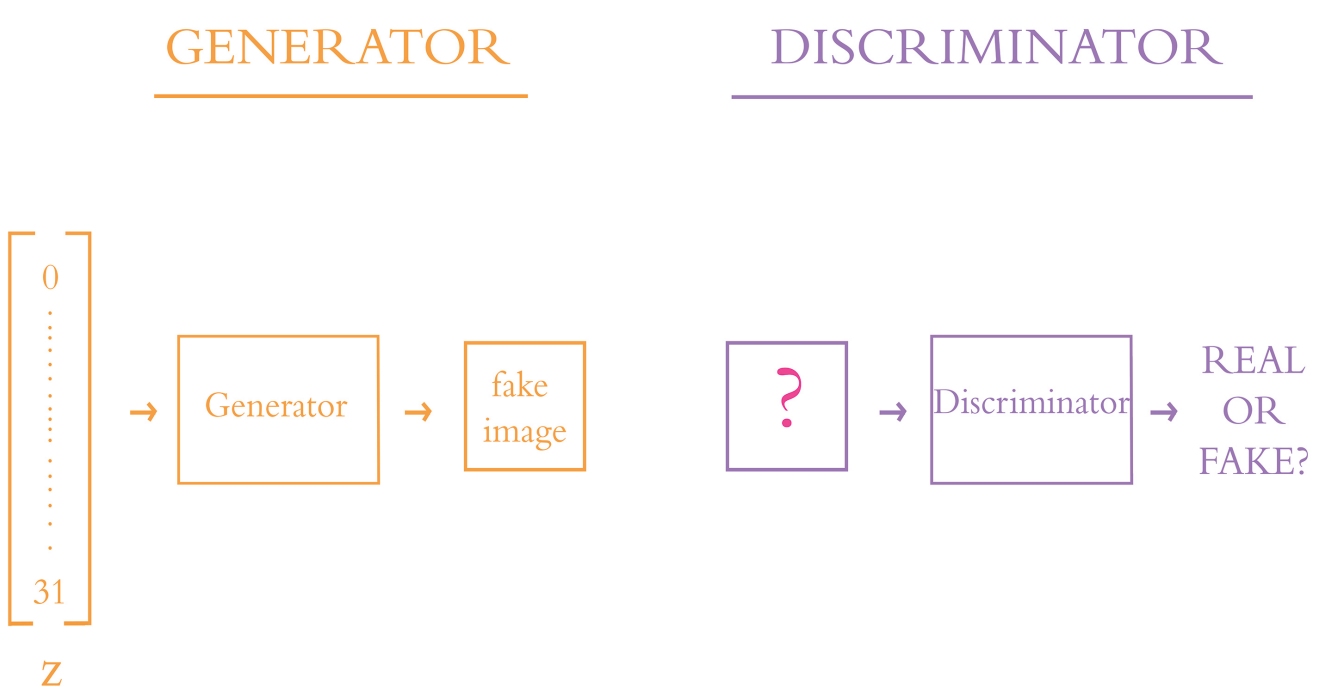

TRAINING THE DISCRIMINATOR

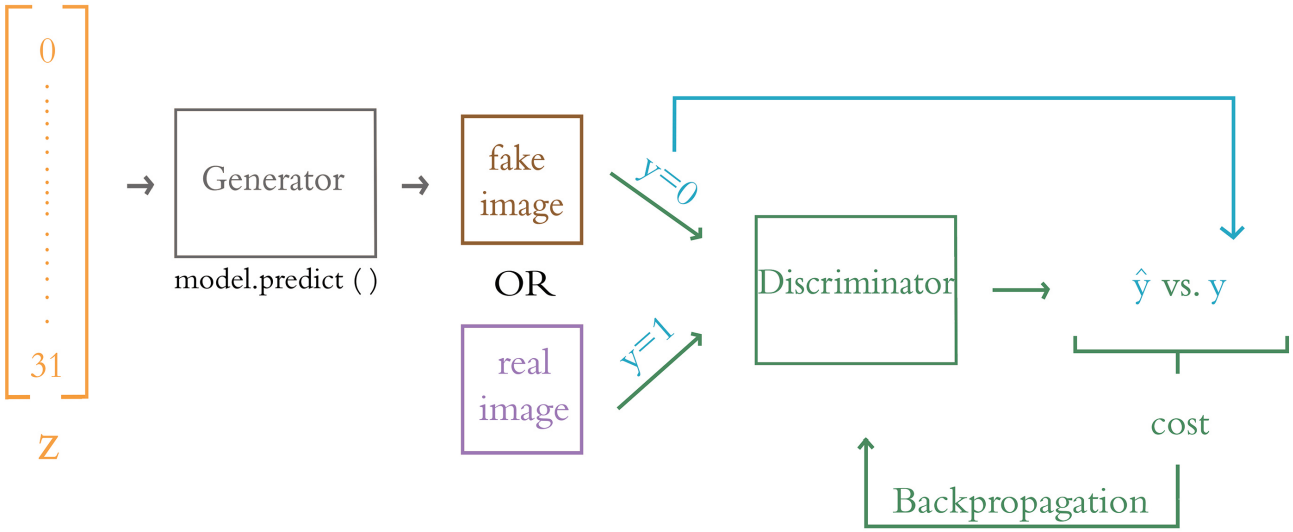

TRAINING THE GENERATOR

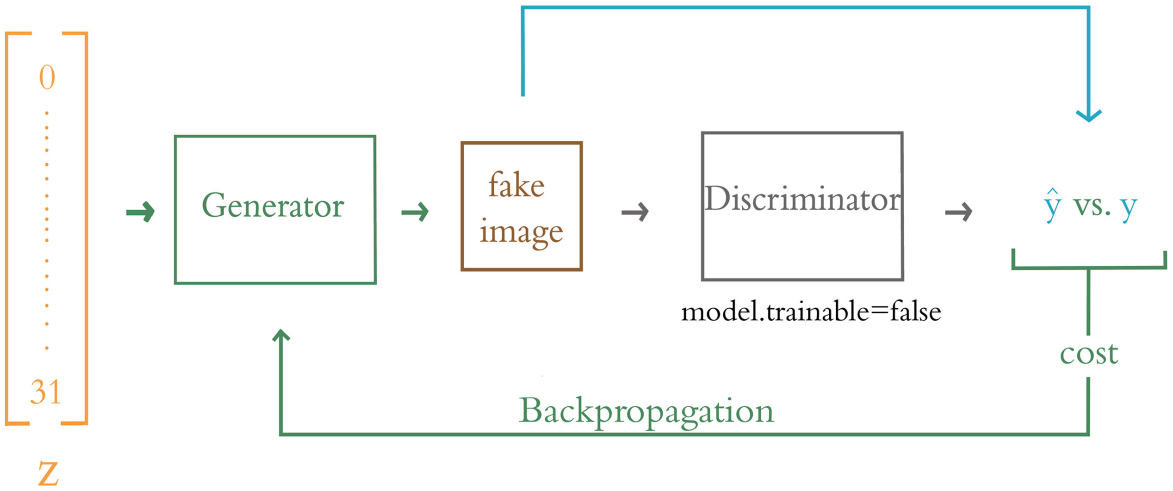

1-D Gaussian

Approximating a Toy Distribution

[video]

Outline

[Quick, Draw!]

GANimation

(Requires Adobe Acrobat Reader)

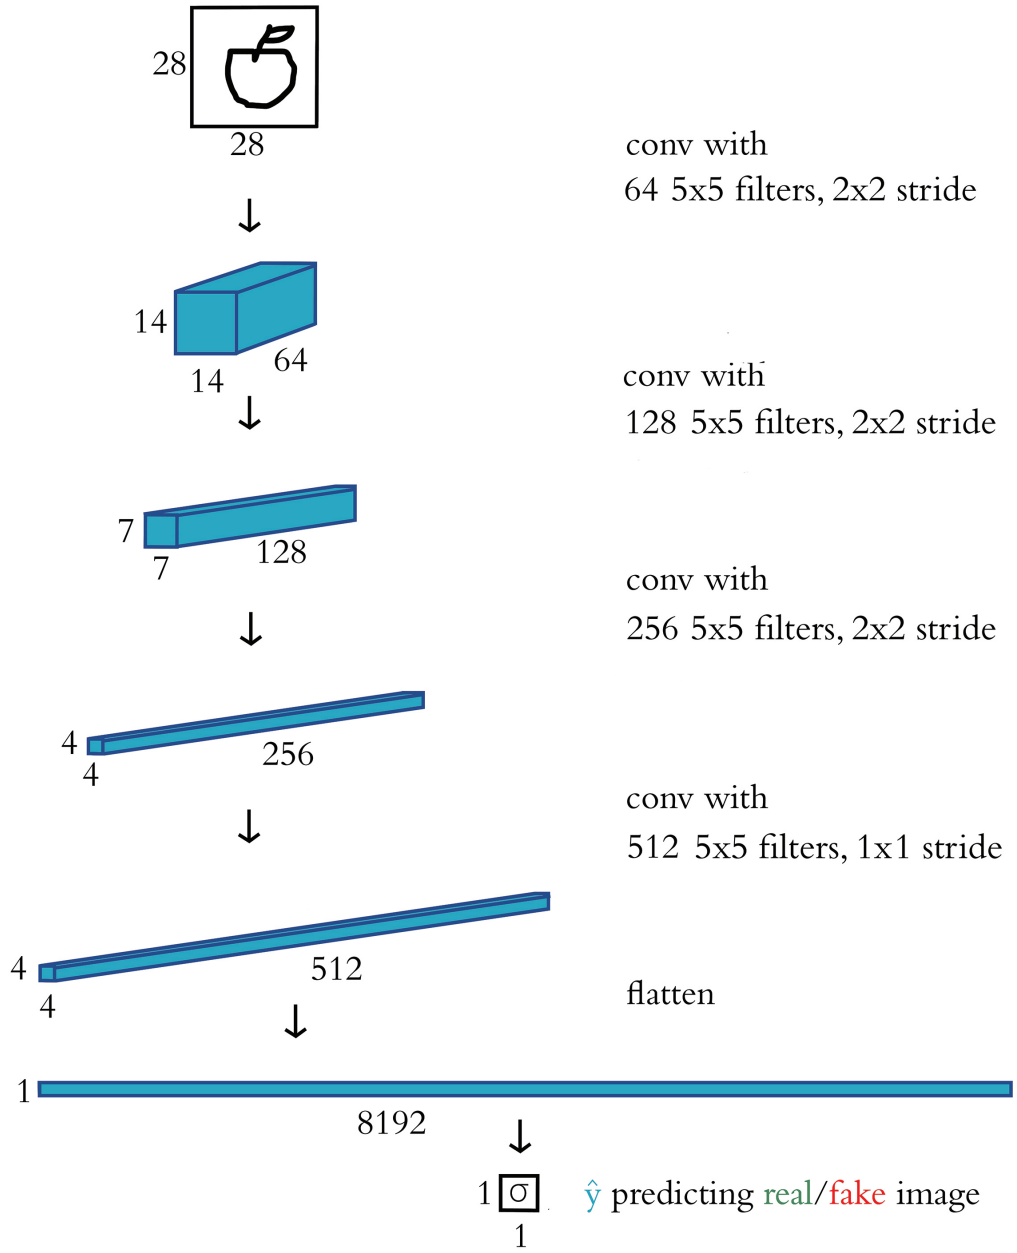

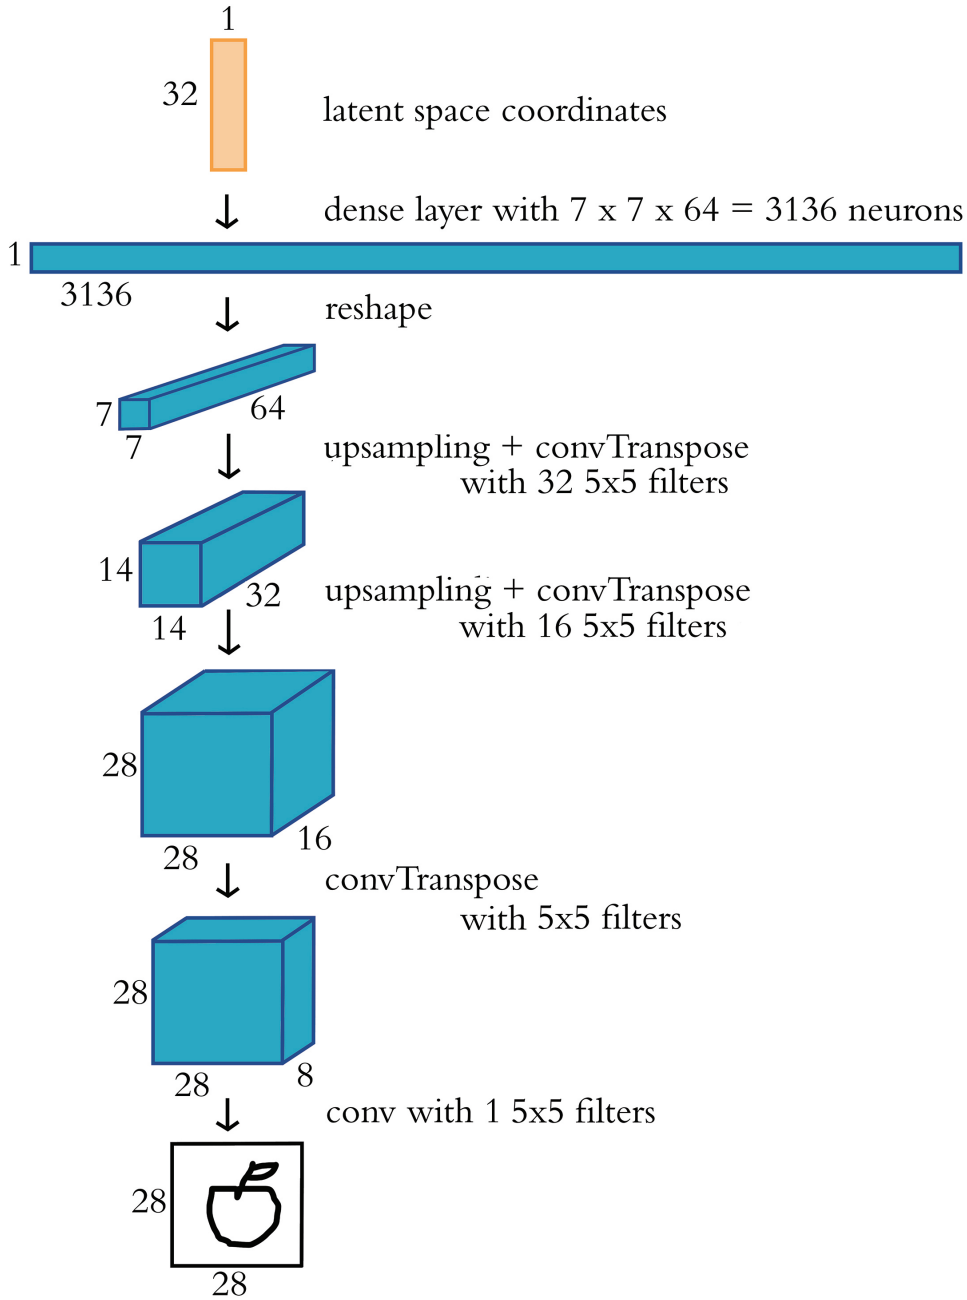

GAN Code

[notebook]