Lecture 8
Generative Adversarial Networks (GANs)

M2 Data Science and AI

Yaohui Wang
http://www-sop.inria.fr/members/Yaohui.Wang/

- Generative Adversarial Networks: *Yaohui Wang*
- DeepFake Detection: *Dr. Antitza Dantcheva*
- Labs (TP): *David Anghelone*

Question: VAE ?

Image Generation

Style Transfer

Video Generation

Outline

- Basic Idea of GAN
- Image Generation
  - Conditional GAN (CGAN, ACGAN)
  - Modern GANs (StyleGAN, BigGAN)
  - Image-to-image translation (Pix2Pix, CycleGAN)
- Video Generation
- GAN interpretablity
- Lab (DCGAN for manga face generation)

Ian Goodfellow

Generative Adversarial Networks [NIPS 2014]


"GANs are the most interesting idea in the last 10 years in ML"
- Yann LeCun

Basic idea of GAN

Basic idea of GAN

$$\begin{bmatrix} 0.3 \\ -0.1 \\ \vdots \\ -0.7 \end{bmatrix} \begin{bmatrix} 0.1 \\ -0.1 \\ \vdots \\ 0.7 \end{bmatrix} \begin{bmatrix} -0.3 \\ 0.1 \\ \vdots \\ 0.9 \end{bmatrix}$$

in a specific range (Gaussian, …)

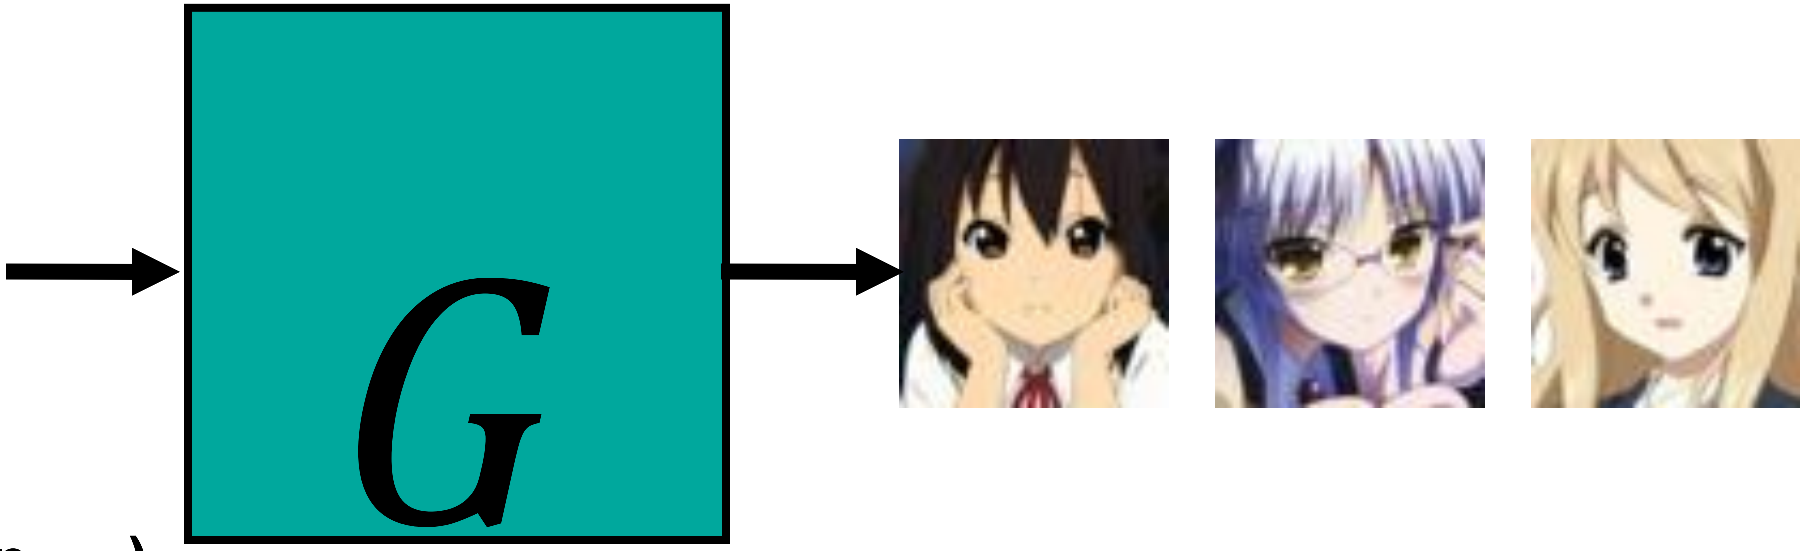

Basic idea of GAN

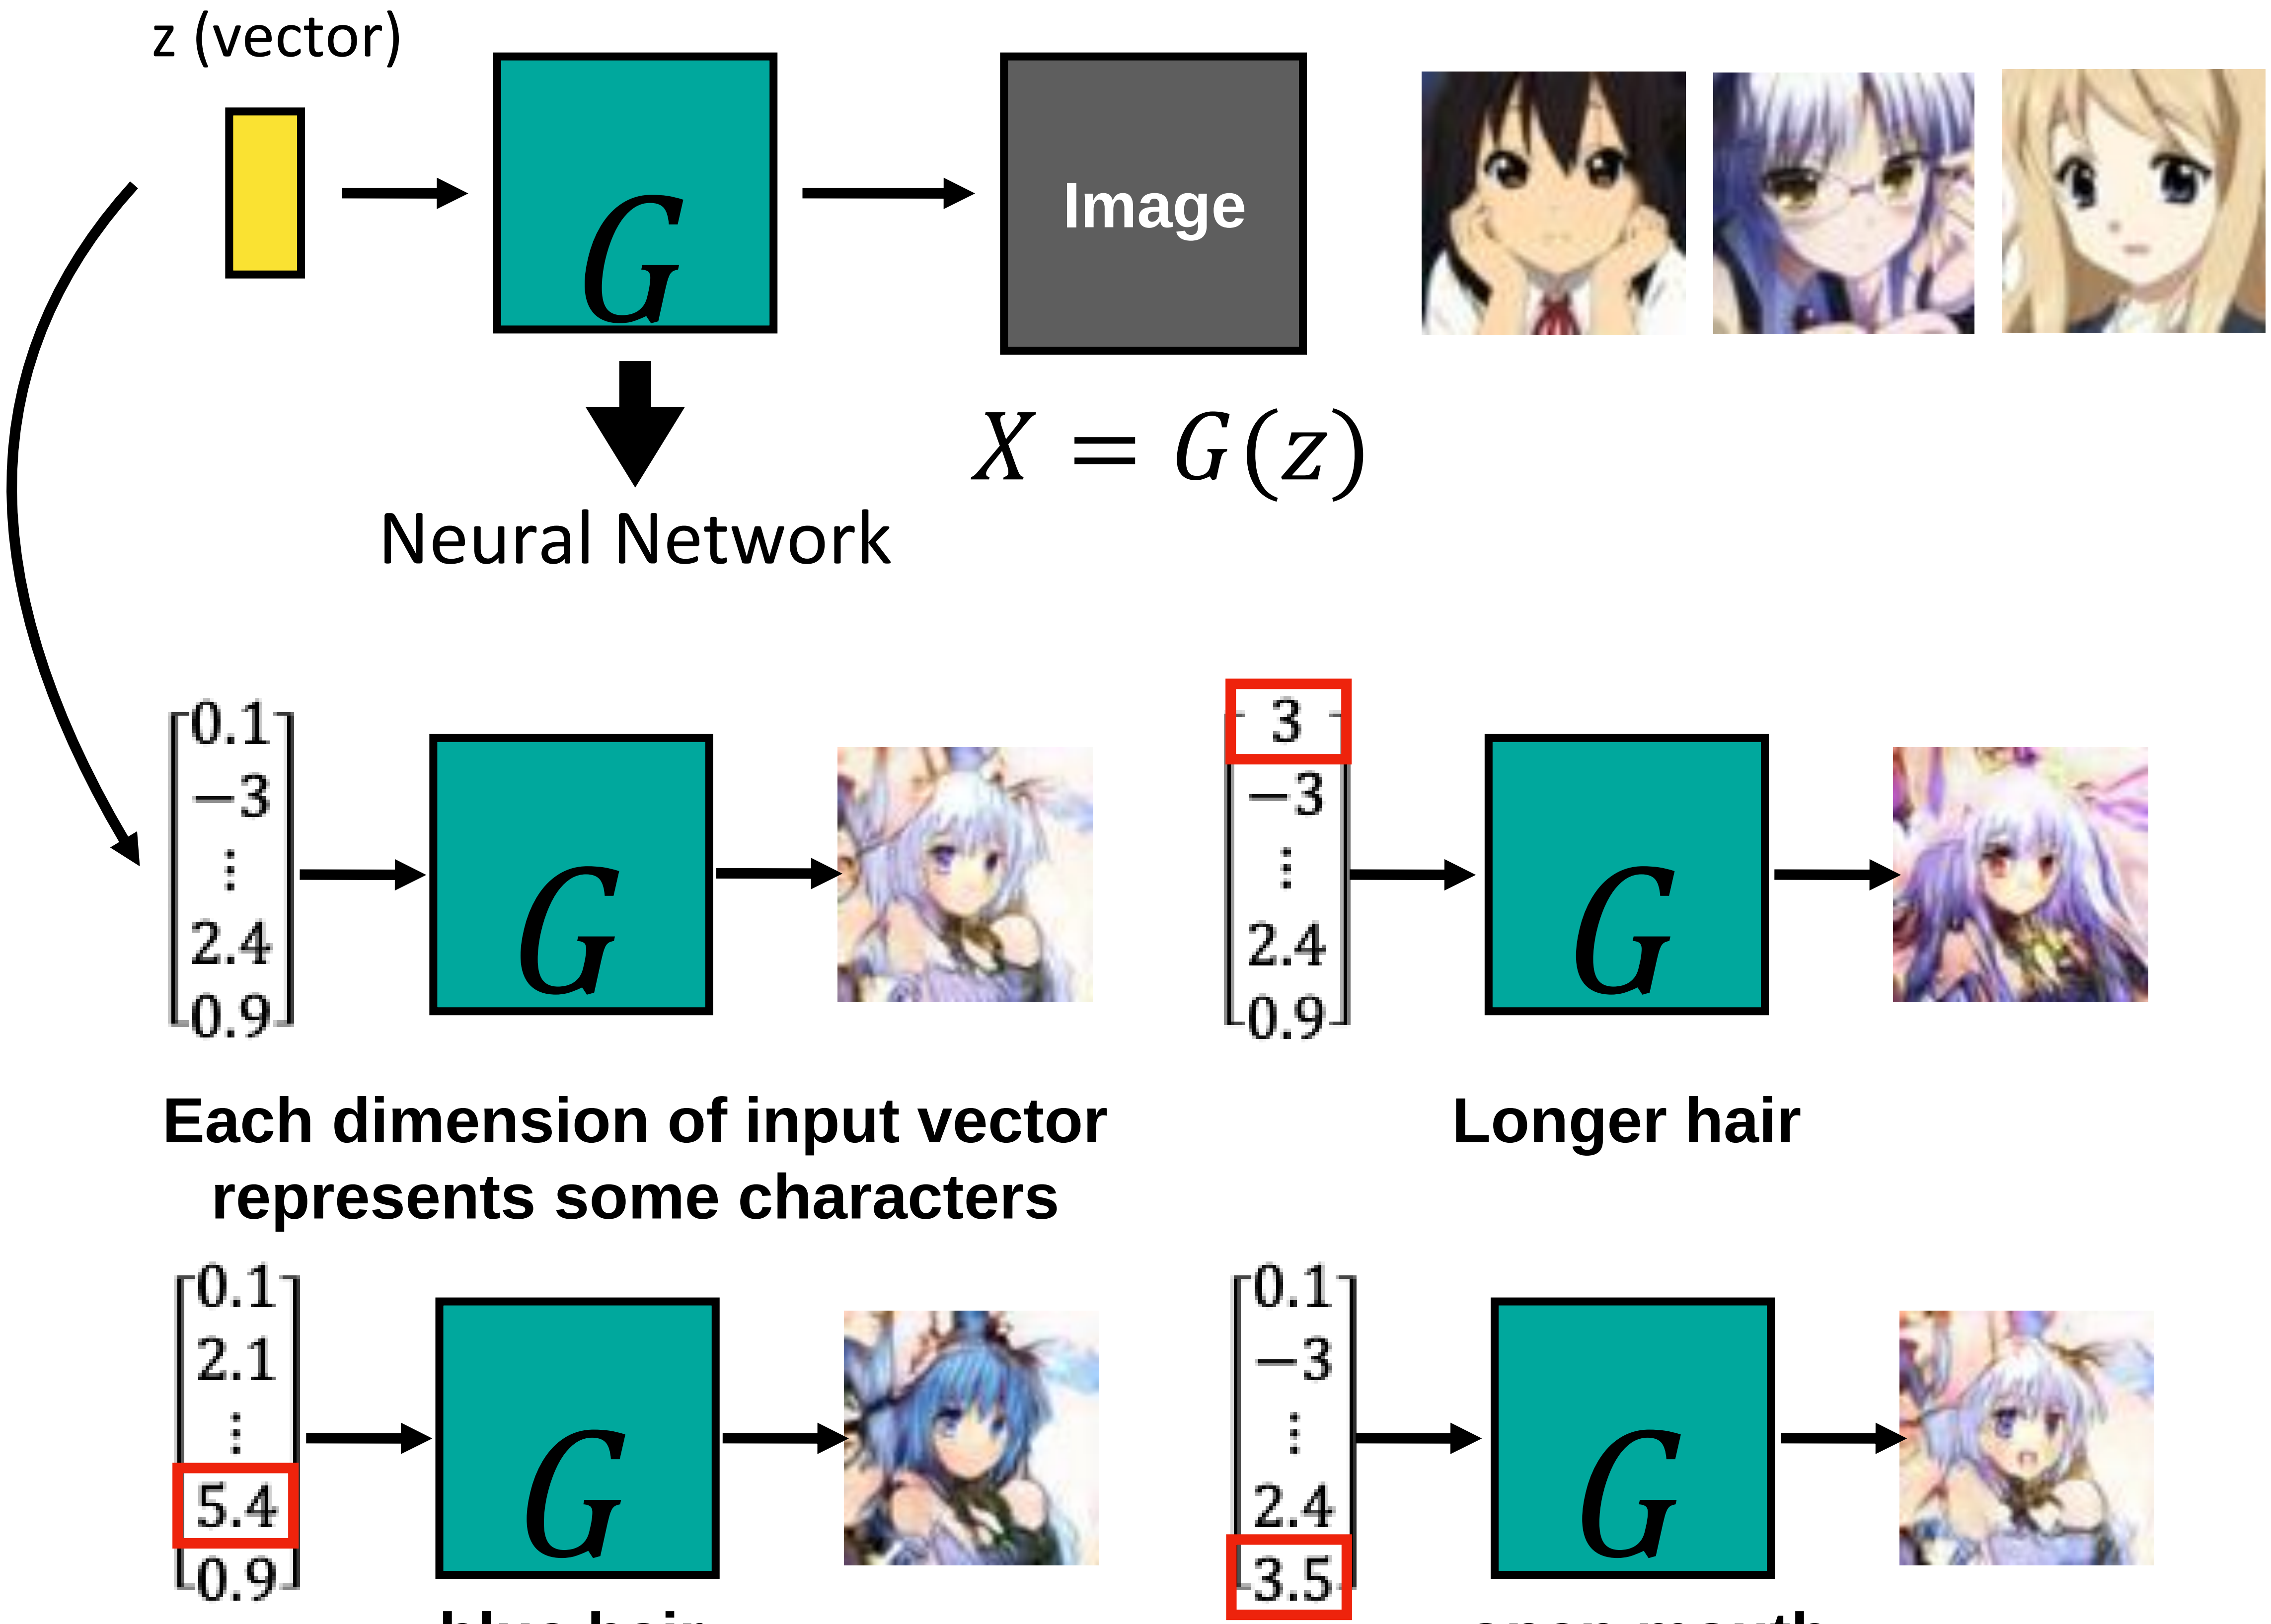

Basic idea of GAN

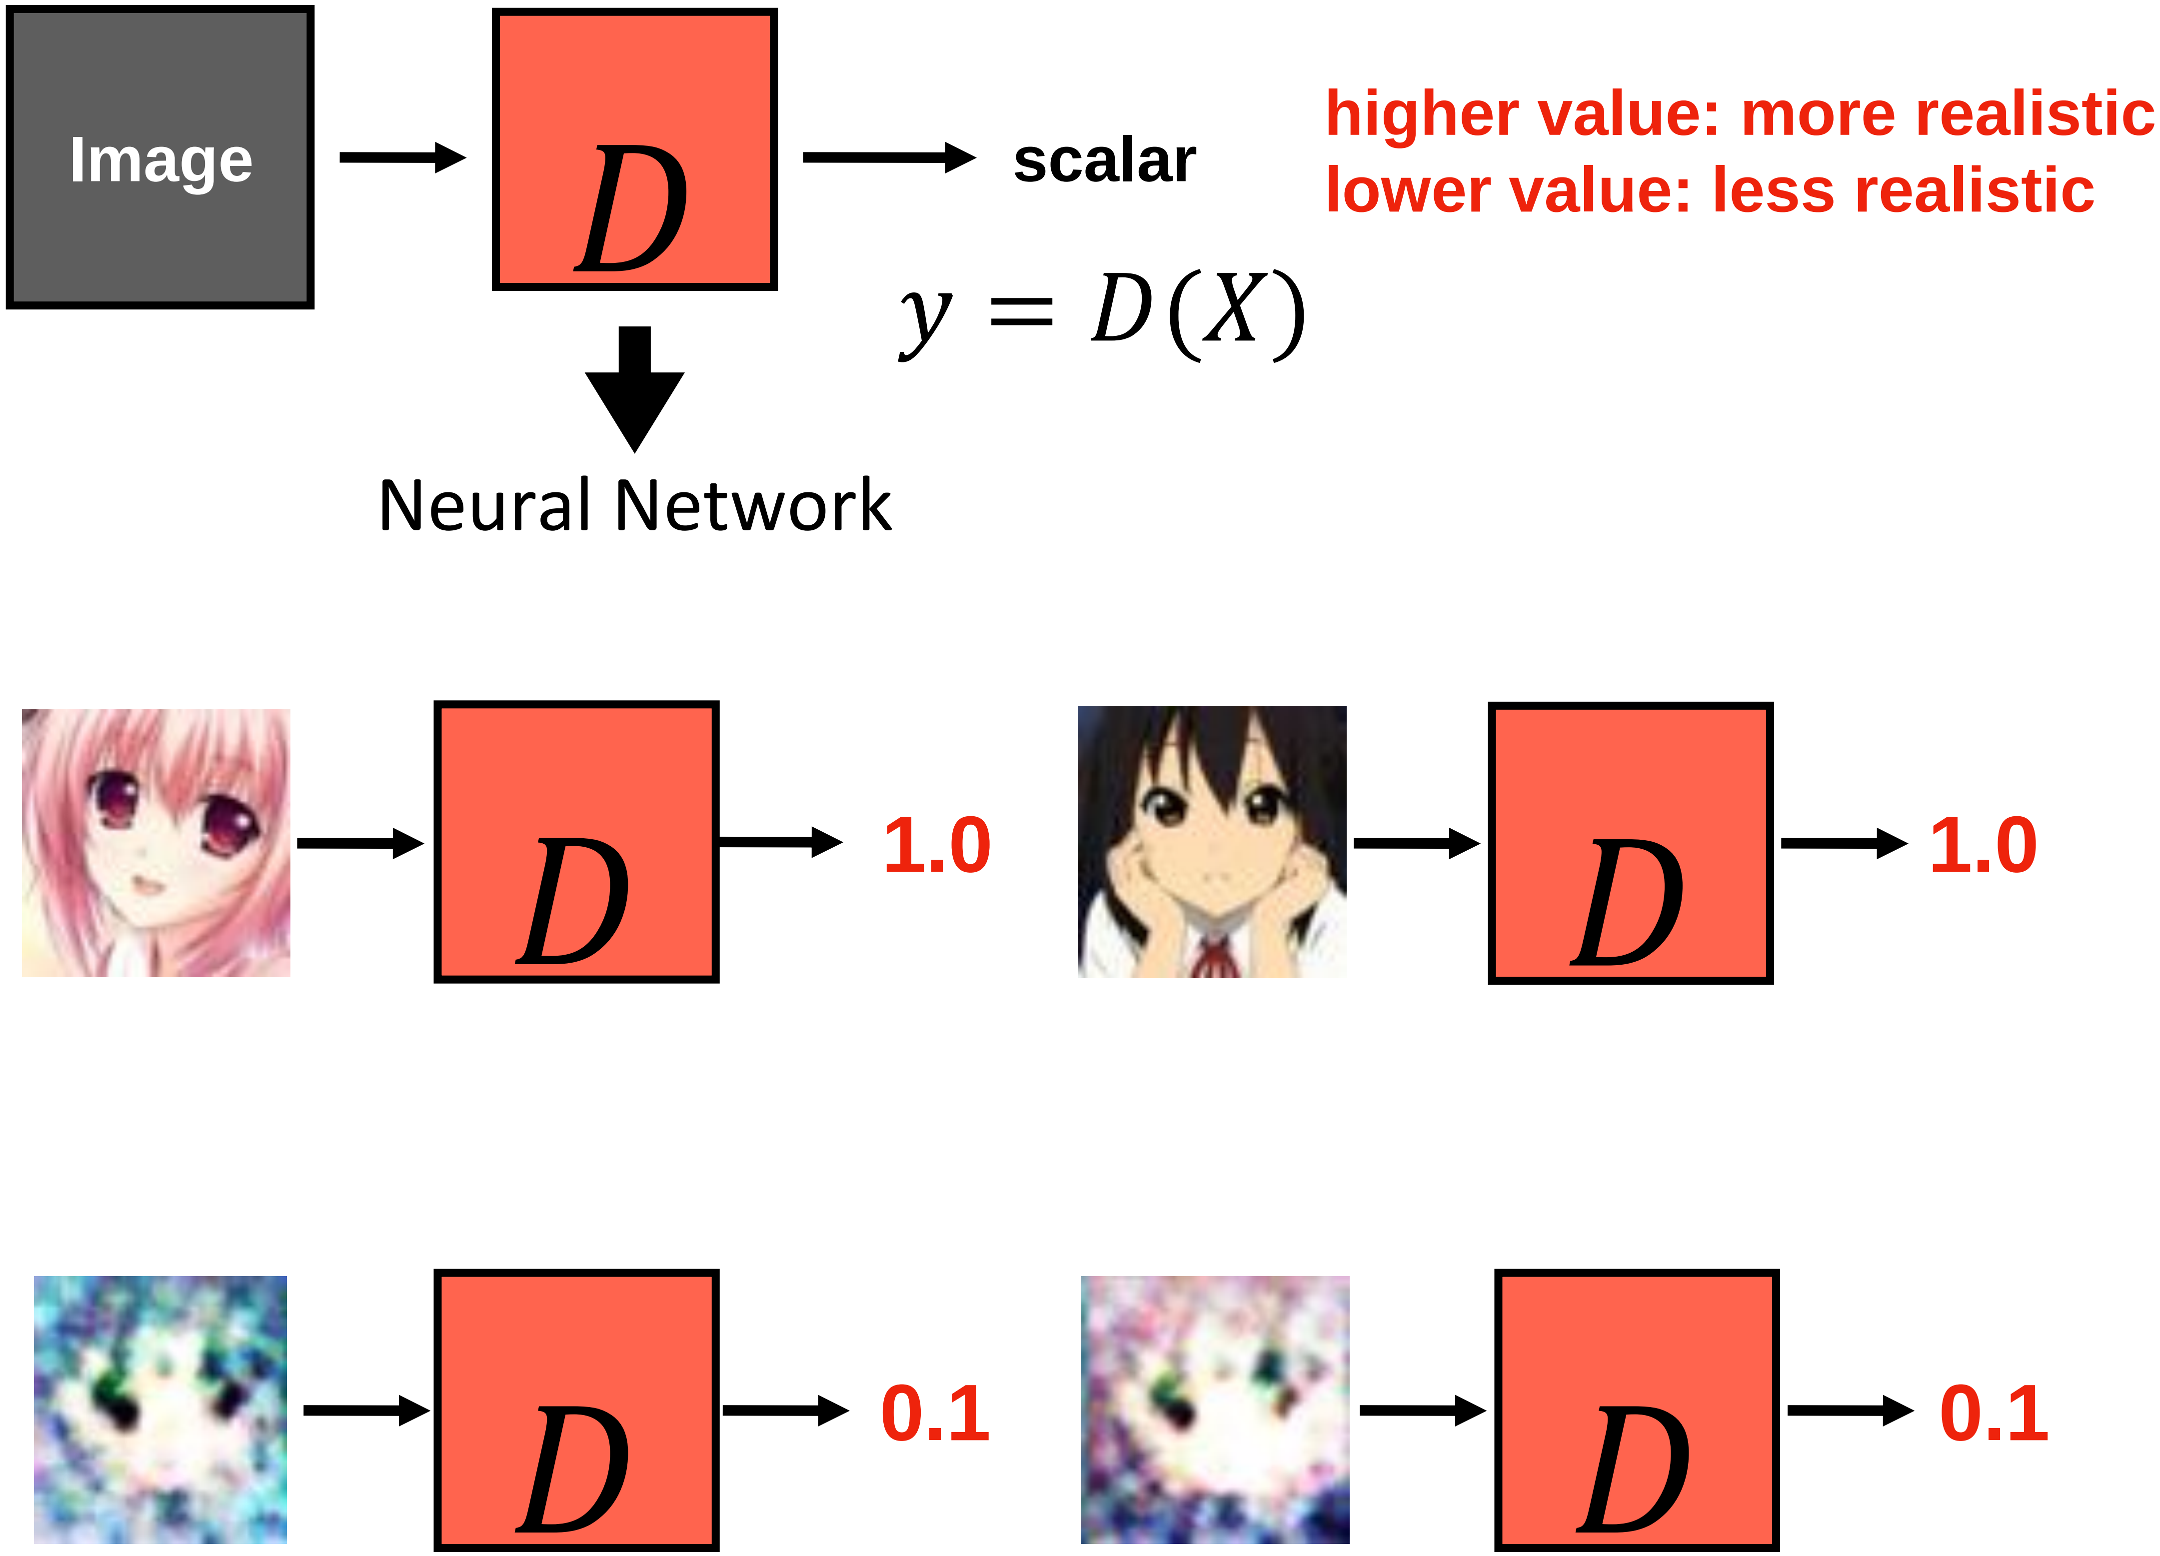

Basic idea of GAN

Adversarial Training (Generative Adversarial Networks)

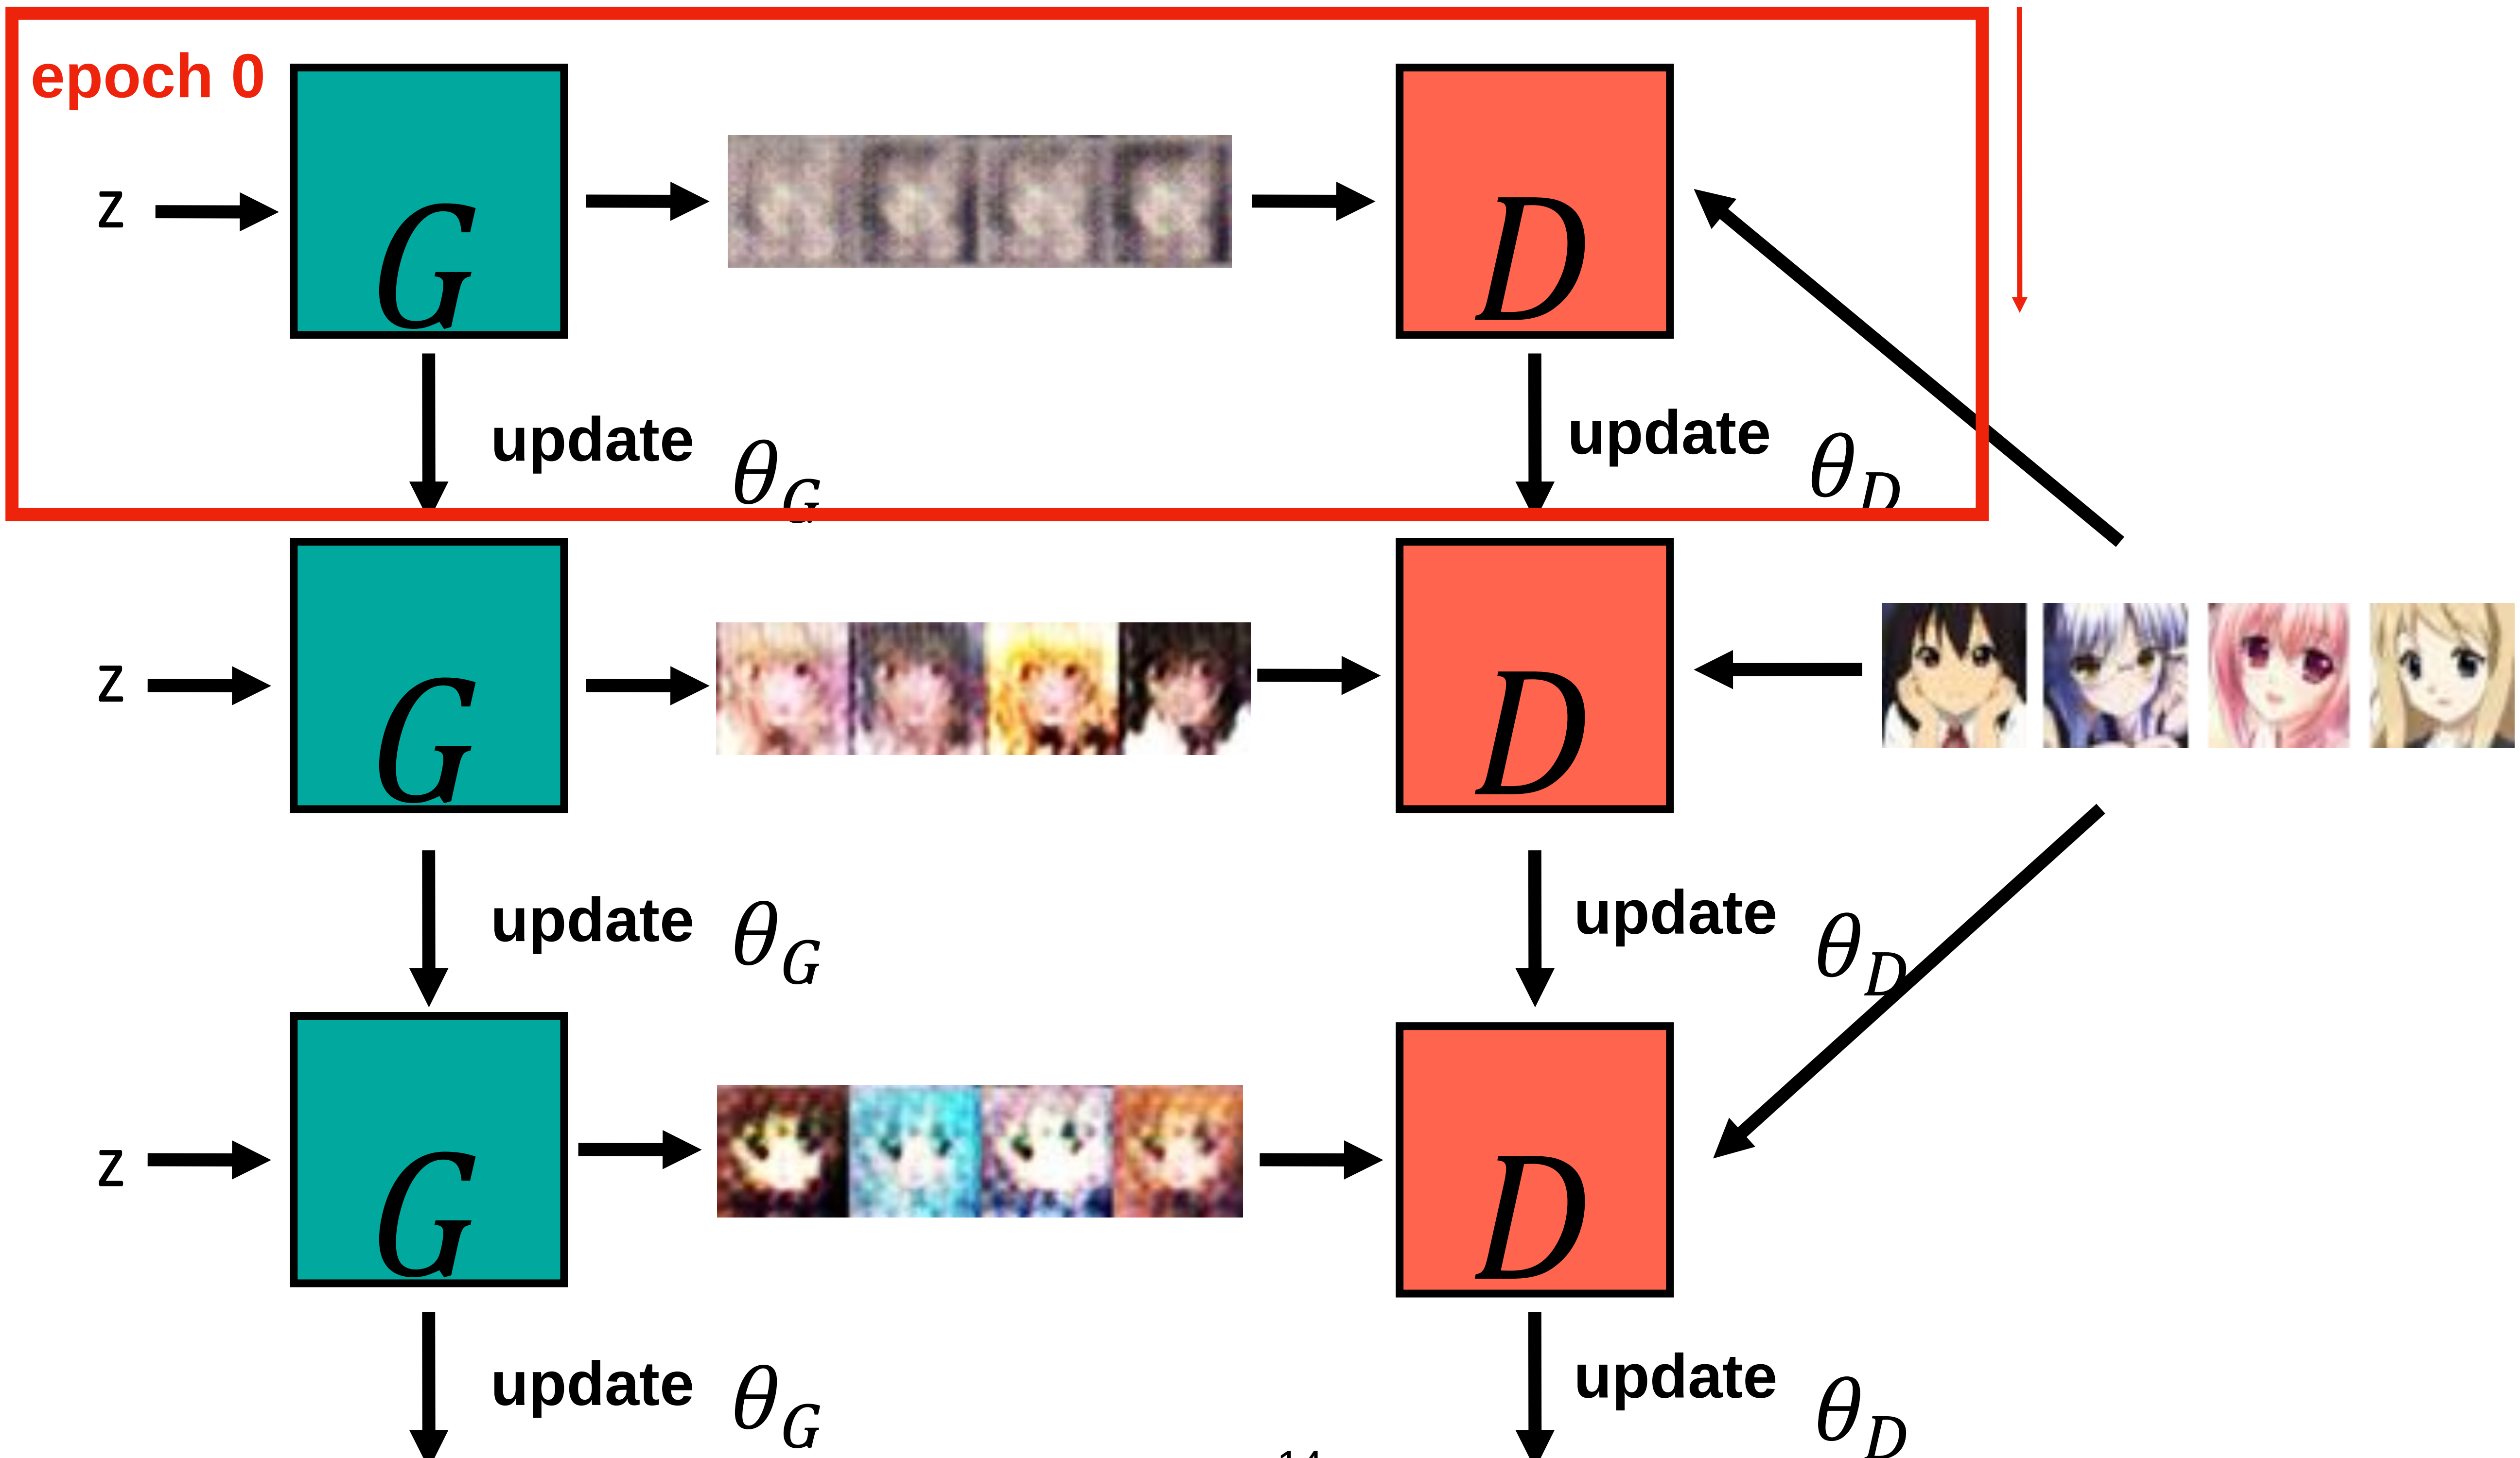

Basic idea of GAN

Adversarial Training (Generative Adversarial Networks)

Algorithm Initialize θ_d for D and θ_g for G

- In each training iteration:

Learning D

  - Sample m examples $\{x^1, x^2, \dots, x^m\}$ from database
  - Sample m noise samples $\{z^1, z^2, \dots, z^m\}$ from a distribution
  - Obtaining generated data $\{\tilde{x}^1, \tilde{x}^2, \dots, \tilde{x}^m\}$, $\tilde{x}^i = G(z^i)$
  - Update discriminator parameters θ_d to maximize
    - $\tilde{V} = \frac{1}{m}\sum_{i=1}^{m} log D(x^i) + \frac{1}{m}\sum_{i=1}^{m} log\left(1 - D(\tilde{x}^i)\right)$
    - $\theta_d \leftarrow \theta_d + \eta \nabla \tilde{V}(\theta_d)$

Learning G

  - Sample m noise samples $\{z^1, z^2, \dots, z^m\}$ from a distribution
  - Update generator parameters θ_g to maximize
    - $\tilde{V} = \frac{1}{m}\sum_{i=1}^{m} log\left(D\left(G(z^i)\right)\right)$
    - $\theta_g \leftarrow \theta_g - \eta \nabla \tilde{V}(\theta_g)$

Basic idea of GAN

Generator: G is a network. It defines a probability distribut P_G

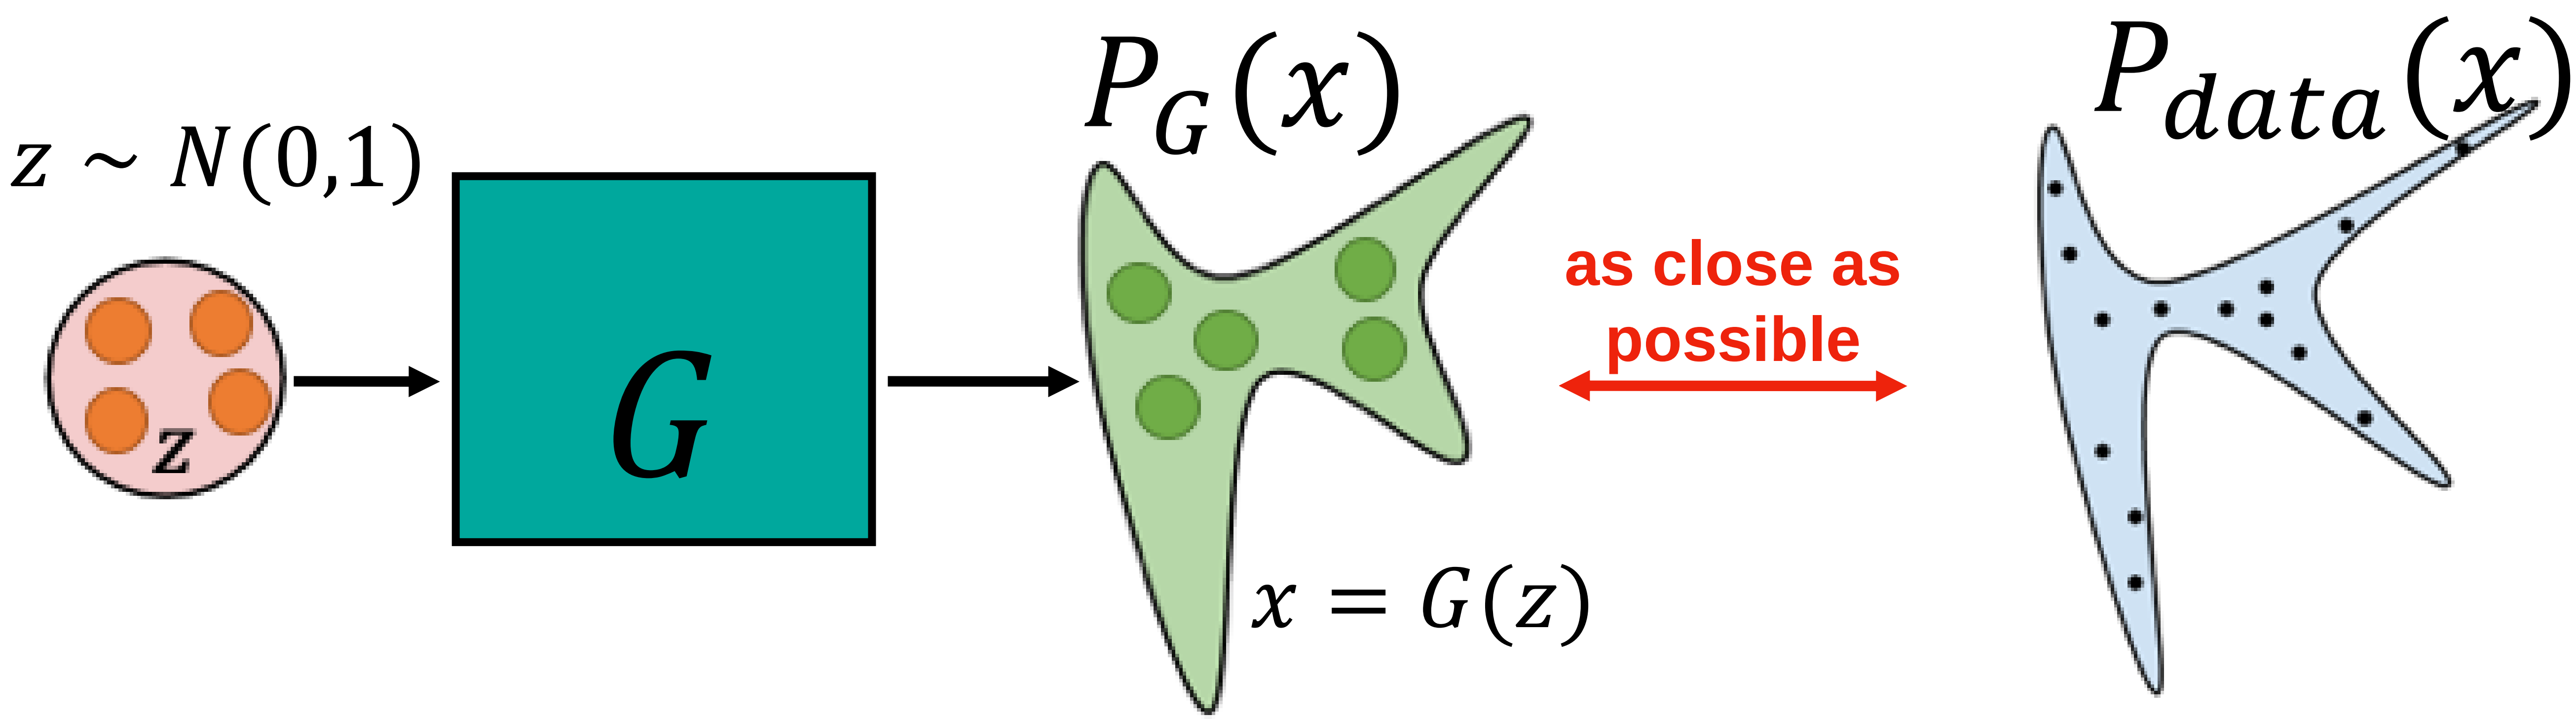

$$G^* = \arg\min_G \boxed{Div(P_G, P_{data})}$$

how to compute the divergence between two distributions ?

Basic idea of GAN

Discriminator $$G^* = \arg\min_G Div(P_G, P_{data})$$

Although we do not know the distributions of $P_G(x)$ and $P_{data}(x)$, we can still sample from them

real images

sampling from $P_{data}(x)$

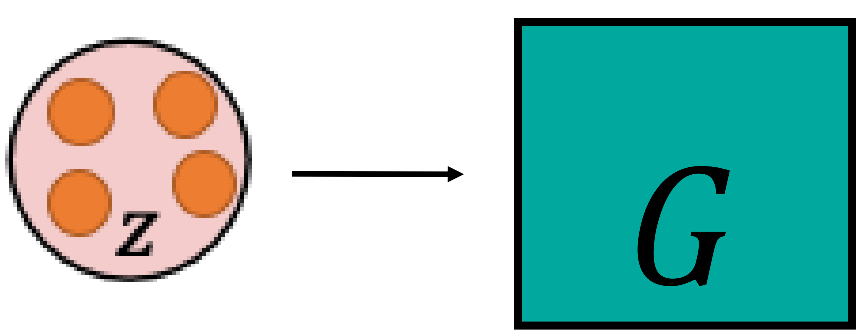

sample from normal distribution

sampling from $P_G(x)$

Basic idea of GAN

Discriminator $G^* = arg\min_G Div(P_G, P_{data})$

Objective function for D

$$V(G,D) = E_{x\sim P_{data}}[logD(x)] + E_{x\sim P_G}[log(1 - D(x))]$$

(G is fixed)

$$D^* = arg\boxed{\max_D V(G,D)}$$

= binary classification

JS Divergence

Basic idea of GAN

Discriminator $G^* = \underset{G}{argmin} Div(P_G, P_{data})$

Objective function for G

$$G^* = \underset{G}{argmin}(\cancel{E_{x\sim P_{data}}[logD(x)]} + \underline{E_{x\sim P_G}[log(1 - D(G(z))}$$

(D is fixed)

$$E_{x\sim P_G}[-log(D(G(z)))])$$

Basic idea of GAN

$$E_{x \sim P_G}[log(1 - D(G(z)))])$$

slow at the beginning

$$E_{x \sim P_G}[-log(D(G(z)))])$$

real implementation

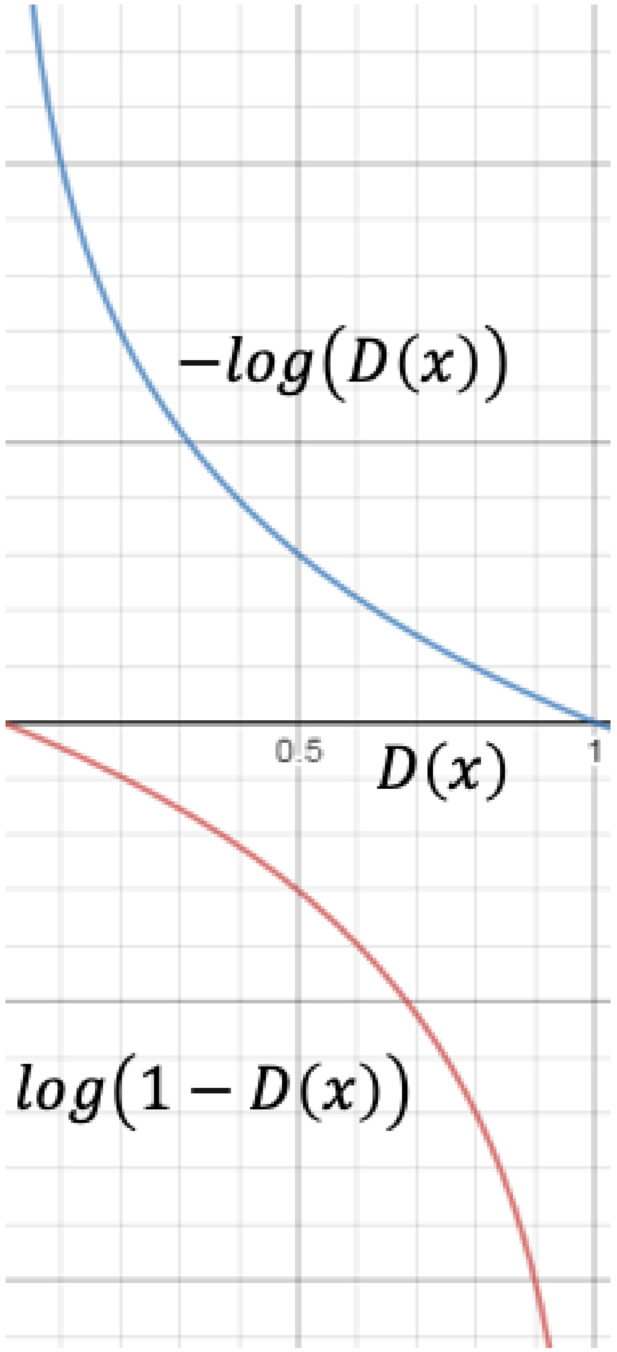

Basic idea of GAN

Different GANs

- Wasserstein GAN
- Wasserstein GAN-GP (gradient penalty)
- LSGAN
- …

Basic idea of GAN

$$V(G,D) = E_{x\sim P_{data}}[logD(x)] + E_{Z\sim P_Z}[log(1 - D(G(z)))$$

$$G^* = arg\min_G \max_D V(G,D)$$

Training Steps:

- Initialize Generator and Discriminator
- In each training iteration:

  Step 1: Fix Generator G, and update Discriminator D

  Step 2: Fix Discriminator D, and update Generator G

Vanilla GAN (unconditional)

Vanilla GAN [Ian Goodfellow, et al, NIPS 2014]

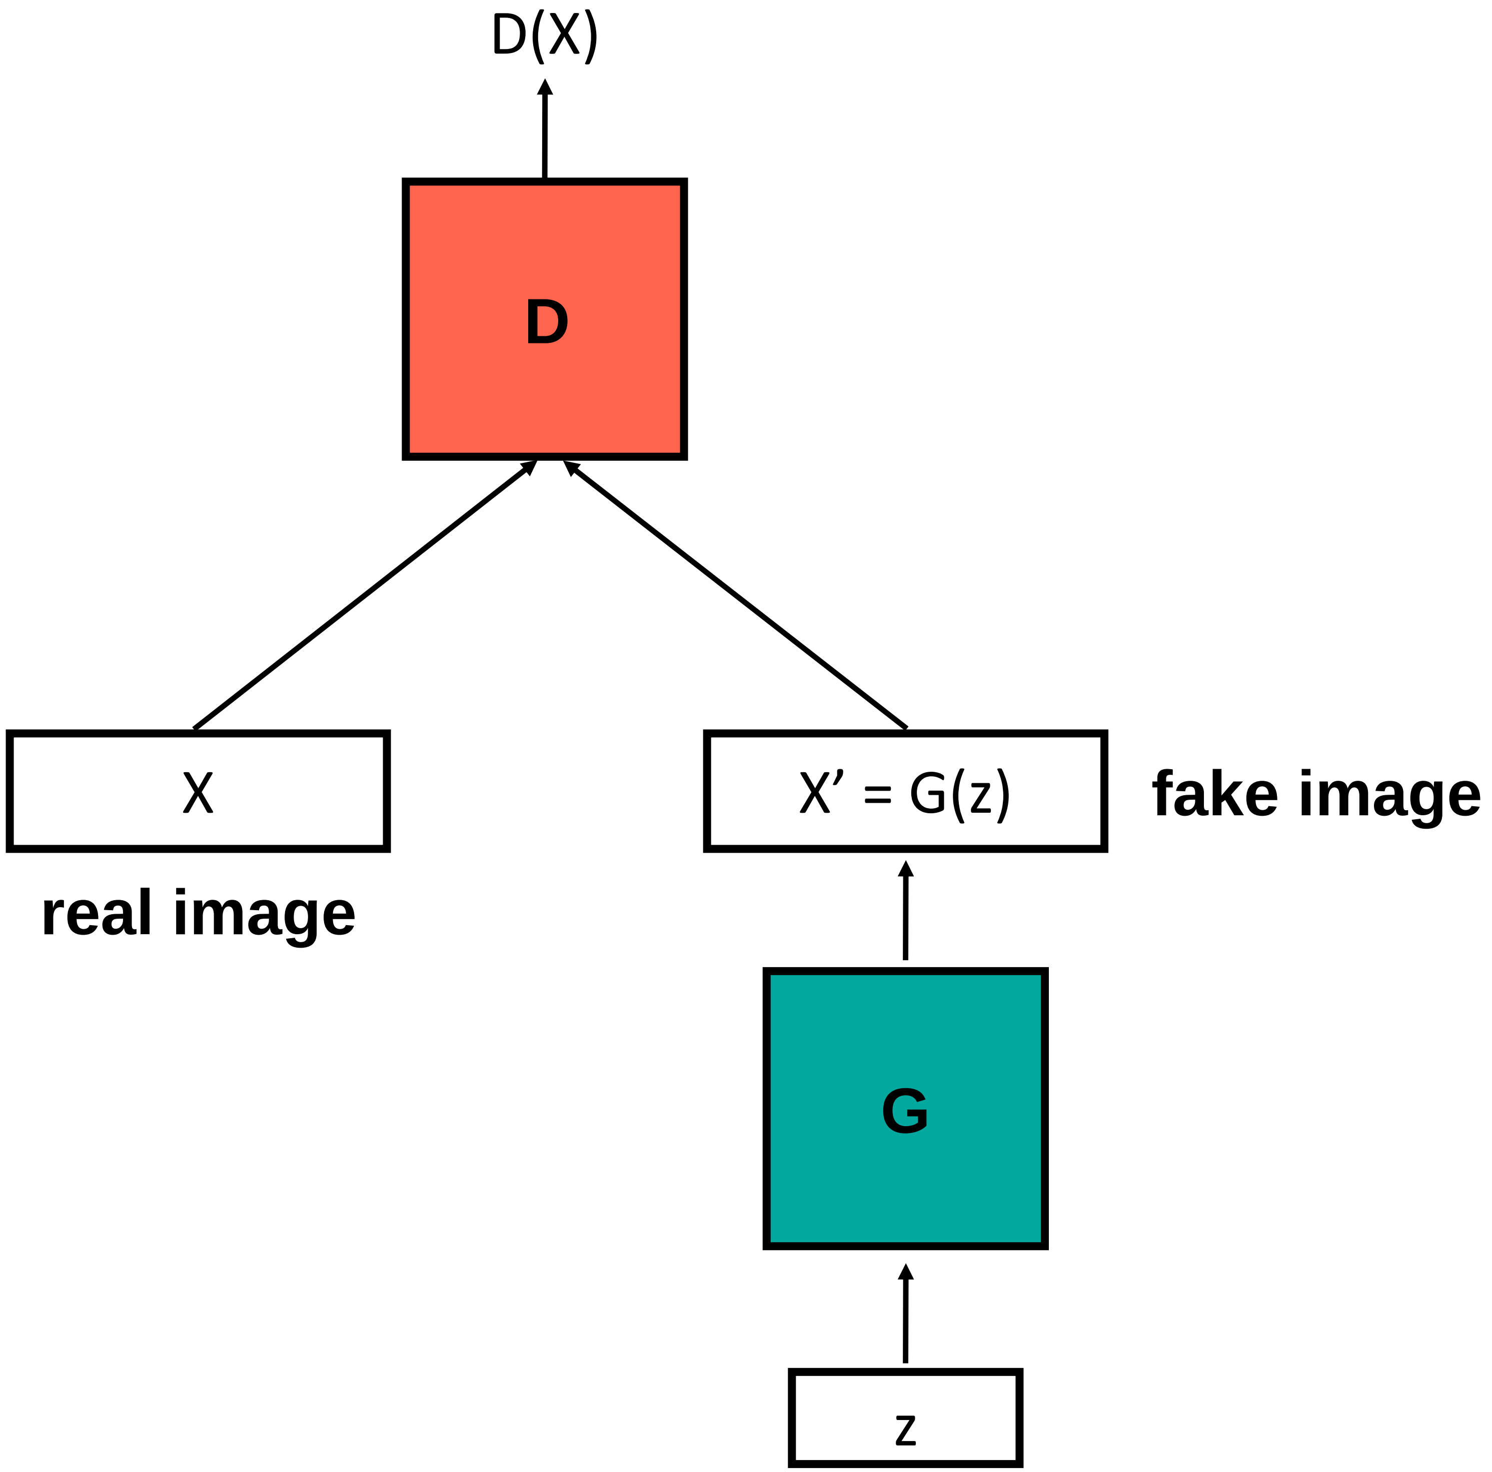

Conditional GAN

Conditional GAN

[M Mirza, et al, arXiv 2014]

CGAN

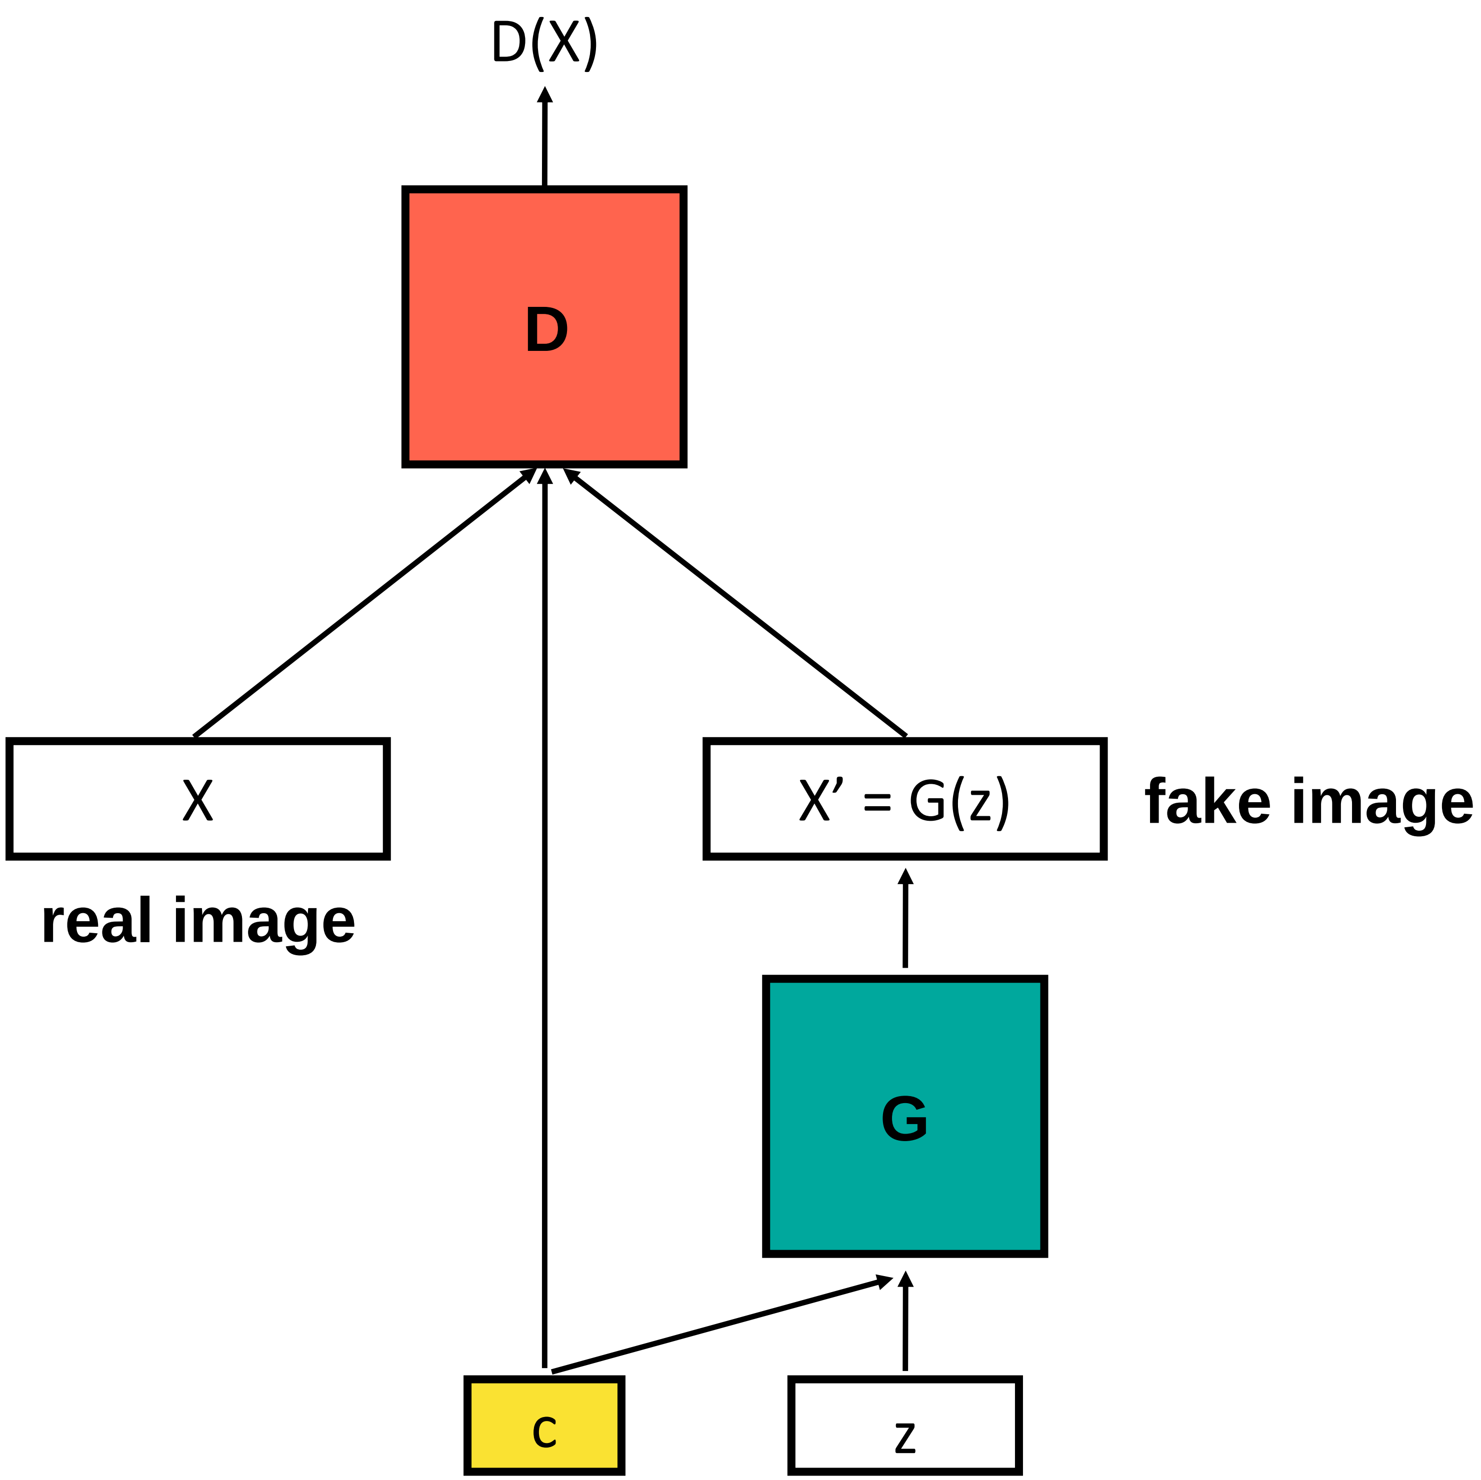

Conditional GAN

[August Odena, et al, ICML 2016]

ACGAN

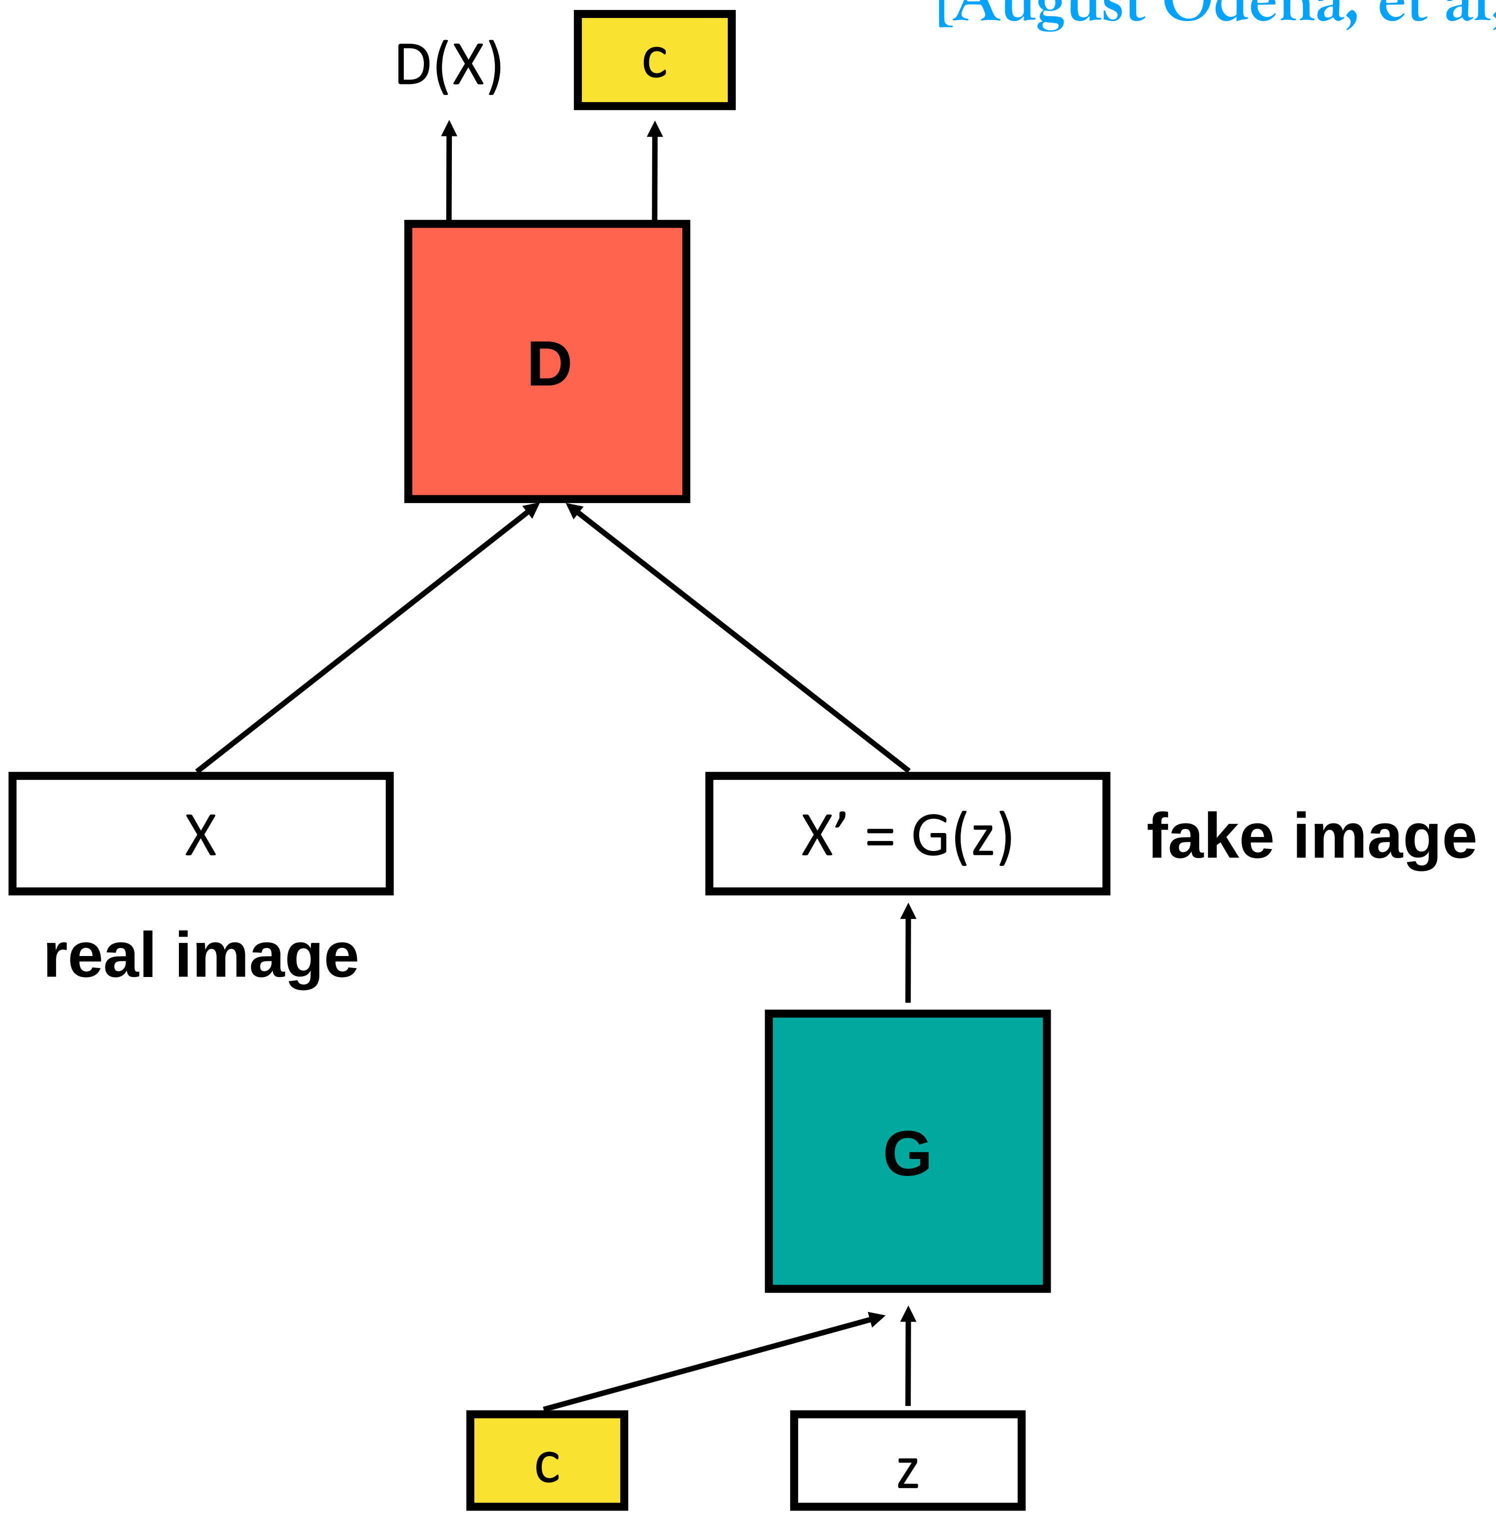

Conditional GAN

male, with glasses

female, with glasses

male, without glasses

female, without glasses

Conditional GAN

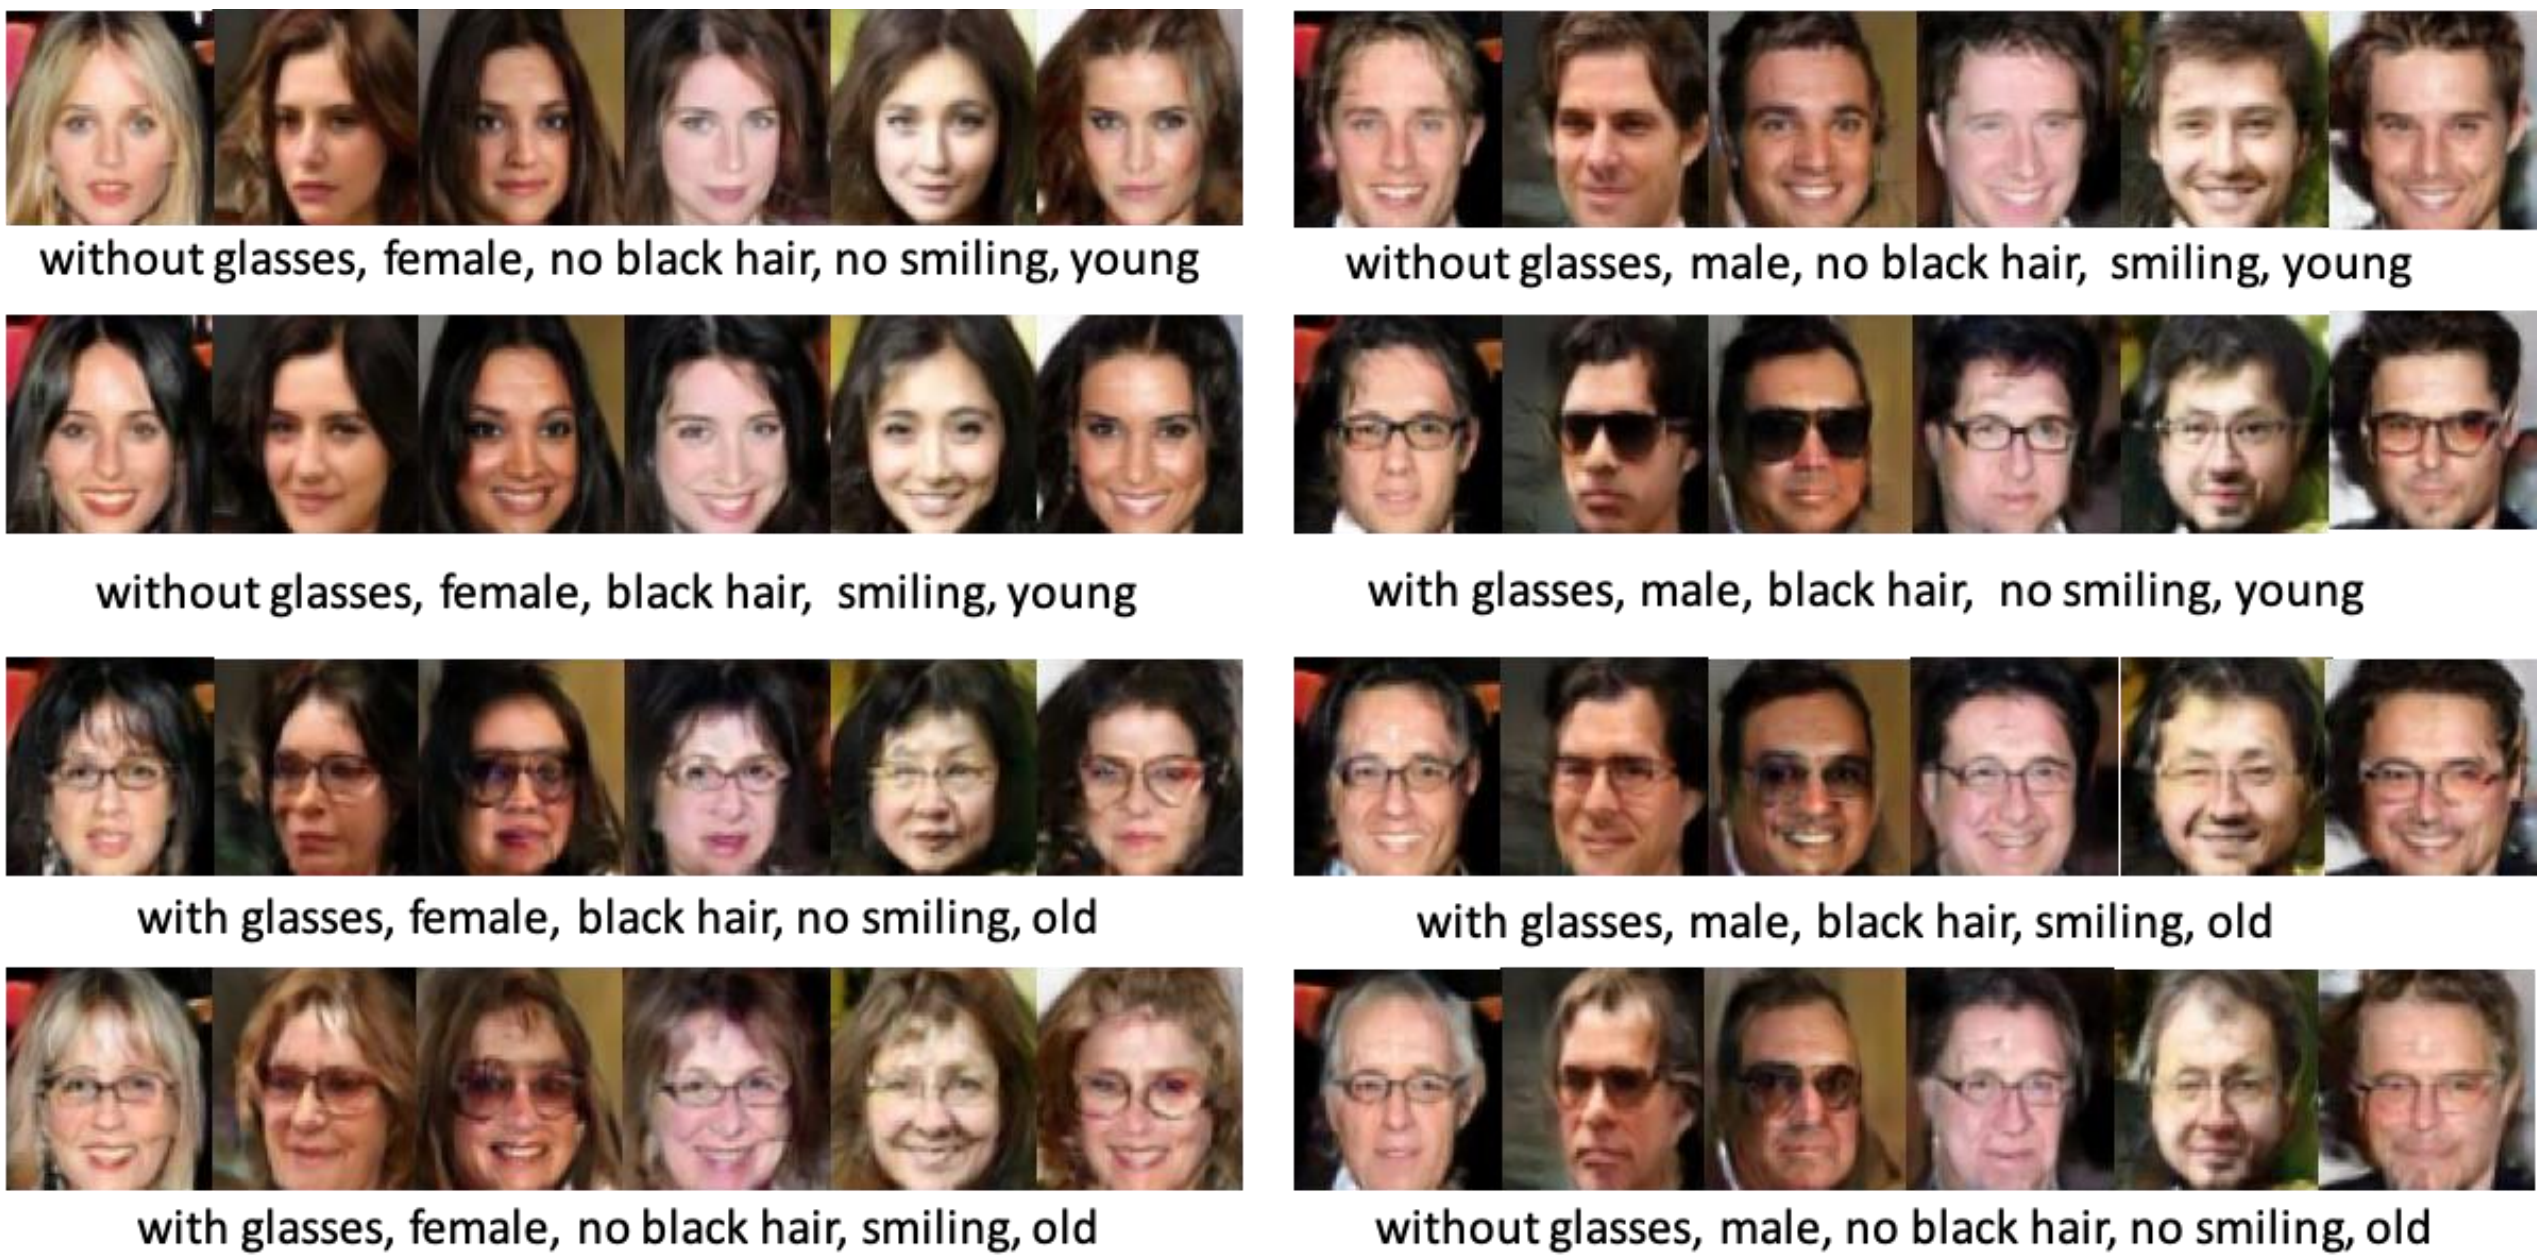

Conditional GAN

[Scott Reed, et al, ICML 2016]

Text-to-image Generation

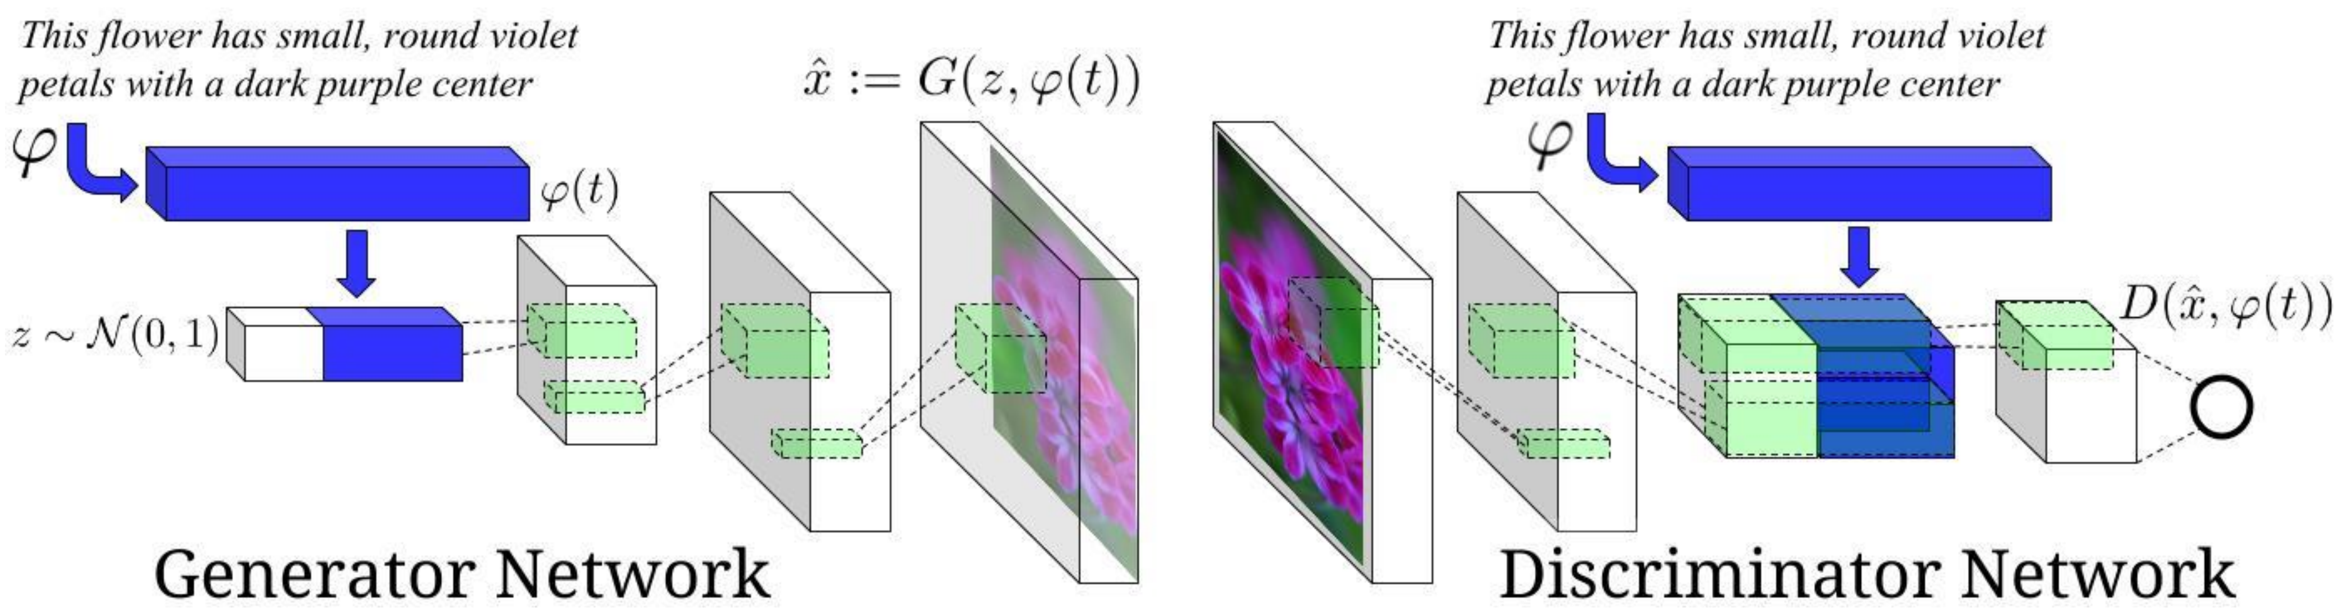

Image-to-image translation

Image-to-image translation

- Traditional method

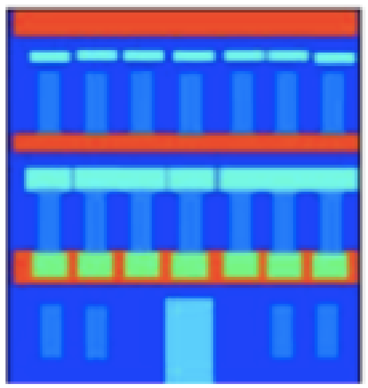

NN

Image

as close as possible

L1 / L2 loss

Testing:

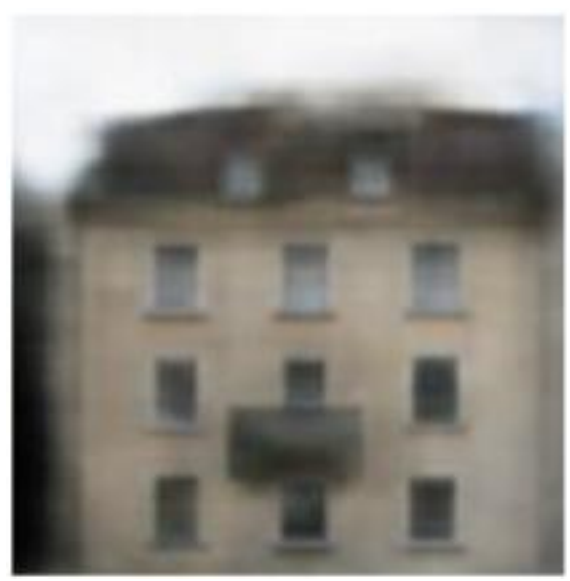

It is blurry,
what is the problem here ?

Image-to-image translation

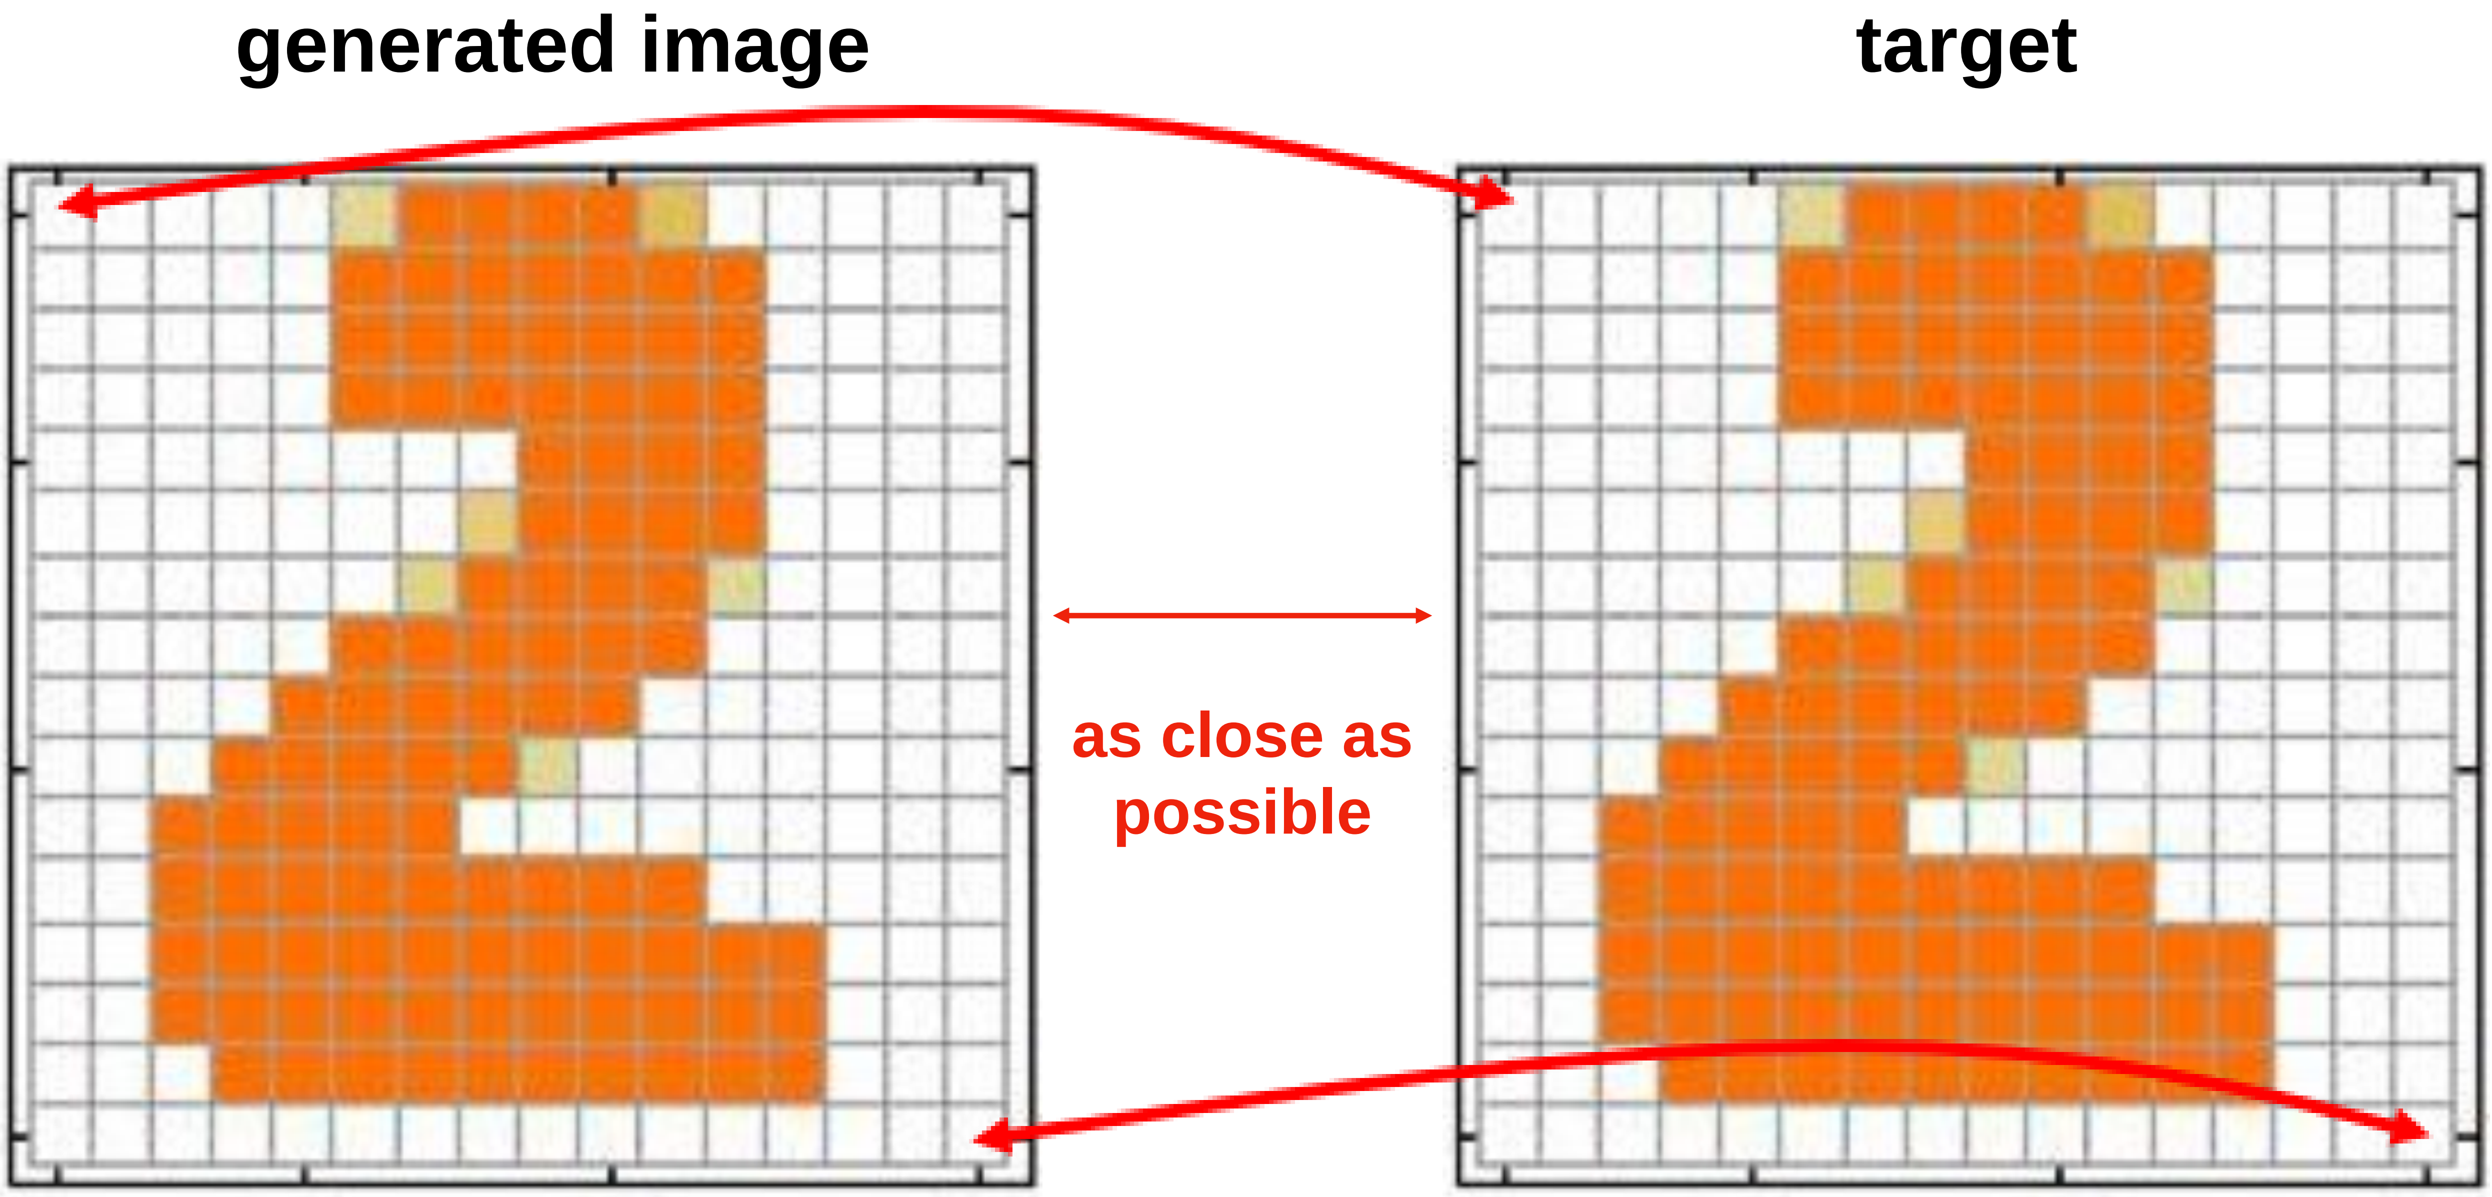

Image-to-image translation

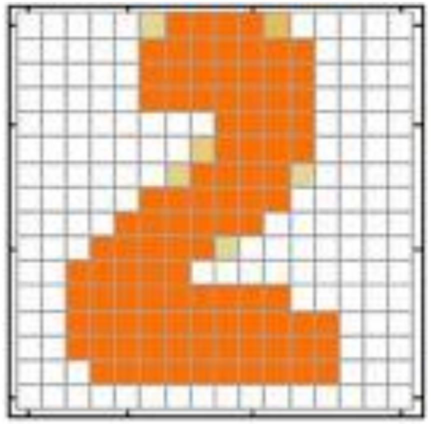

target

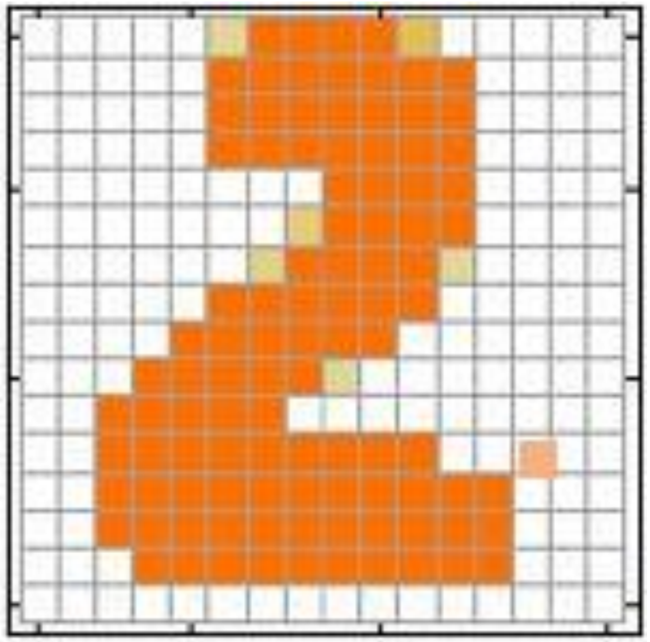

1 pixel error
not realistic

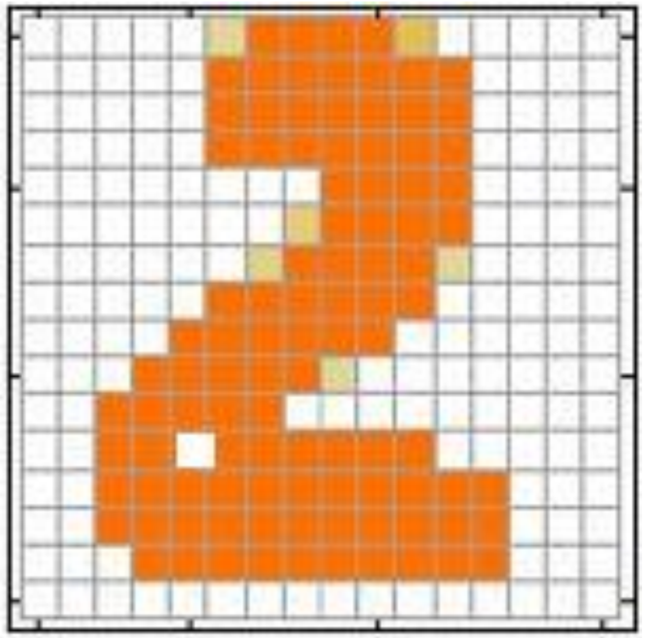

1 pixel error
not realistic

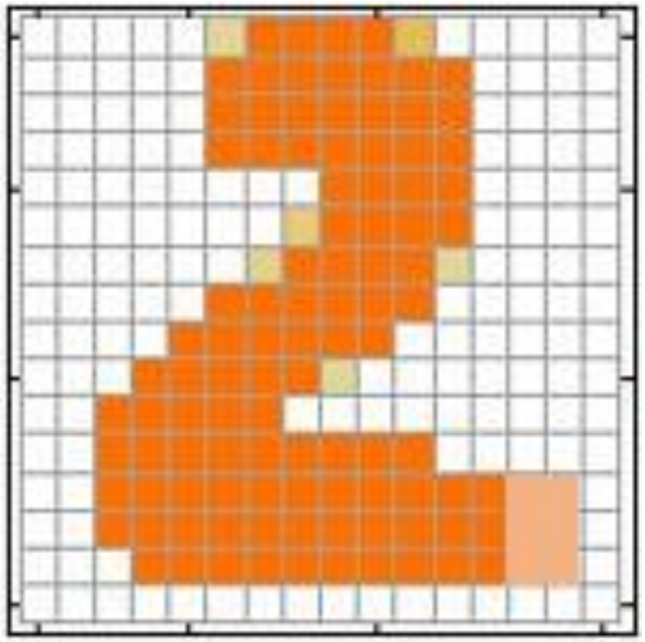

6 pixel error
realistic

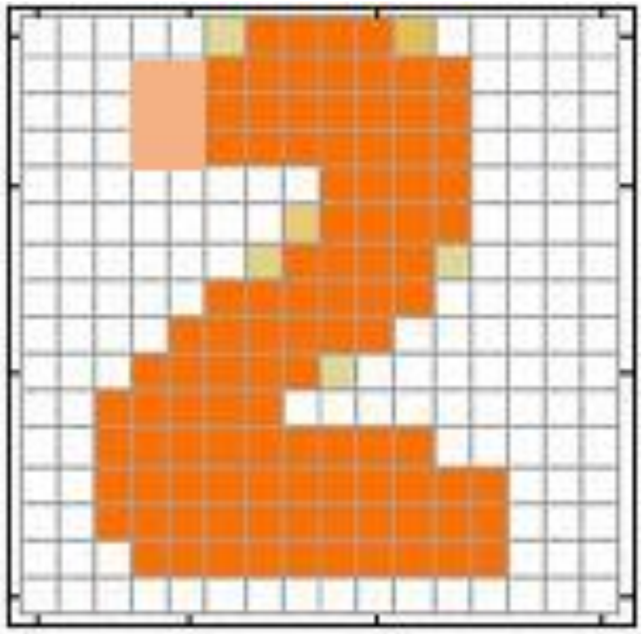

6 pixel error
realistic

Image-to-image translation

Reconstruction loss can not provide a sharp generation, what should be the solution ?

Since we can not find a good metric, we can use GAN to learn the metric !

Image-to-image translation

[Phillip Isola, et al, CVPR 2017]

- GAN method (Pix2Pix)

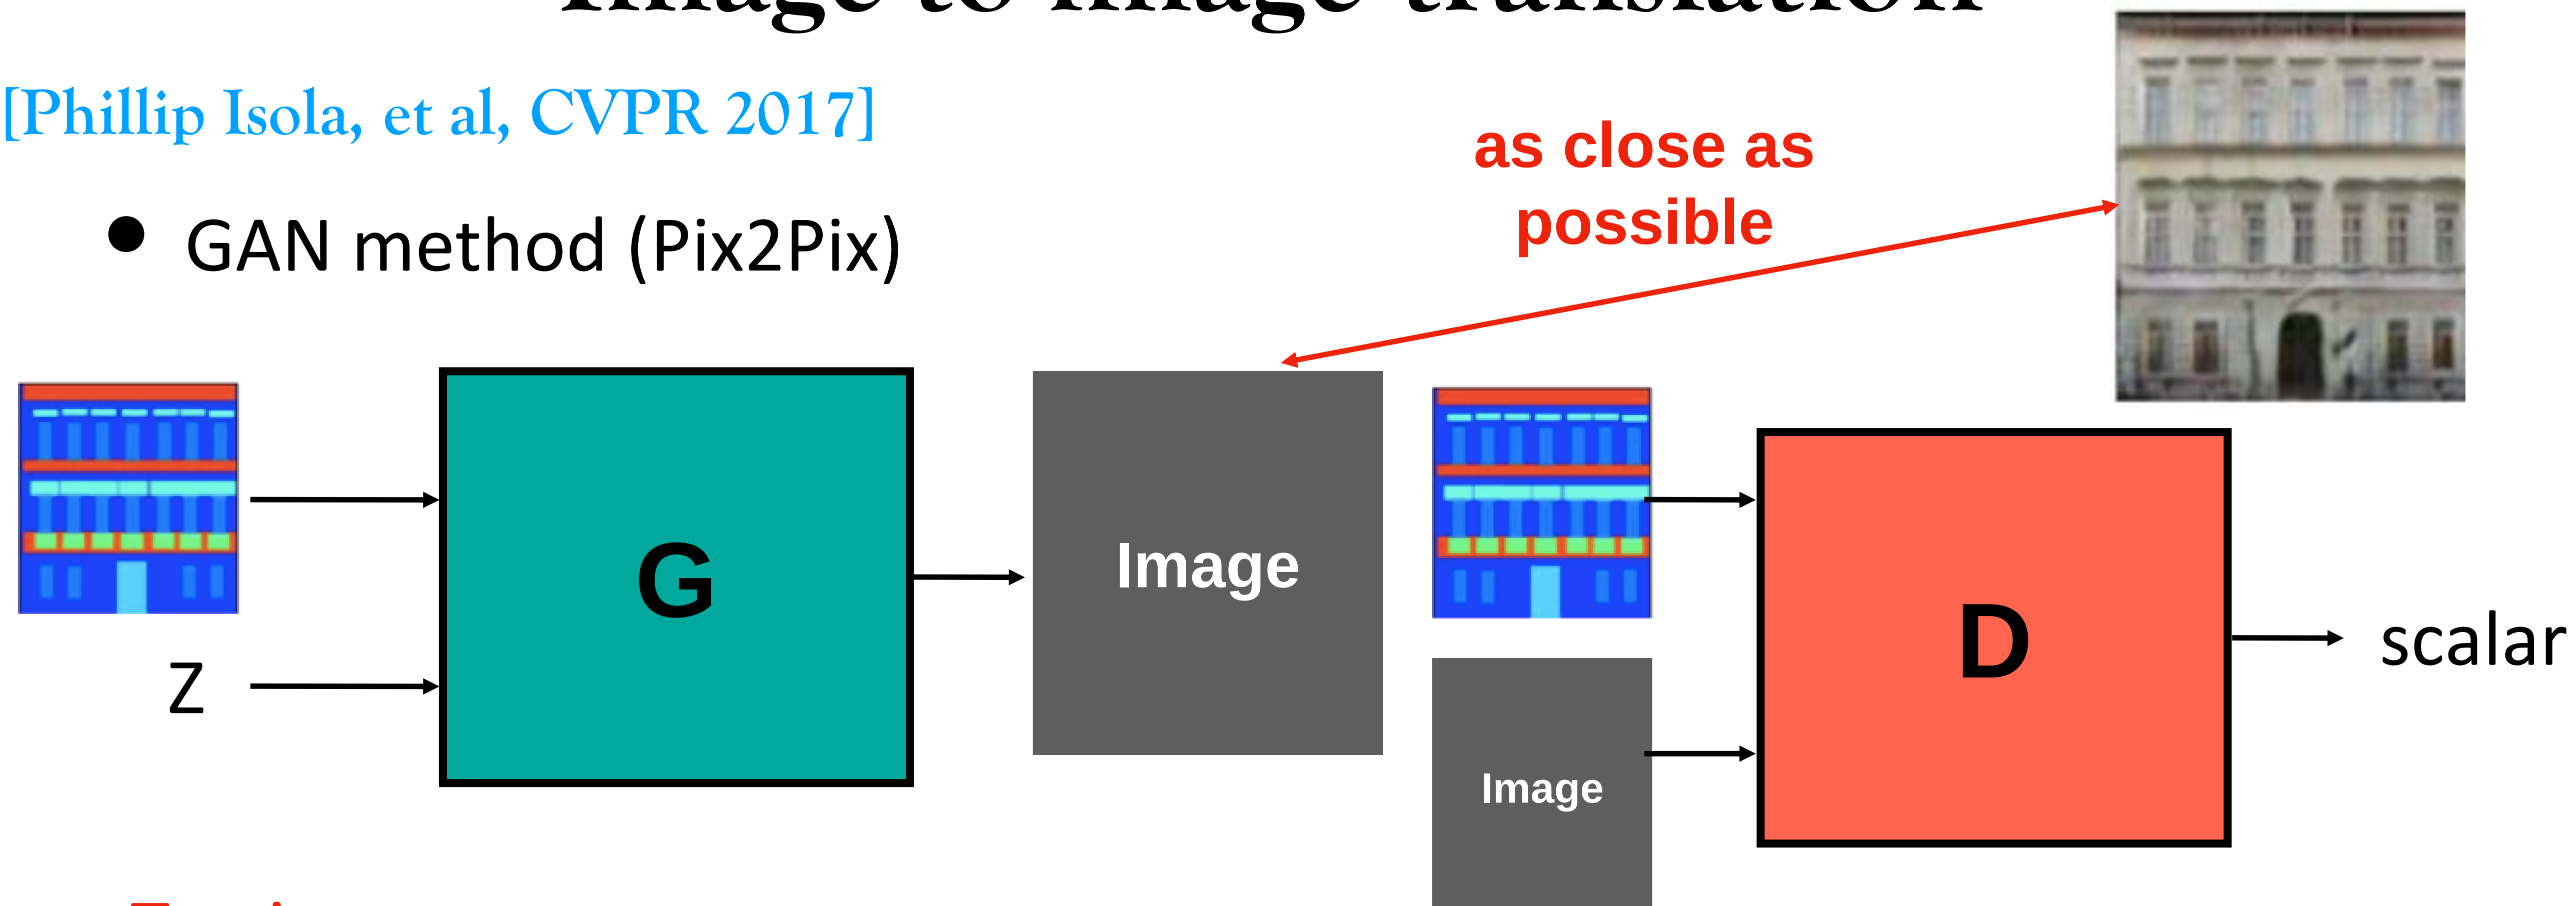

Testing:

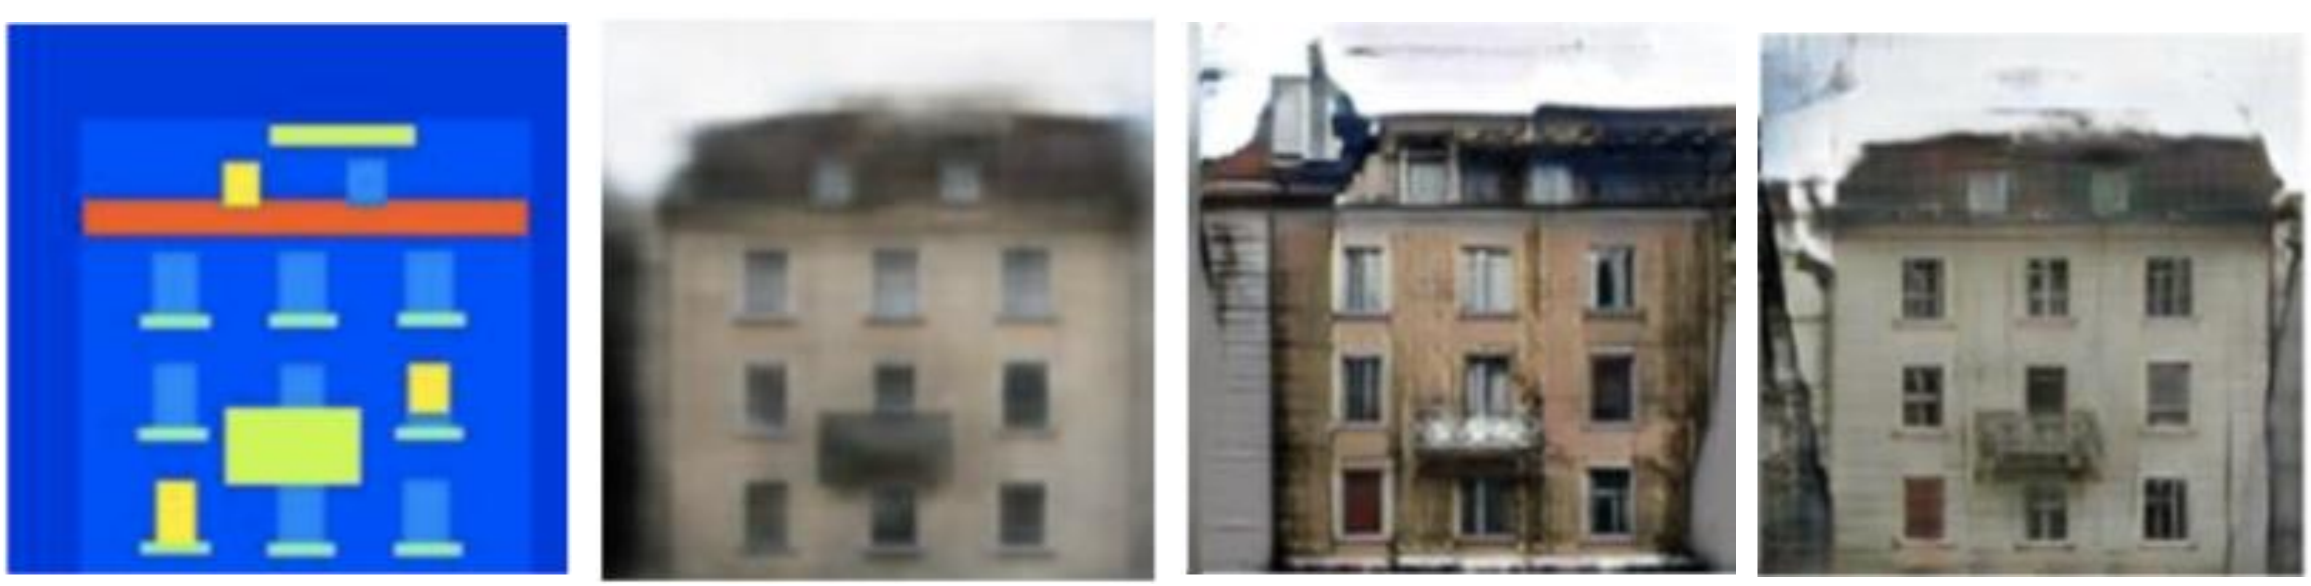

Input | Reconstruct | GAN | GAN + Reconstruct

Image-to-image translation

- What about unpaired data (no ground truth of target image) ?

X: zebra

Y: horse

X: summer

Y: winter

Image-to-image translation

[Jun-yan Zhu, et al, ICCV 2017]

- CycleGAN

Image-to-image translation

[Jun-yan Zhu, et al, ICCV 2017]

- CycleGAN

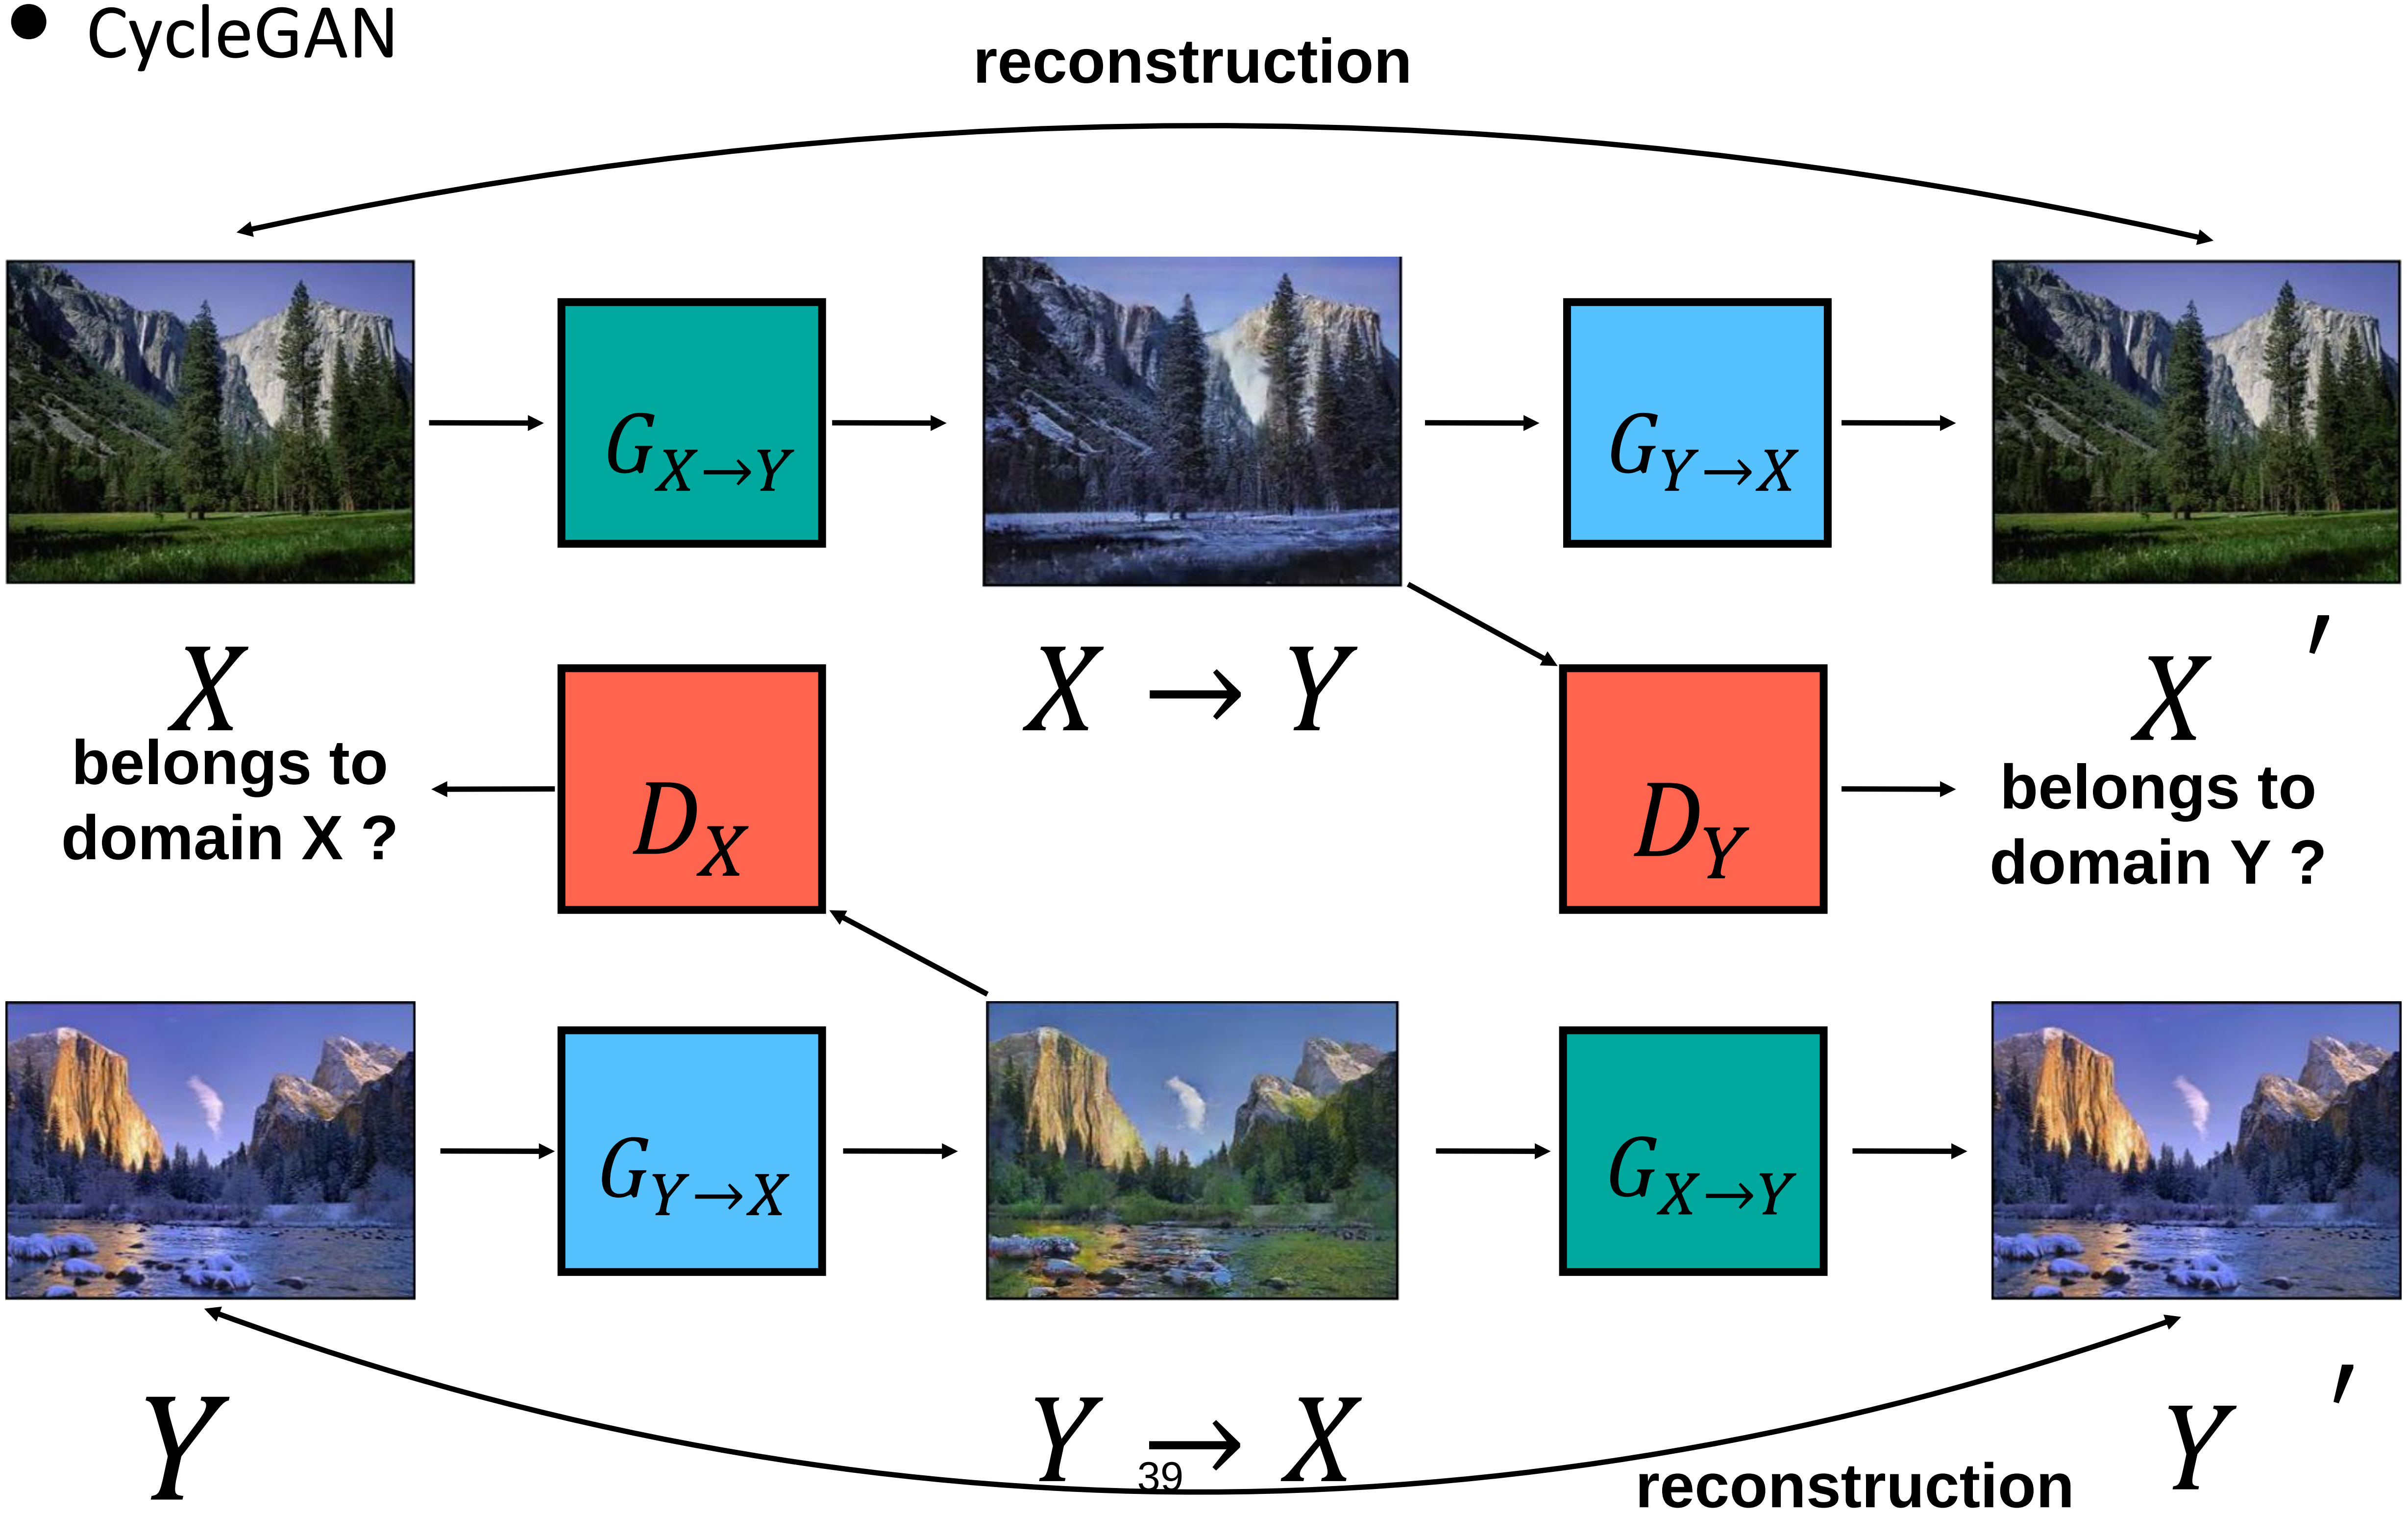

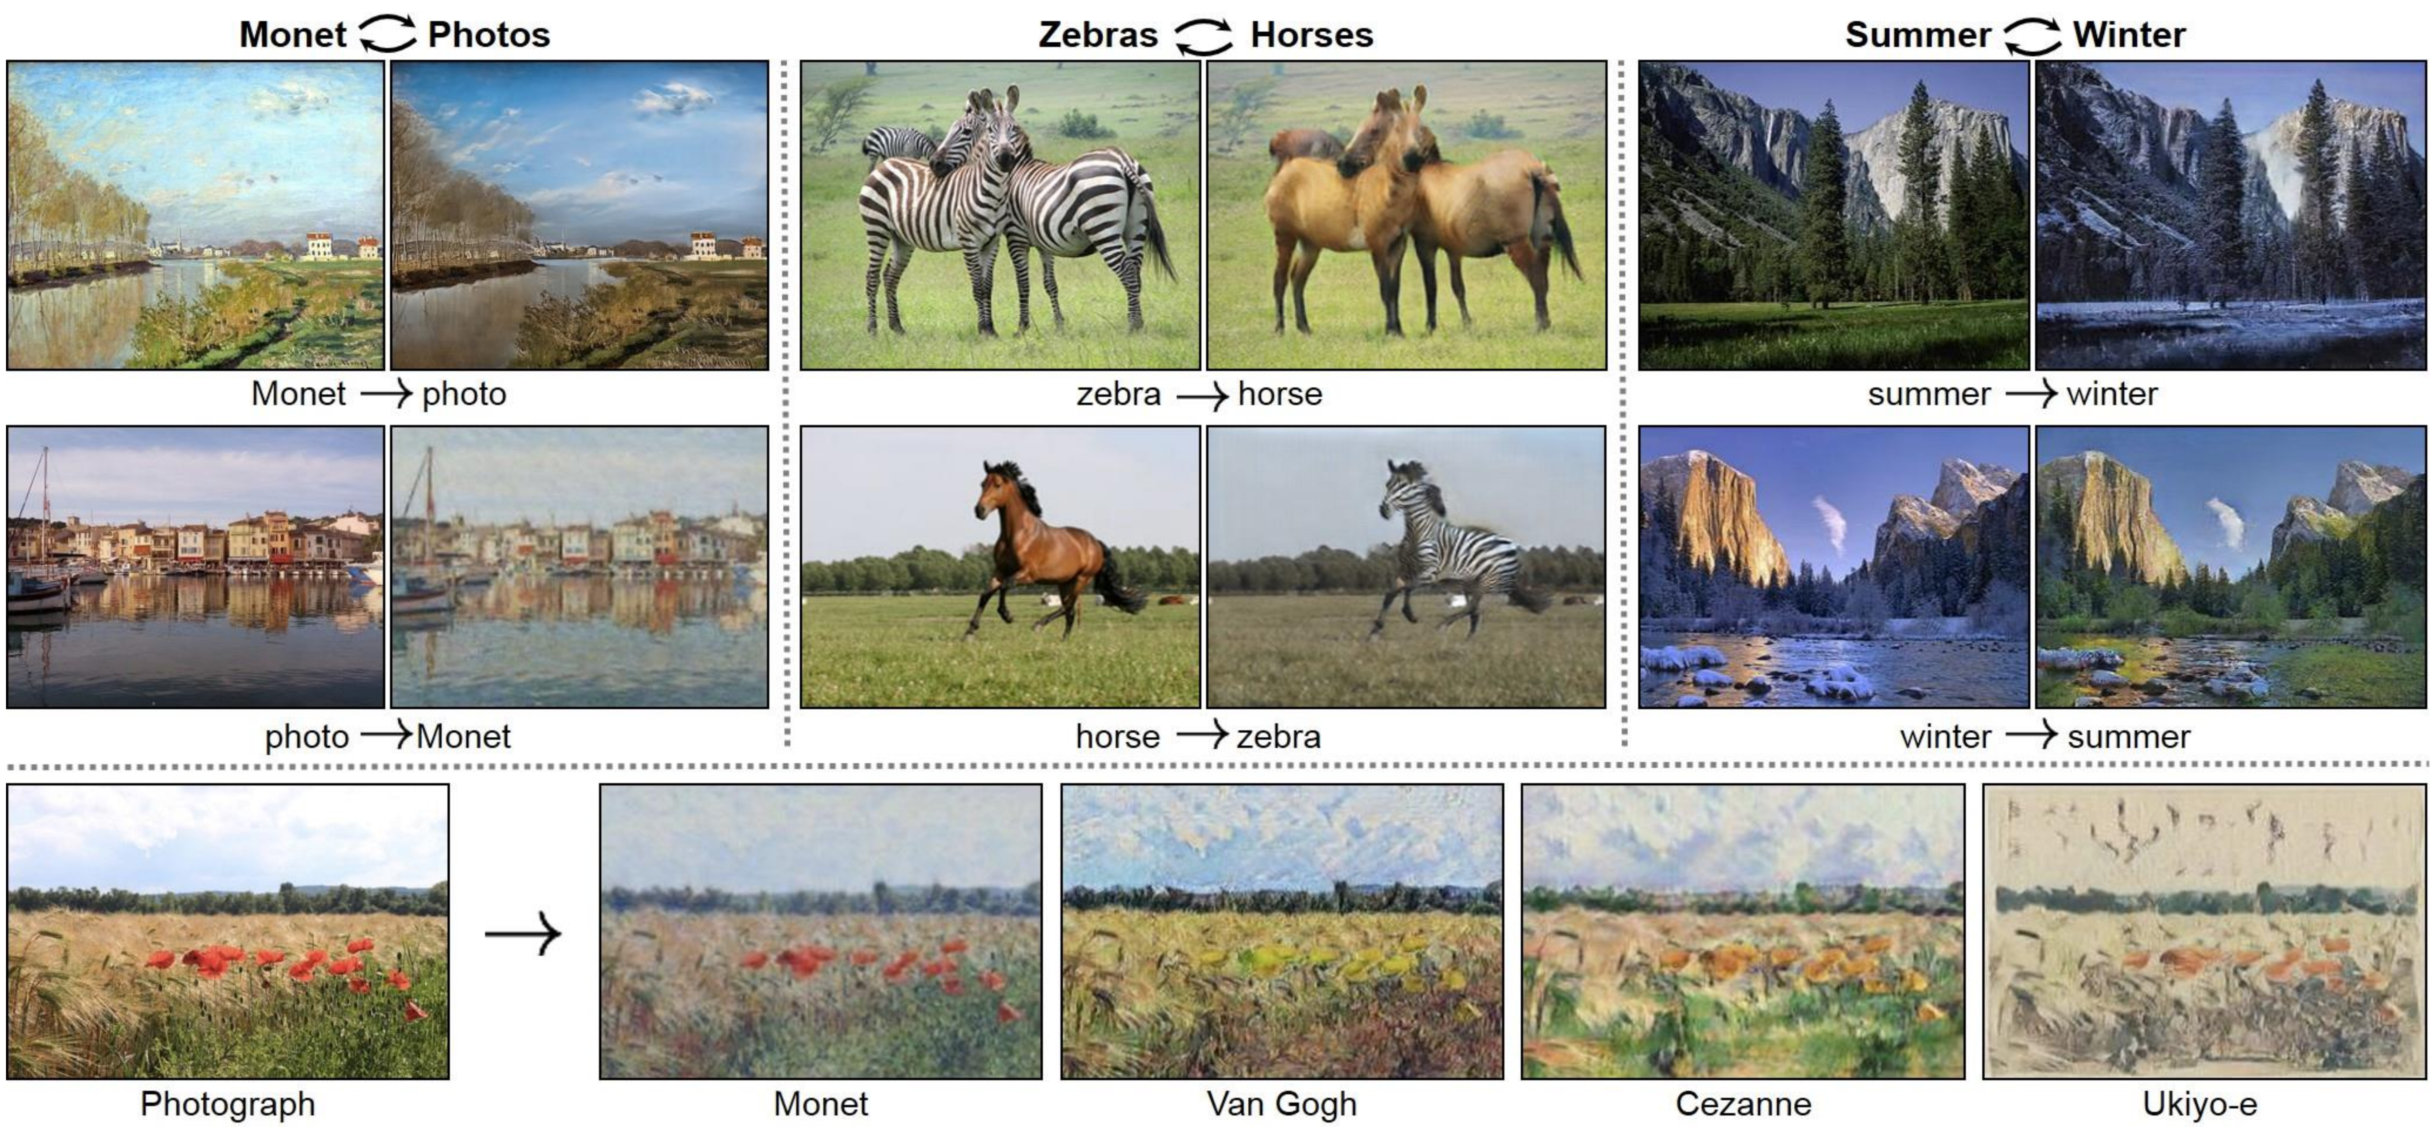
Monet ⇄ Photos
Zebras ⇄ Horses
Summer ⇄ Winter
Monet → photo
zebra → horse
summer → winter
photo → Monet
horse → zebra
winter → summer
Photograph
Monet
Van Gogh
Cezanne
Ukiyo-e

Image-to-image translation

- UNIT
- MUNIT
- …

Modern GAN Architectures

Modern Architectures

DCGAN

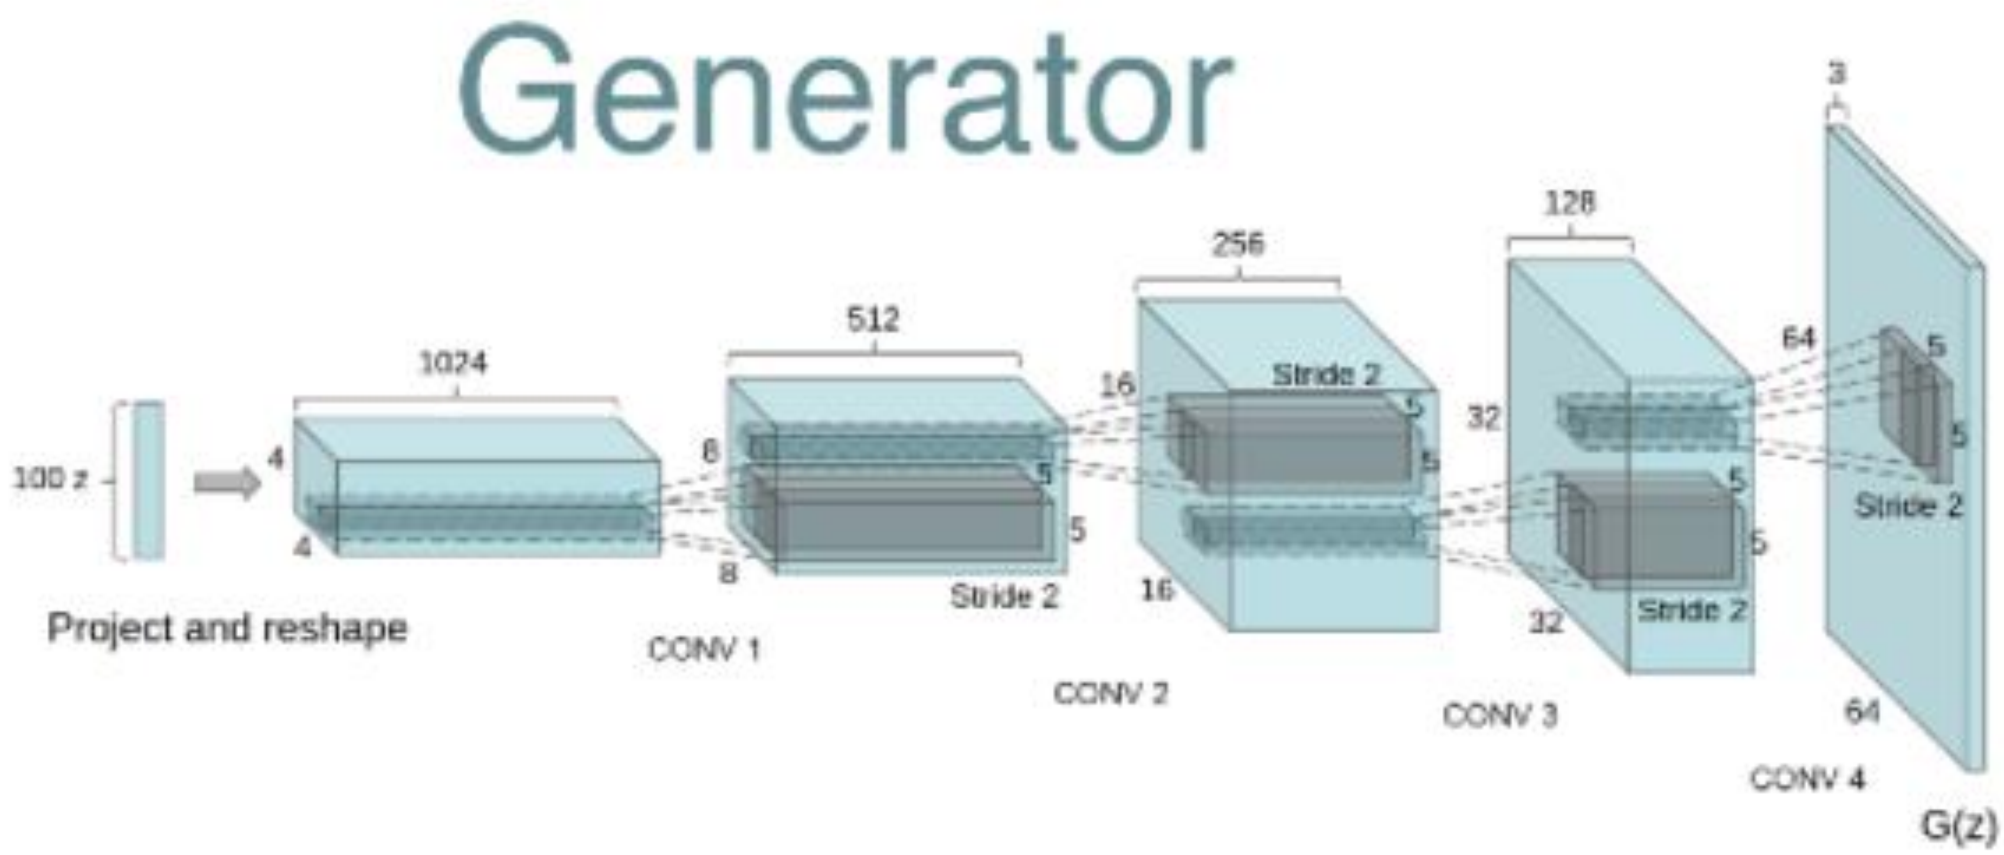

https://github.com/vdumoulin/conv_arithmetic

[A Radford, et al, arXiv 2015]

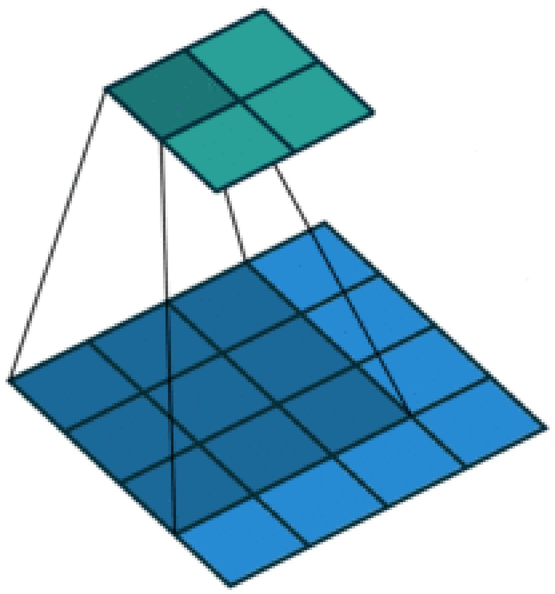

convolution

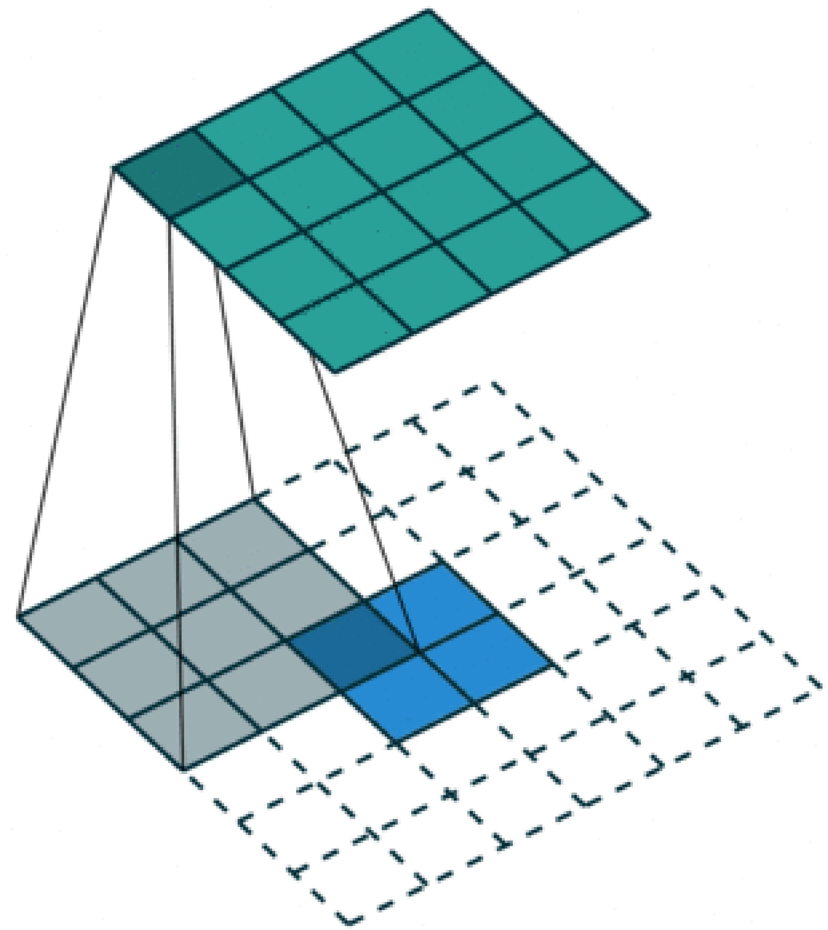

transposed convolution

Modern Architectures

https://github.com/vdumoulin/conv_arithmetic

[A Radford, et al, arXiv 2015]

Results - MNIST

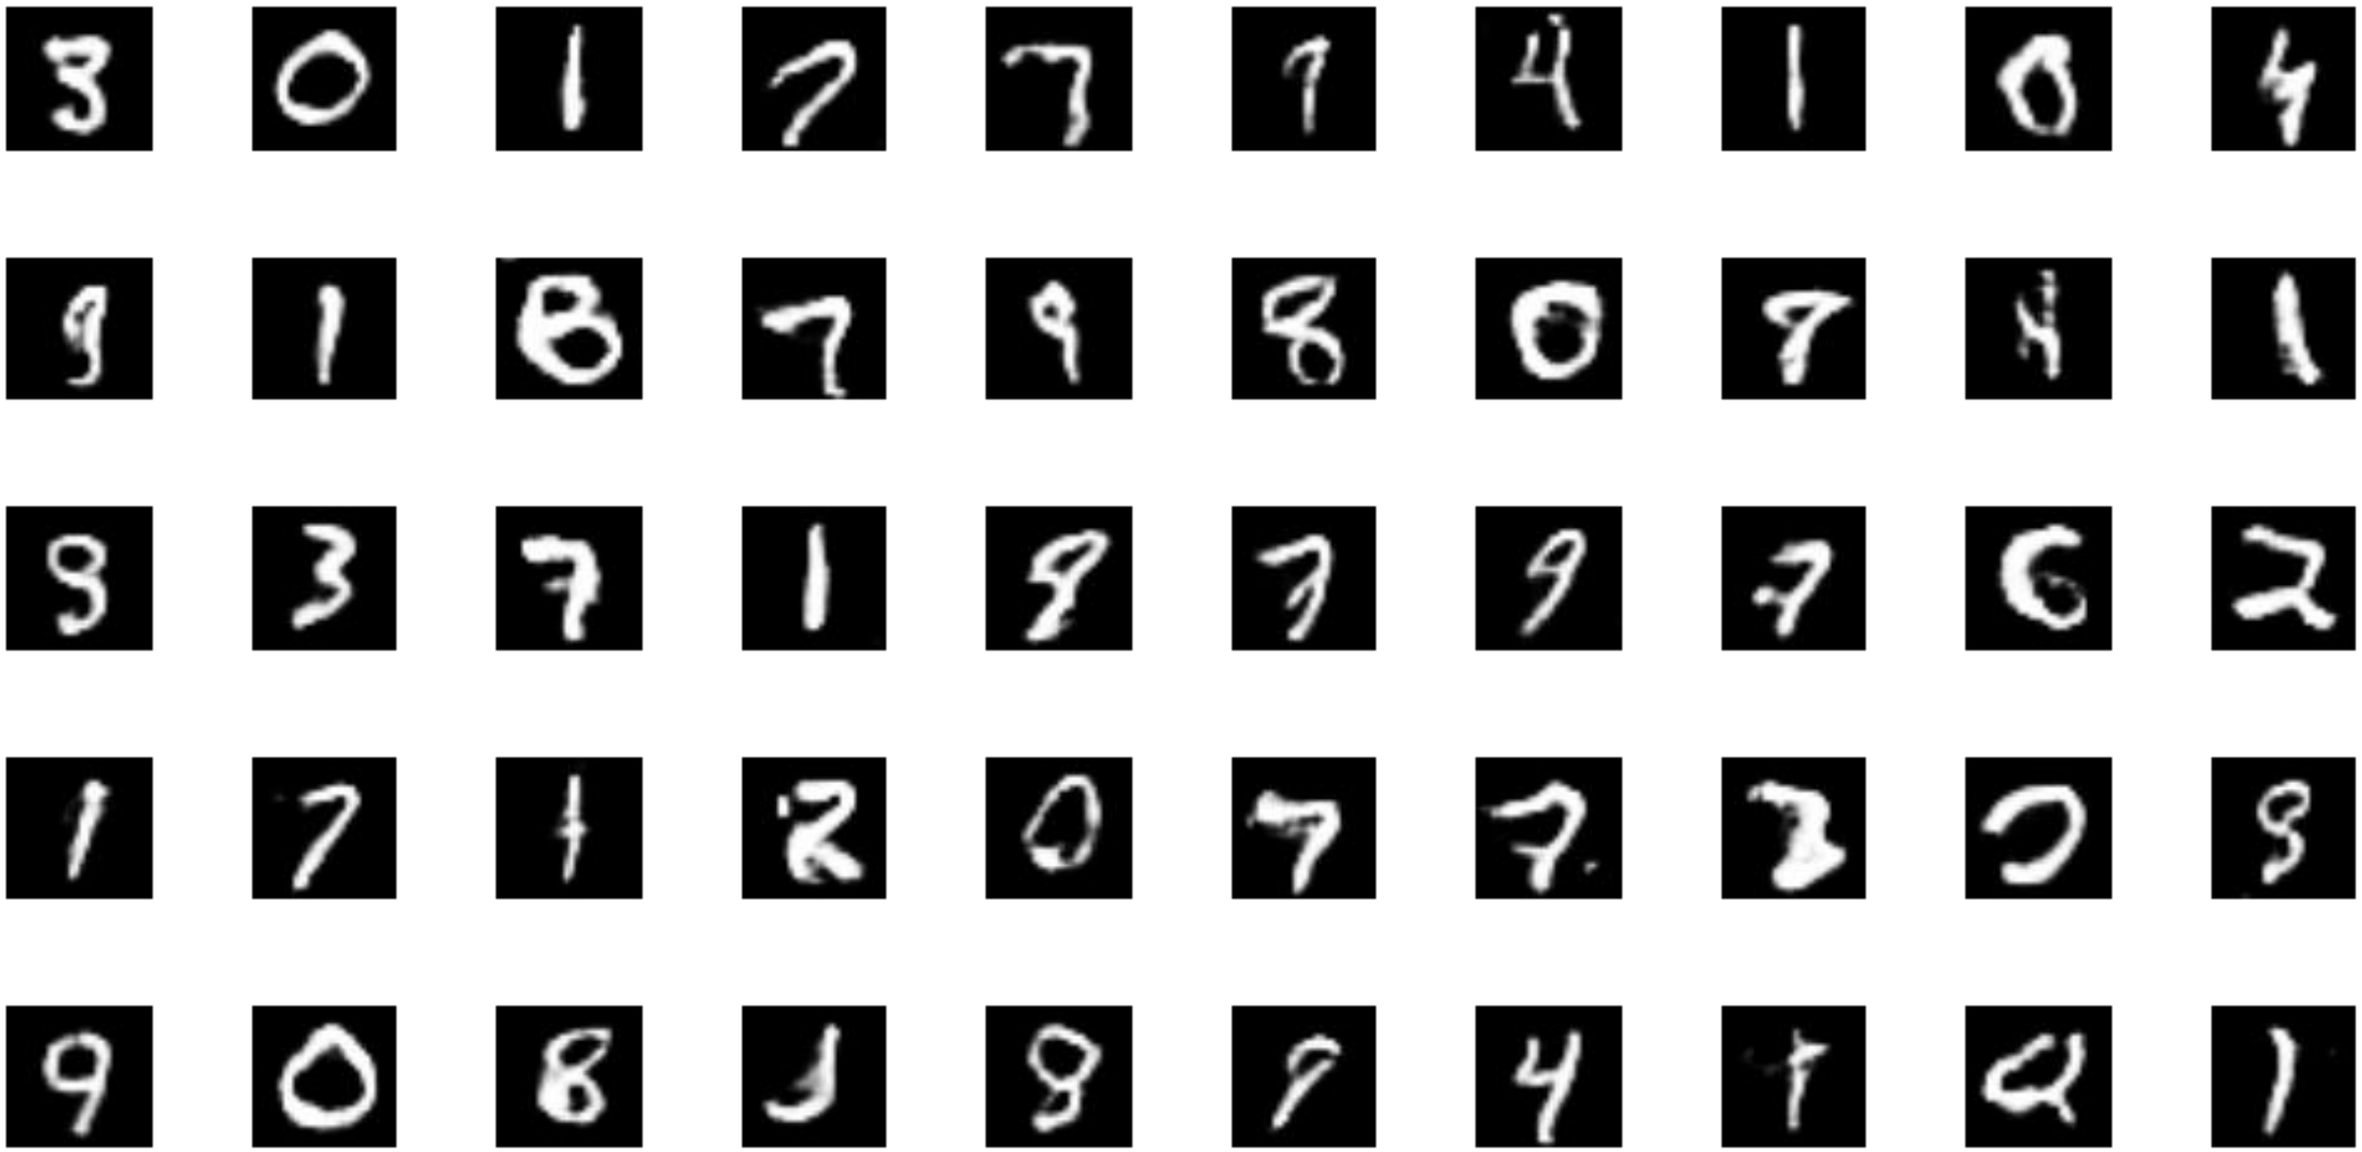

Results - CelebA (faces)

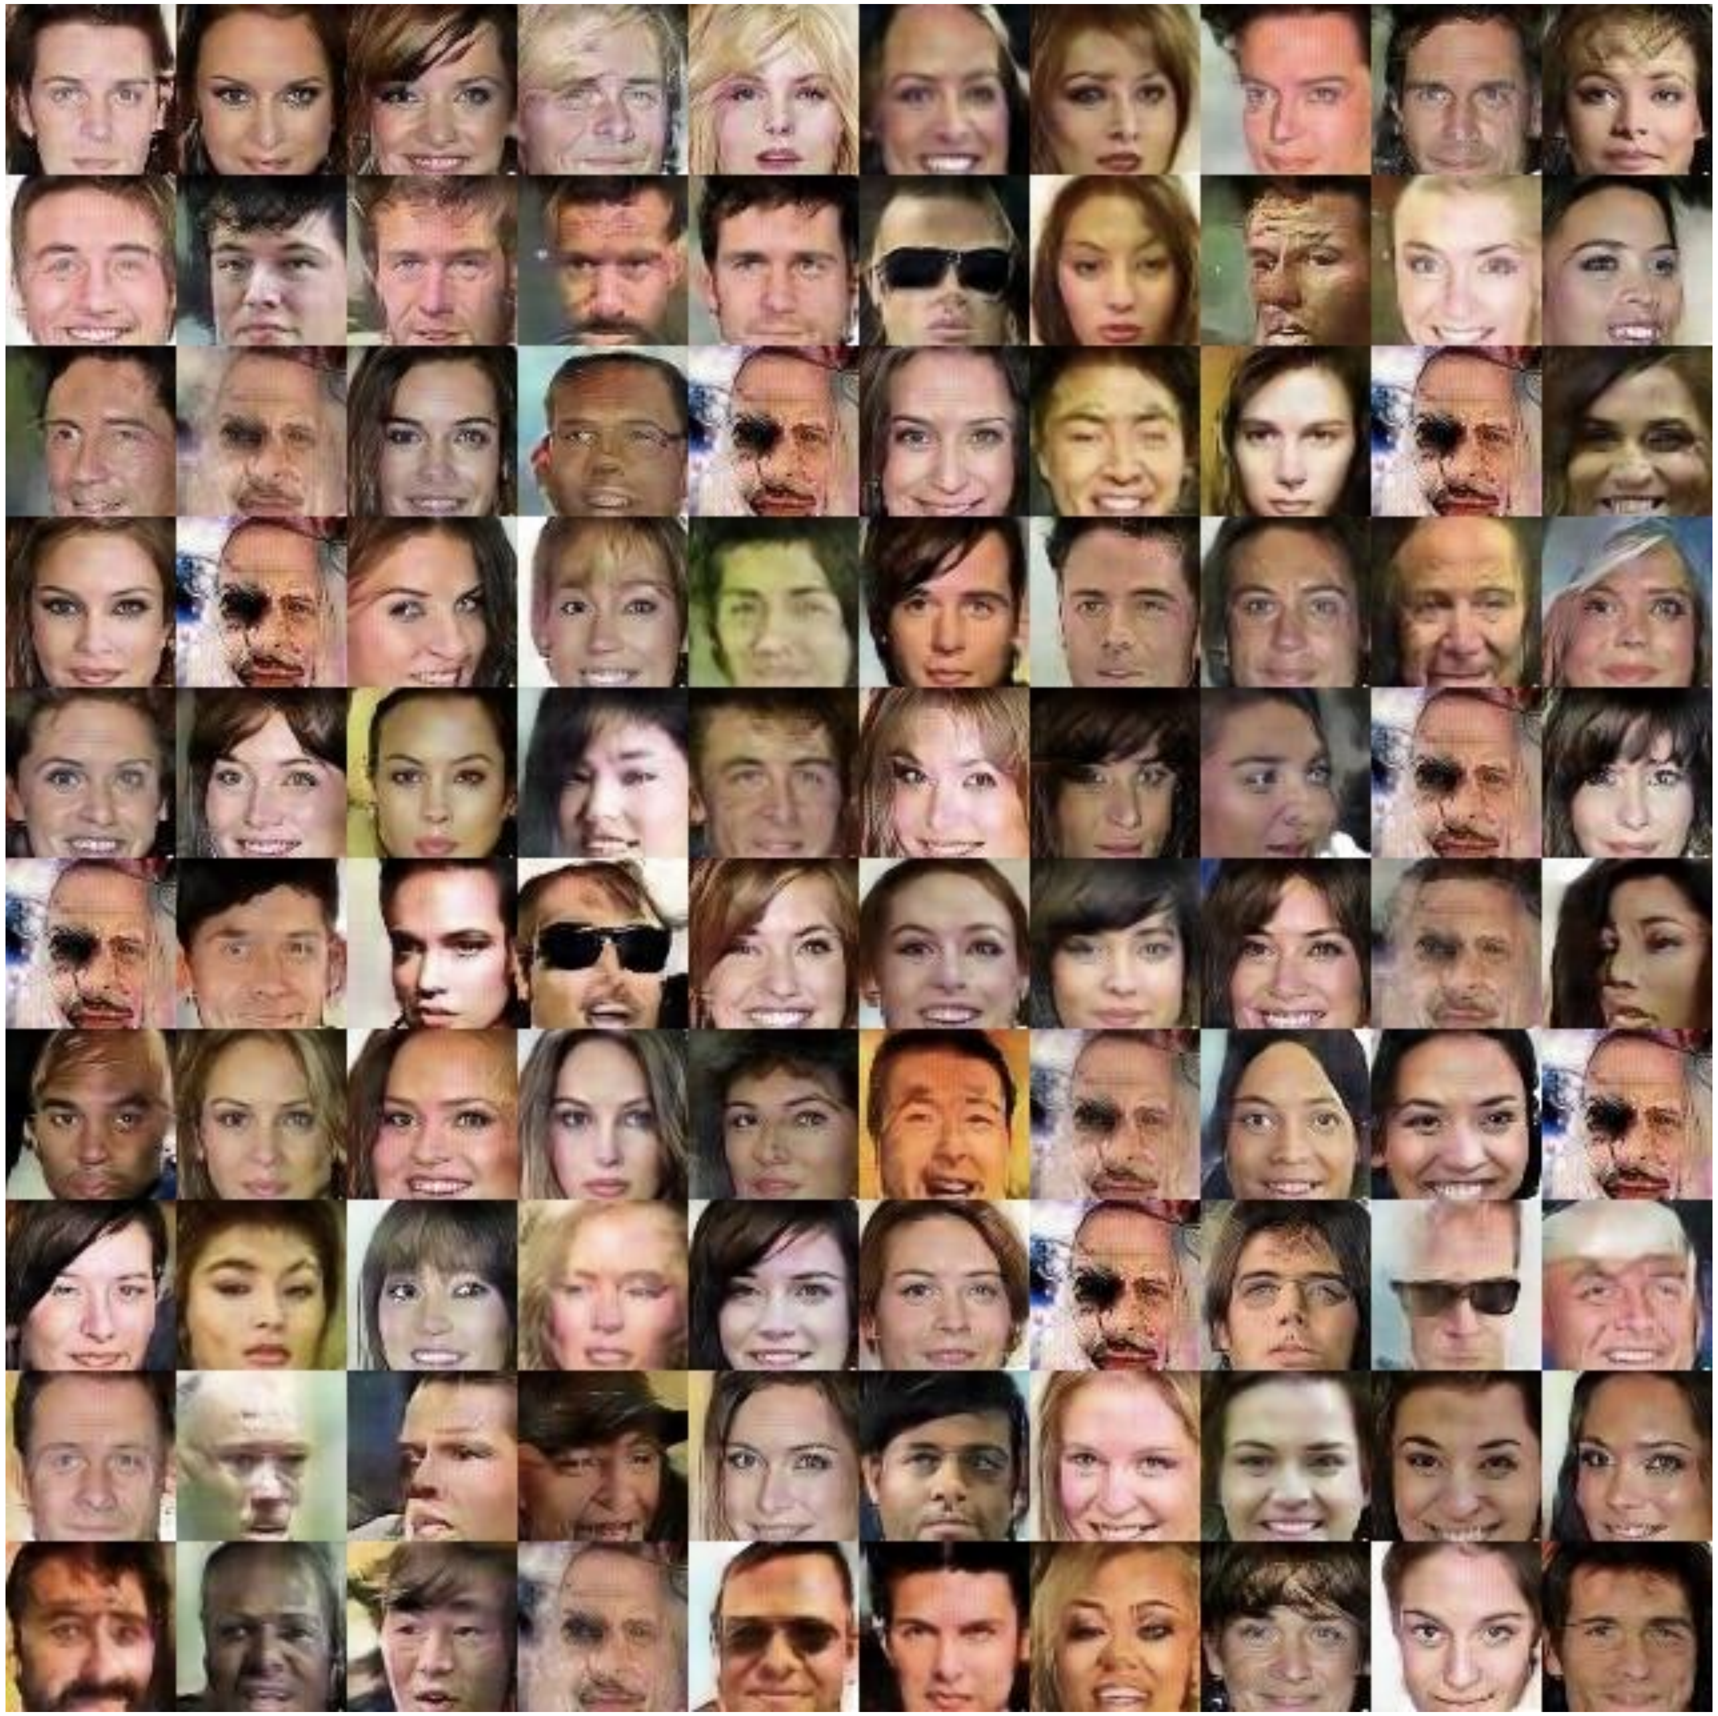

Results - LSUN (bedrooms)

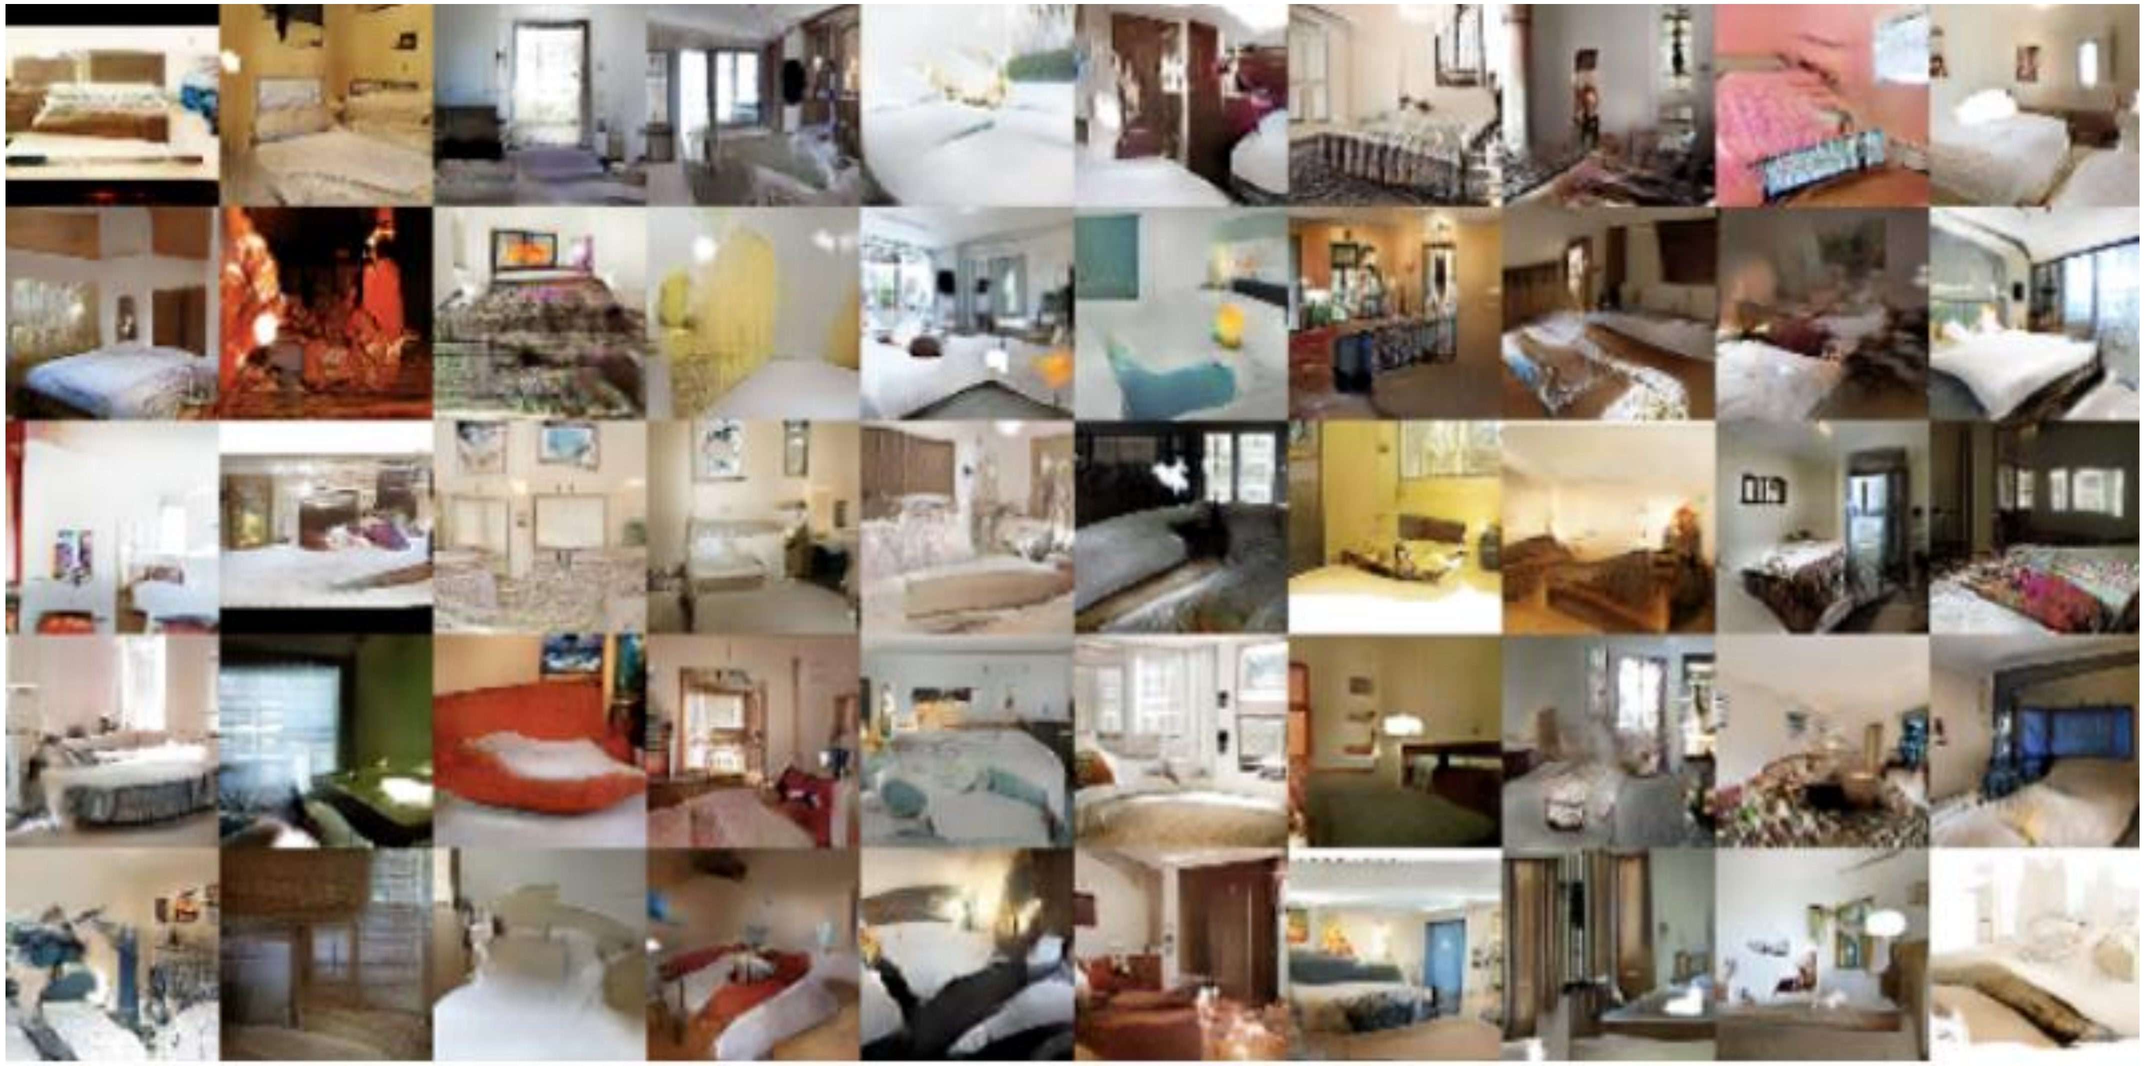

Modern Architectures

StyleGAN (NVIDIA)

https://github.com/NVlabs/stylegan

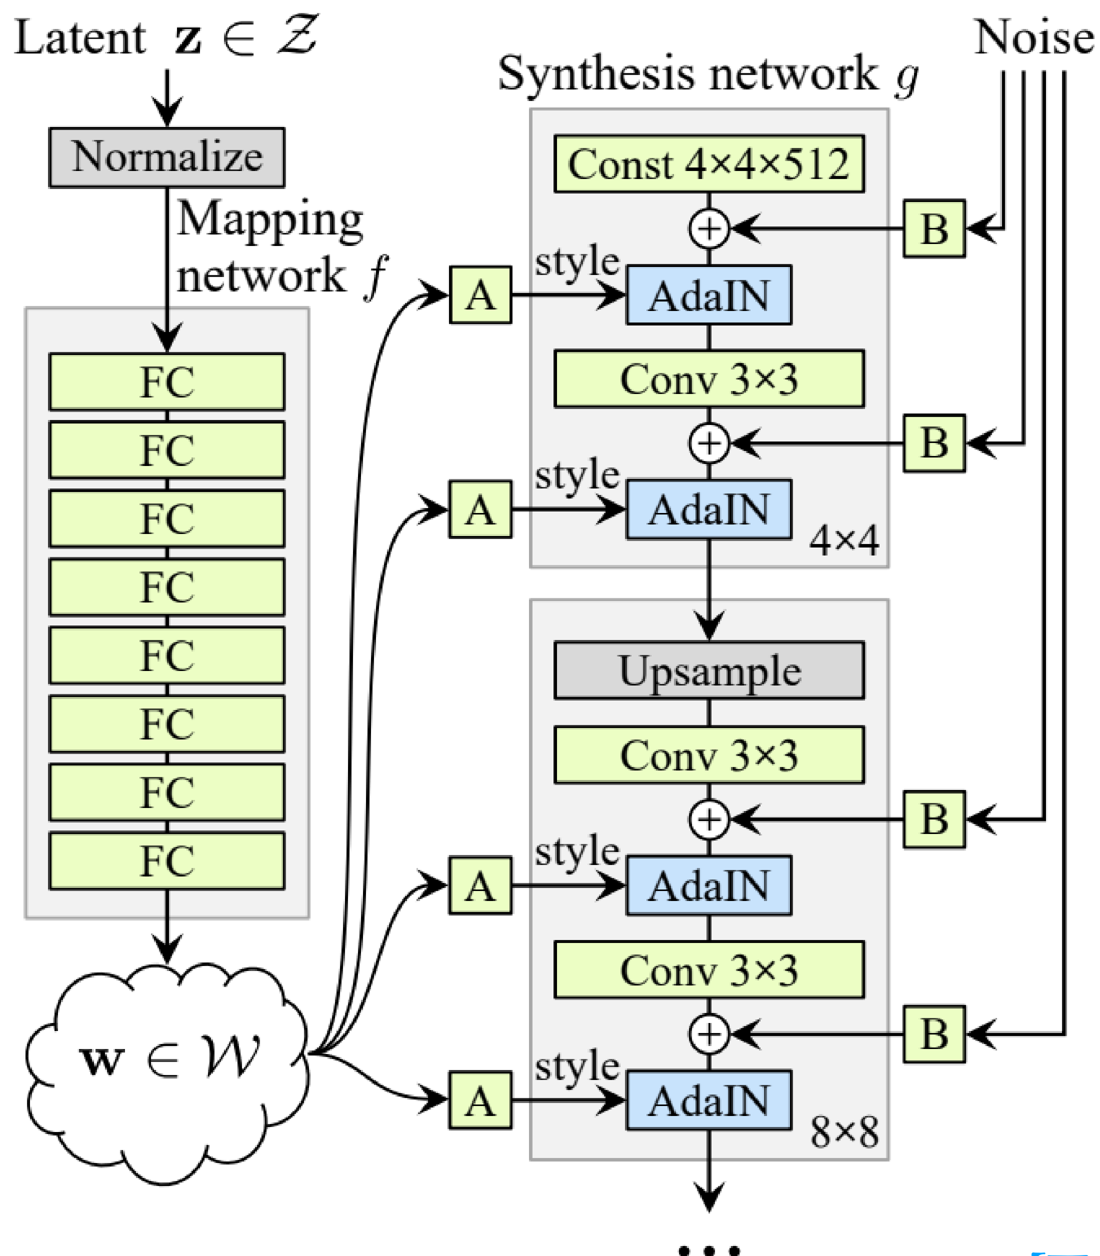

[T Karras, et al, CVPR 2019]

Modern Architectures

StyleGAN

https://www.youtube.com/watch?v=kSLJriaOumA

Karras et al, A Style-Based Generator Architecture for Generative Adversarial Networks, CVPR 2019

Modern Architectures

StyleGAN

GPUs	1024×1024	512×512	256×256
1	41 days 4 hours	24 days 21 hours	14 days 22 hours
2	21 days 22 hours	13 days 7 hours	9 days 5 hours
4	11 days 8 hours	7 days 0 hours	4 days 21 hours
8	6 days 14 hours	4 days 10 hours	3 days 8 hours

Modern Architectures

BigGAN (DeepMind)

https://github.com/ajbrock/BigGAN-PyTorch

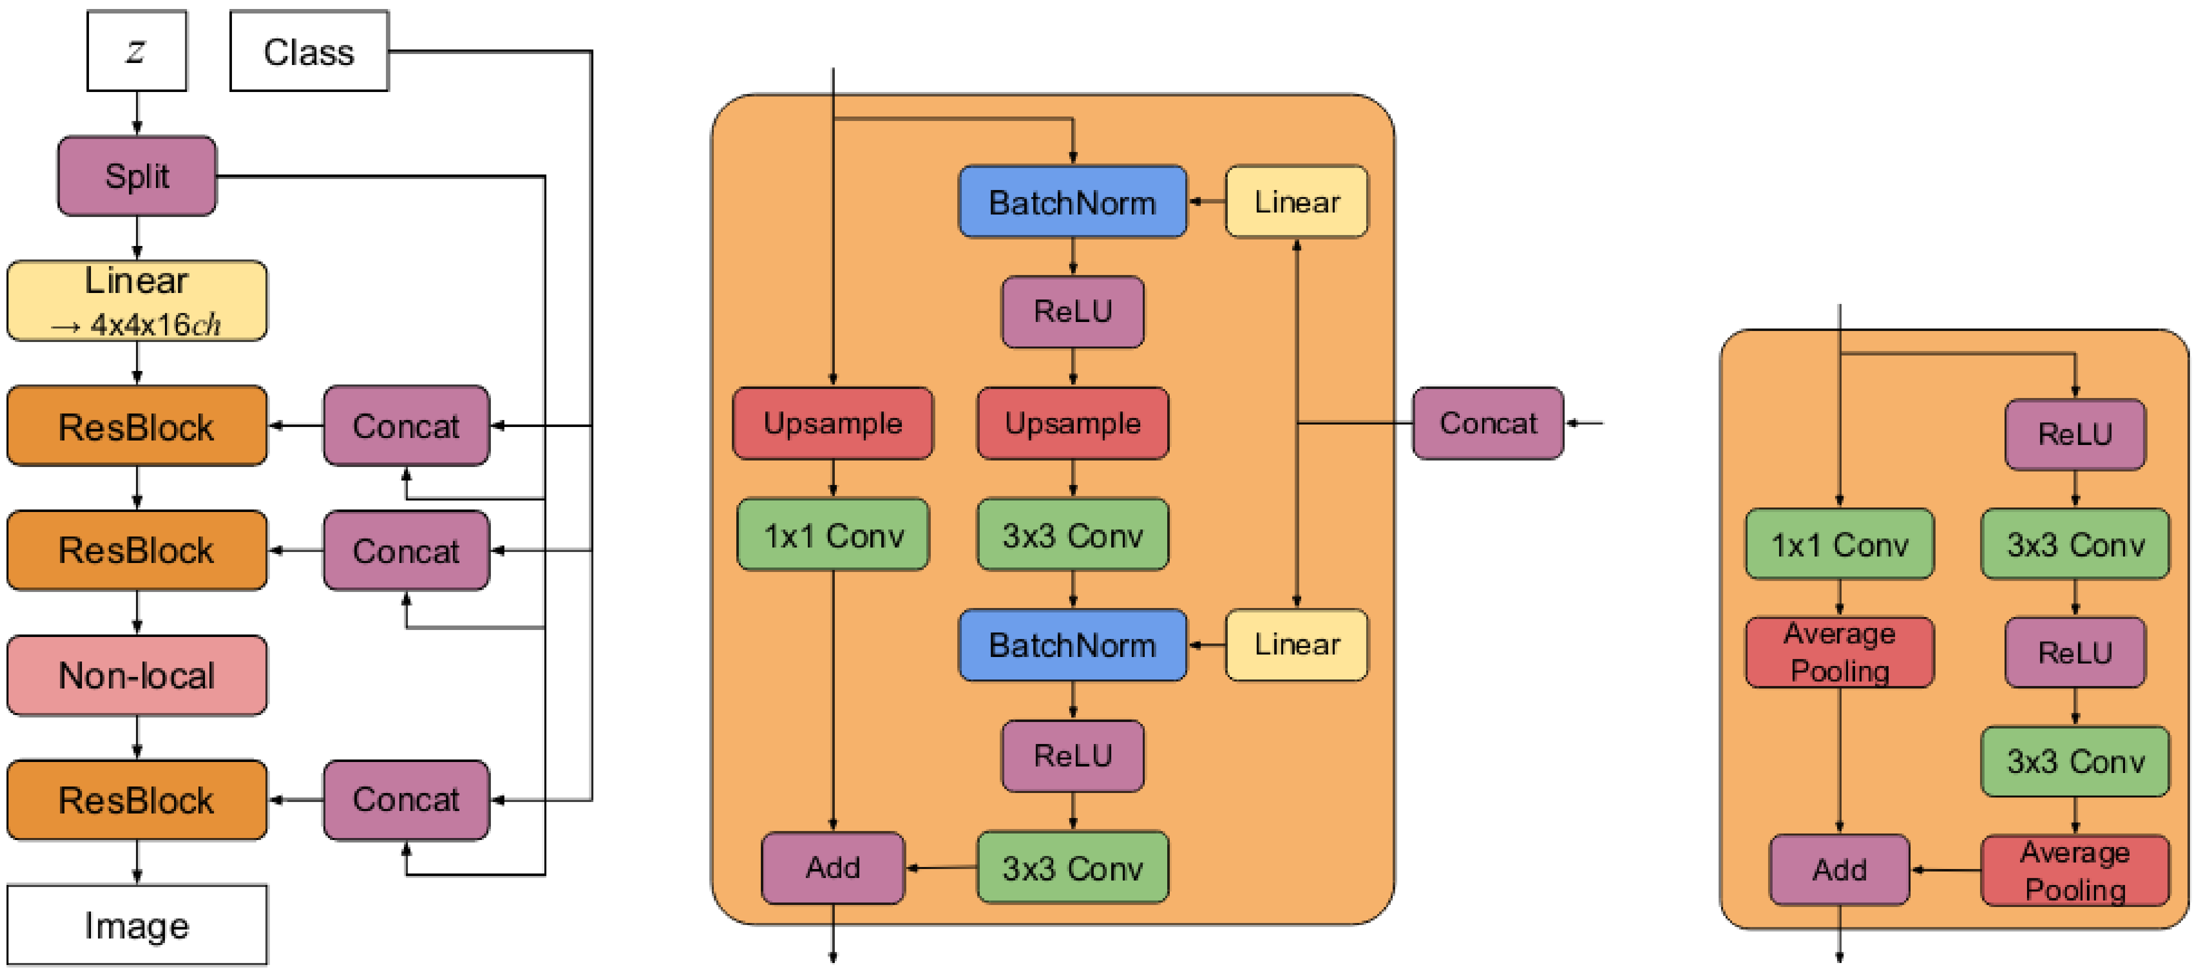

(a) (b) (c)

[A Brock, et al, ICLR 2019]

Modern Architectures

BigGAN

On 8xV100 with full-precision training (no Tensor cores), this script takes 15 days to train to 150k iterations.

Vid-to-vid translation


DEPTH
AHAM
SINGER

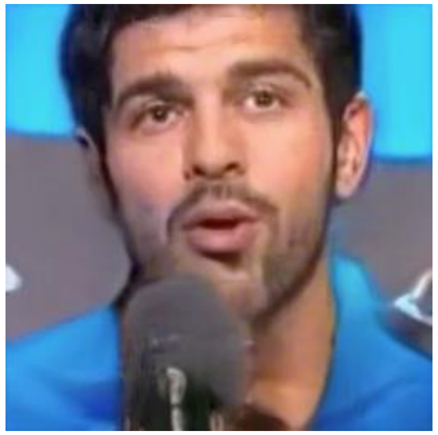

Vid-to-vid translation

[Carolin Chan, et al, ICCV 2019]

- Everybody dance now

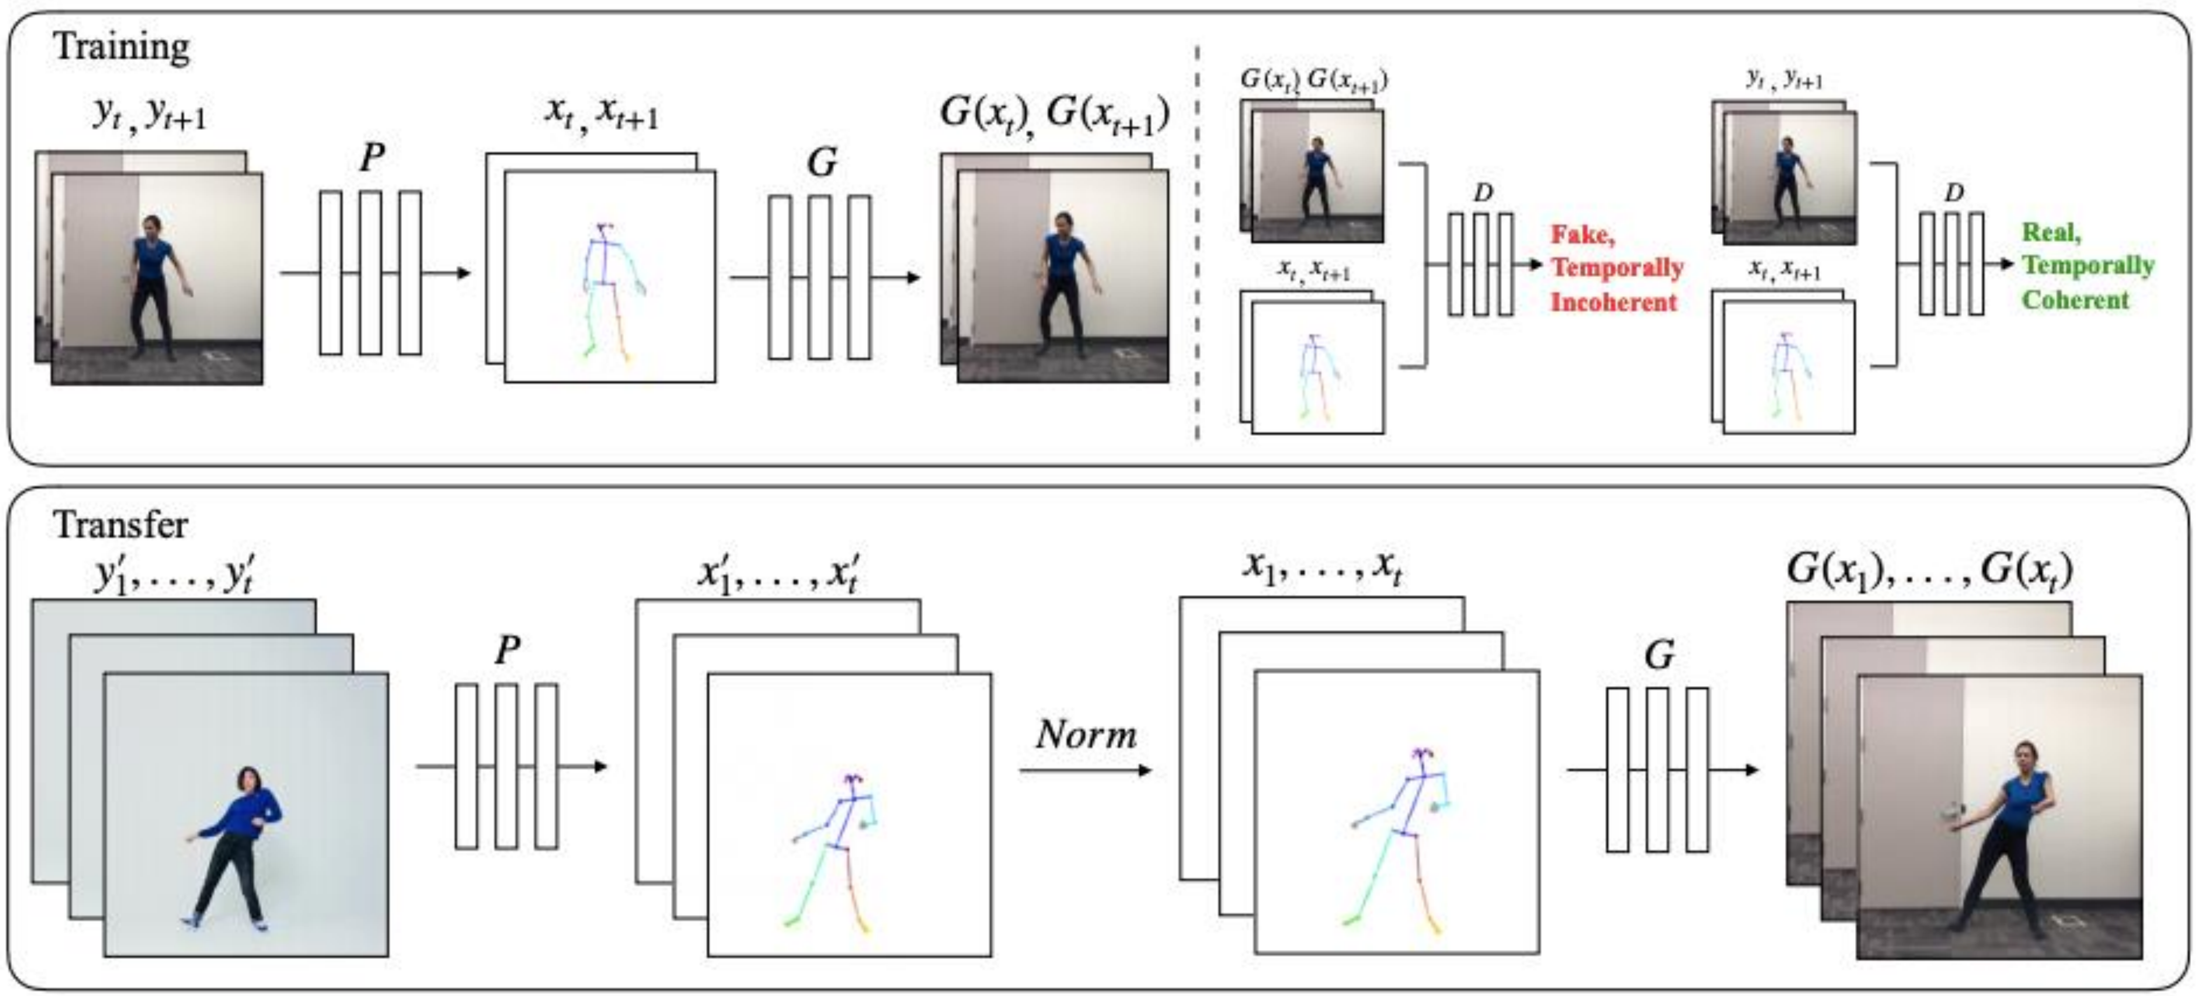

https://www.youtube.com/watch?v=PCBTZh41Ris

Vid-to-vid translation

- Video-to-video synthesis

[Ting-chun Wang, et al, NIPS 2018]

https://github.com/NVIDIA/vid2vid

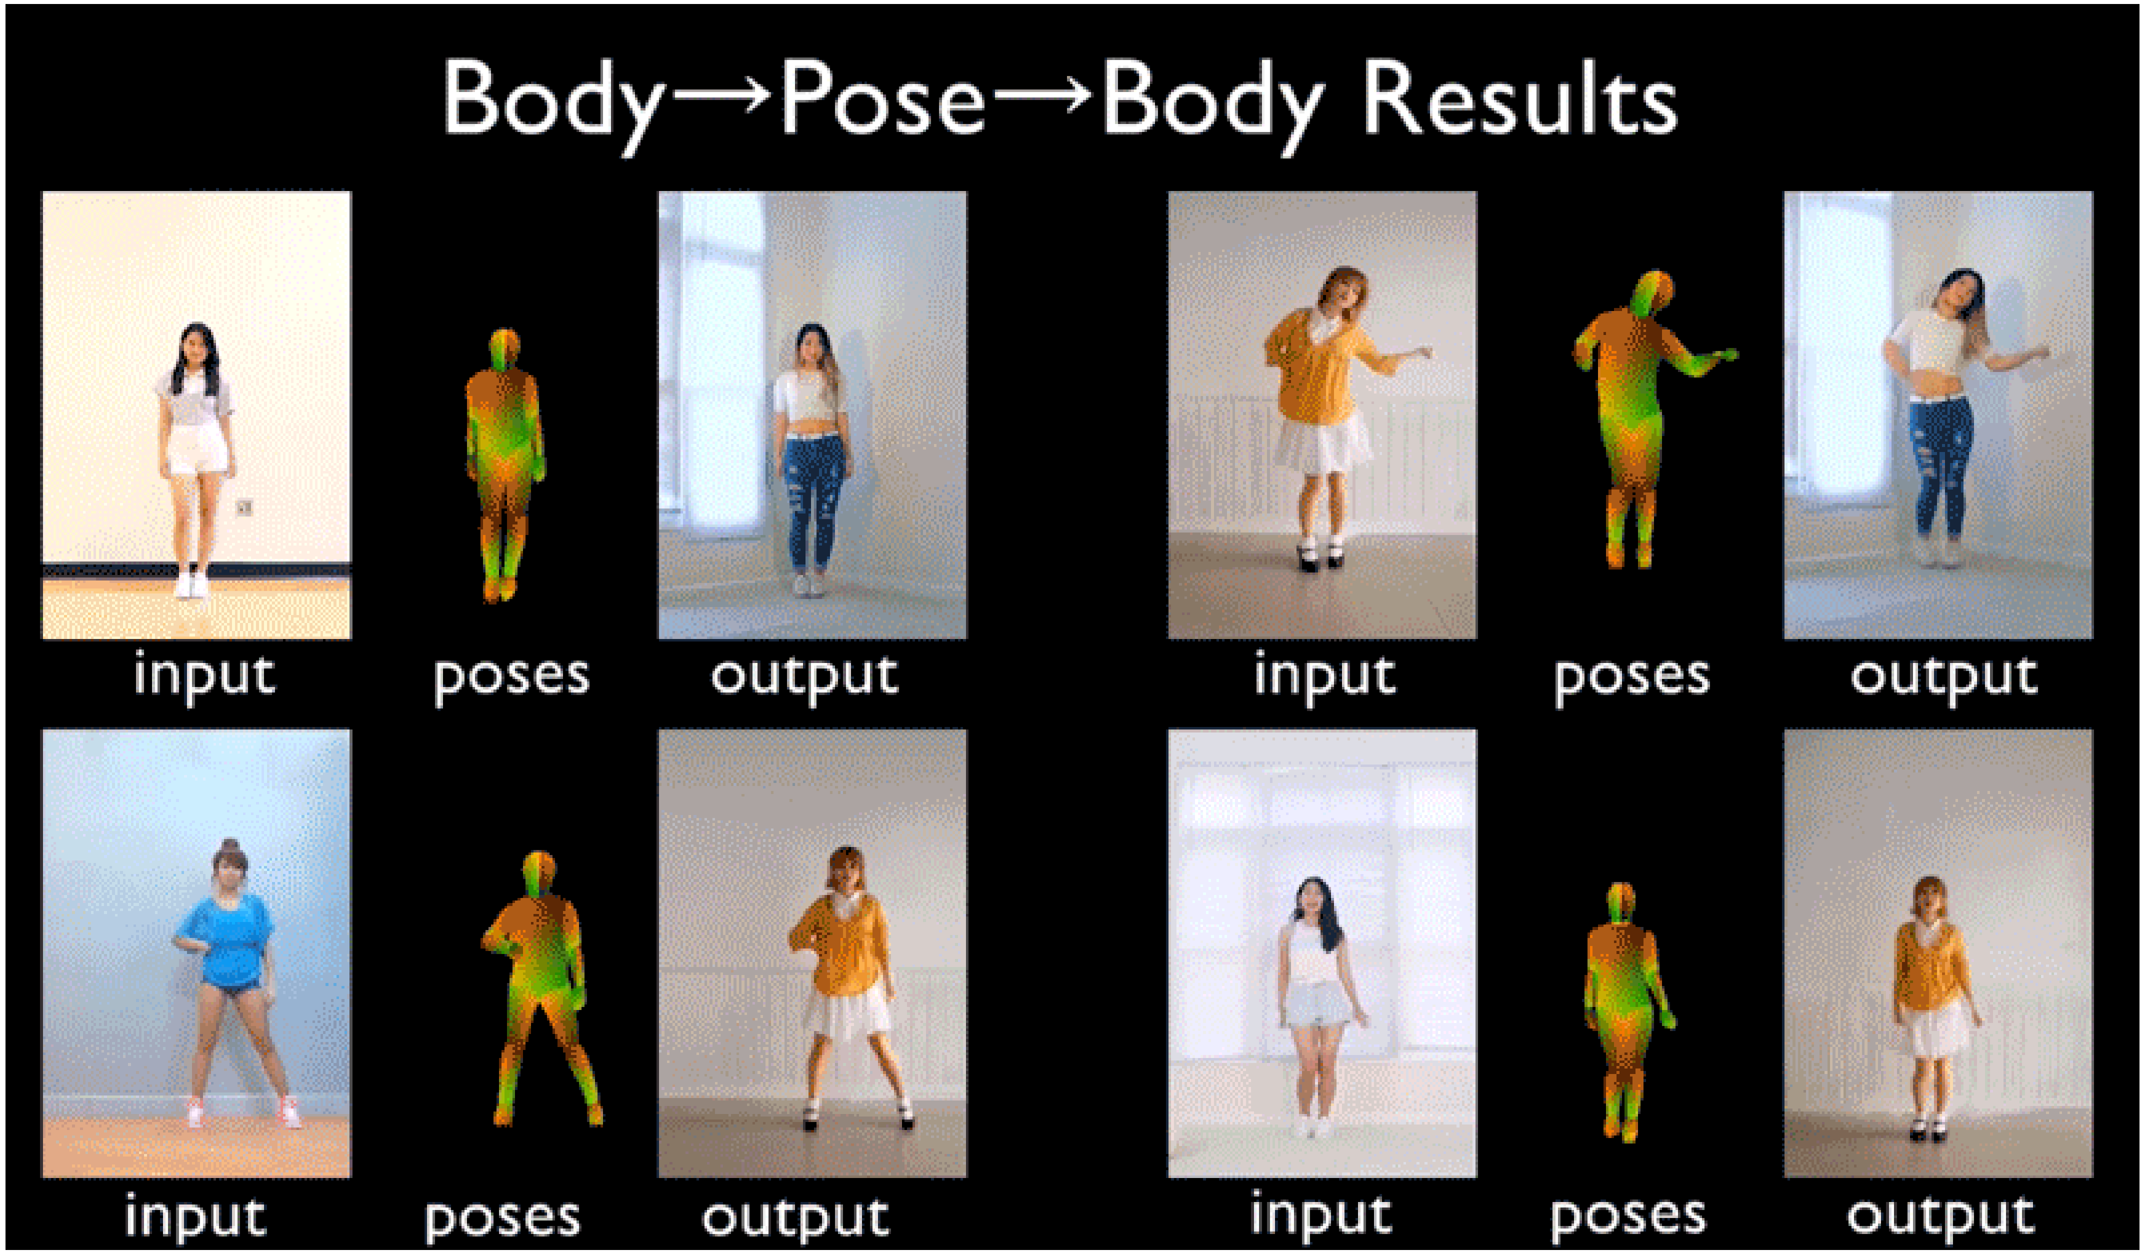

Video Generation

Video Generation

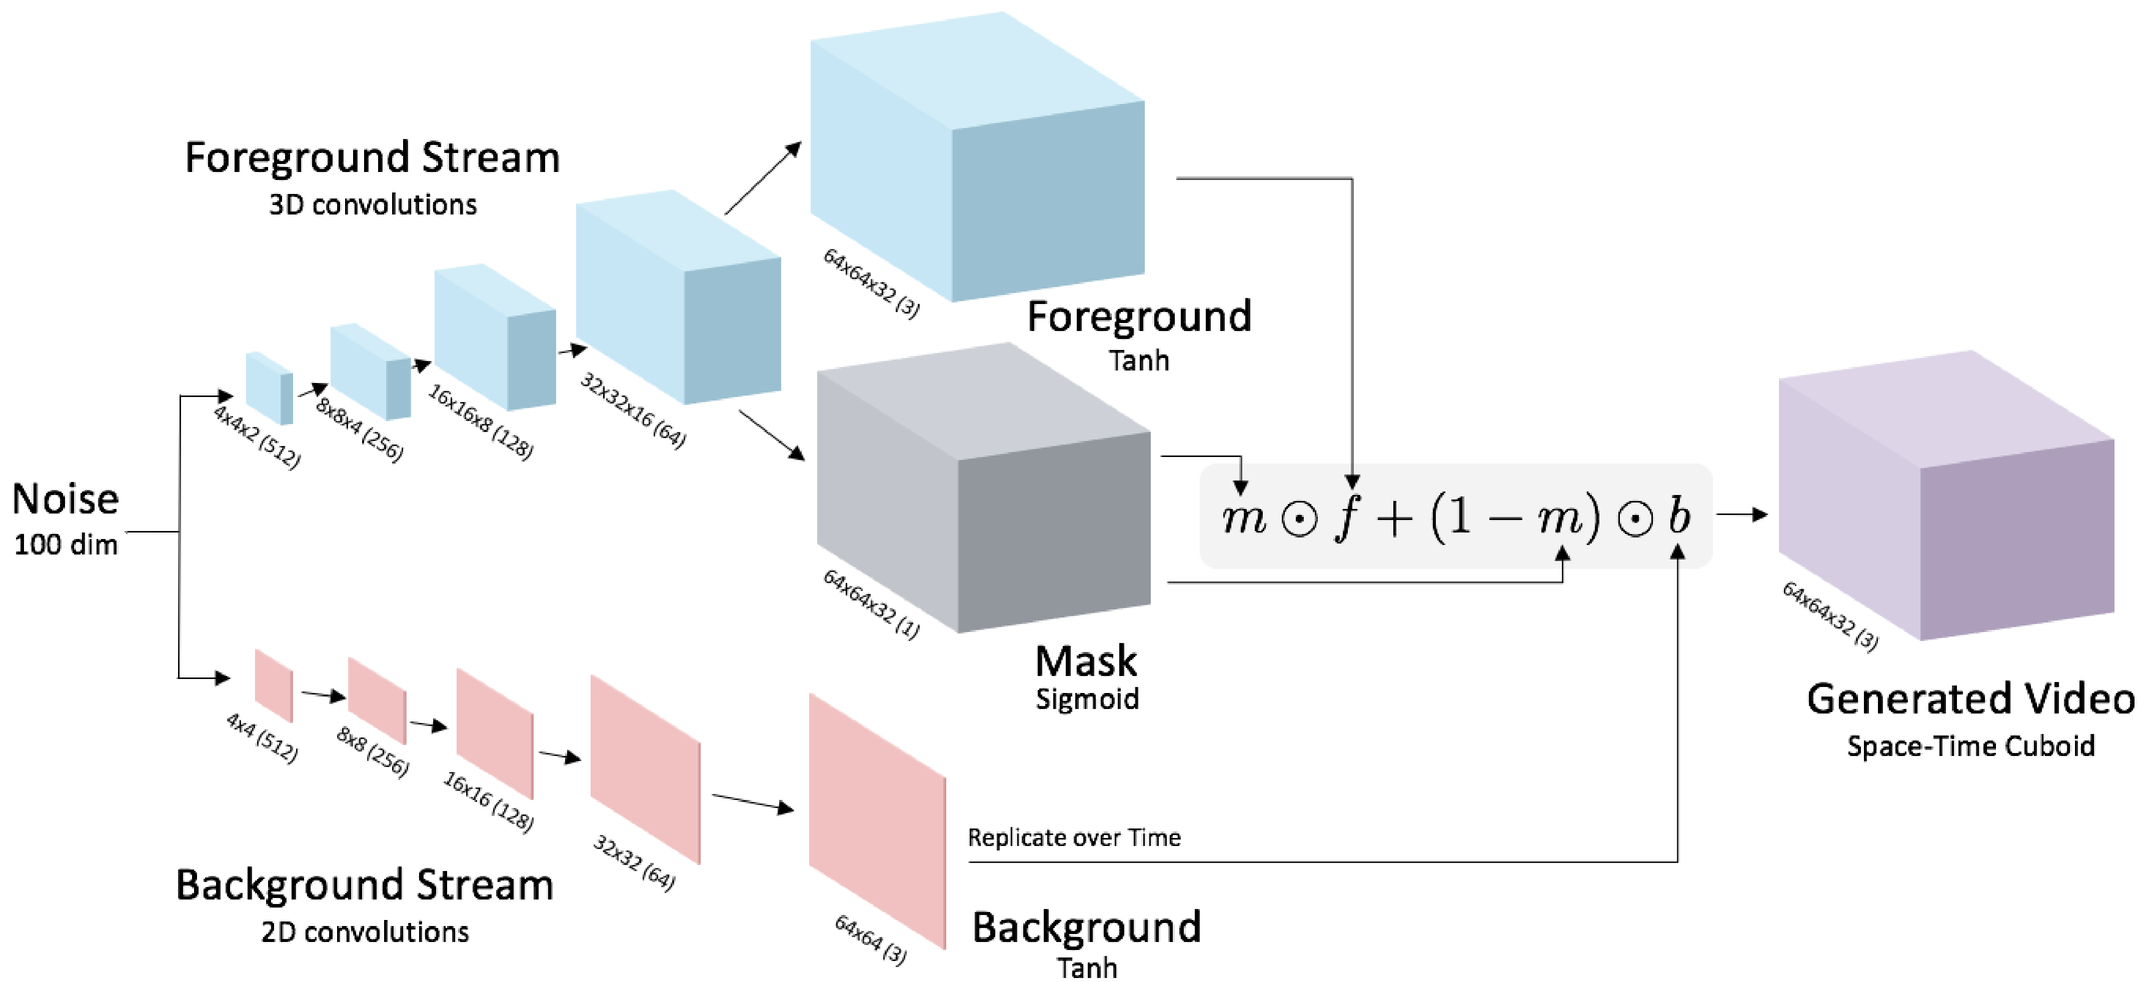

VGAN [NeurIPS'16]

Video Generation

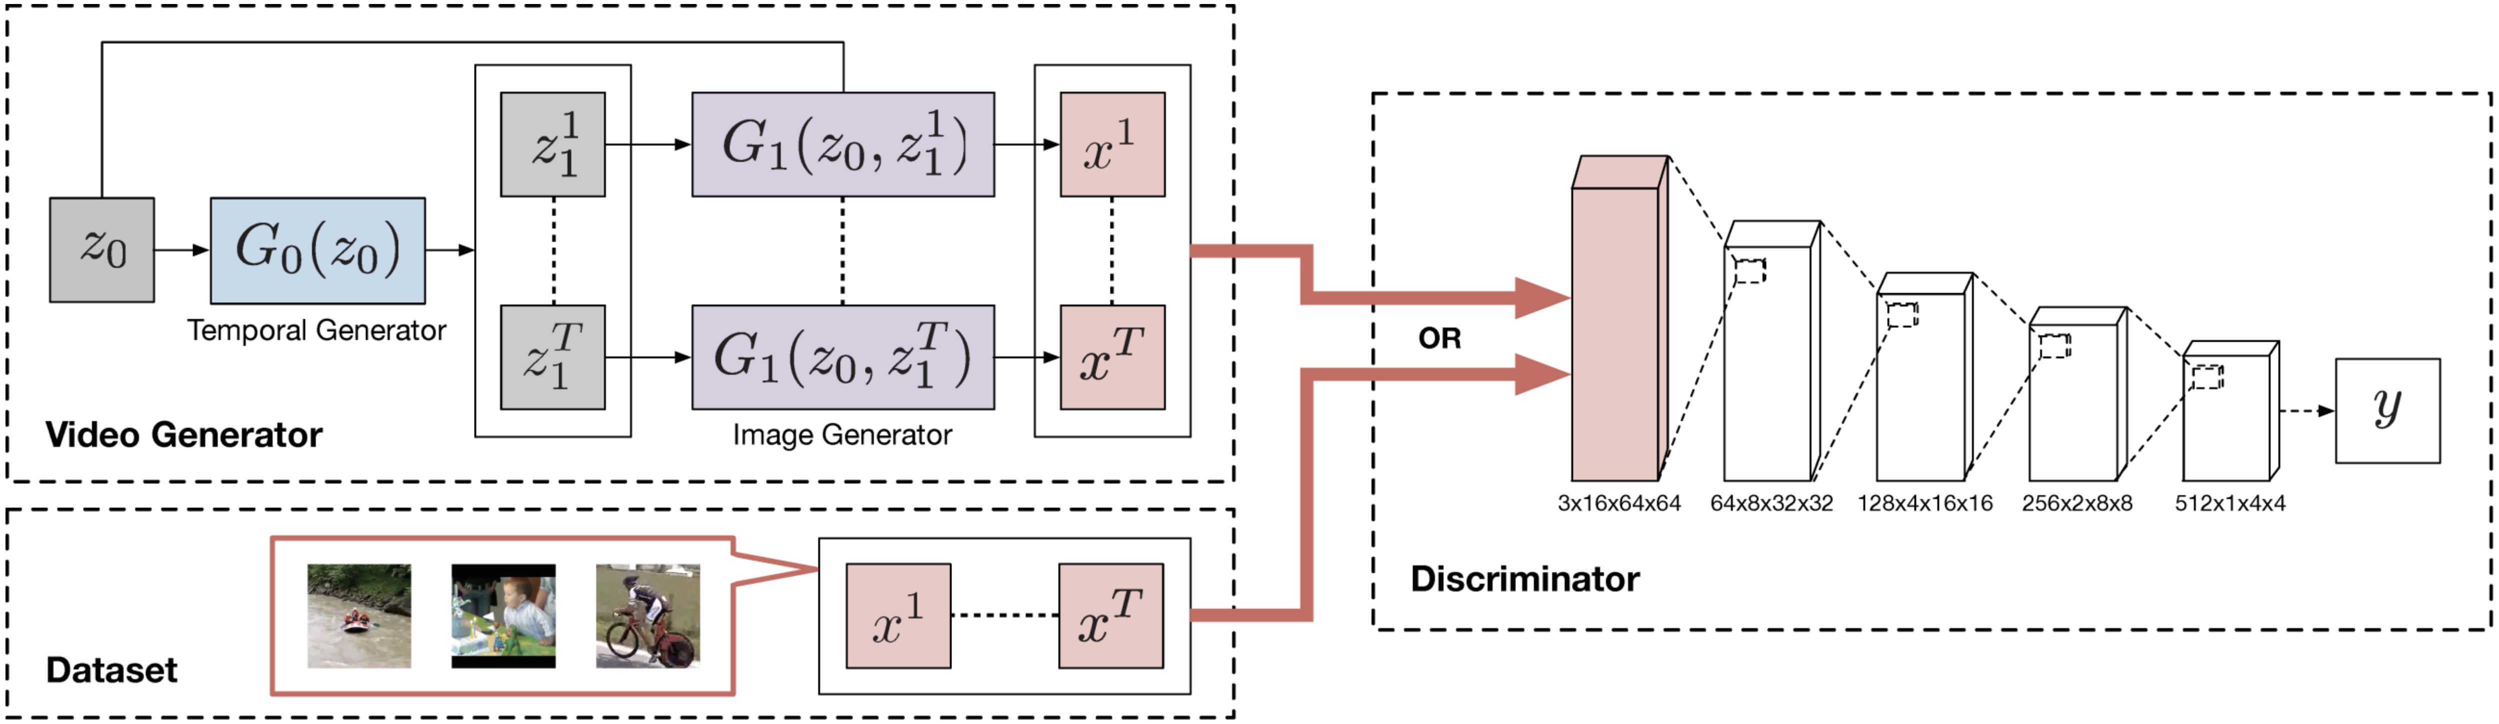

TGAN [ICCV'19]

Video Generation

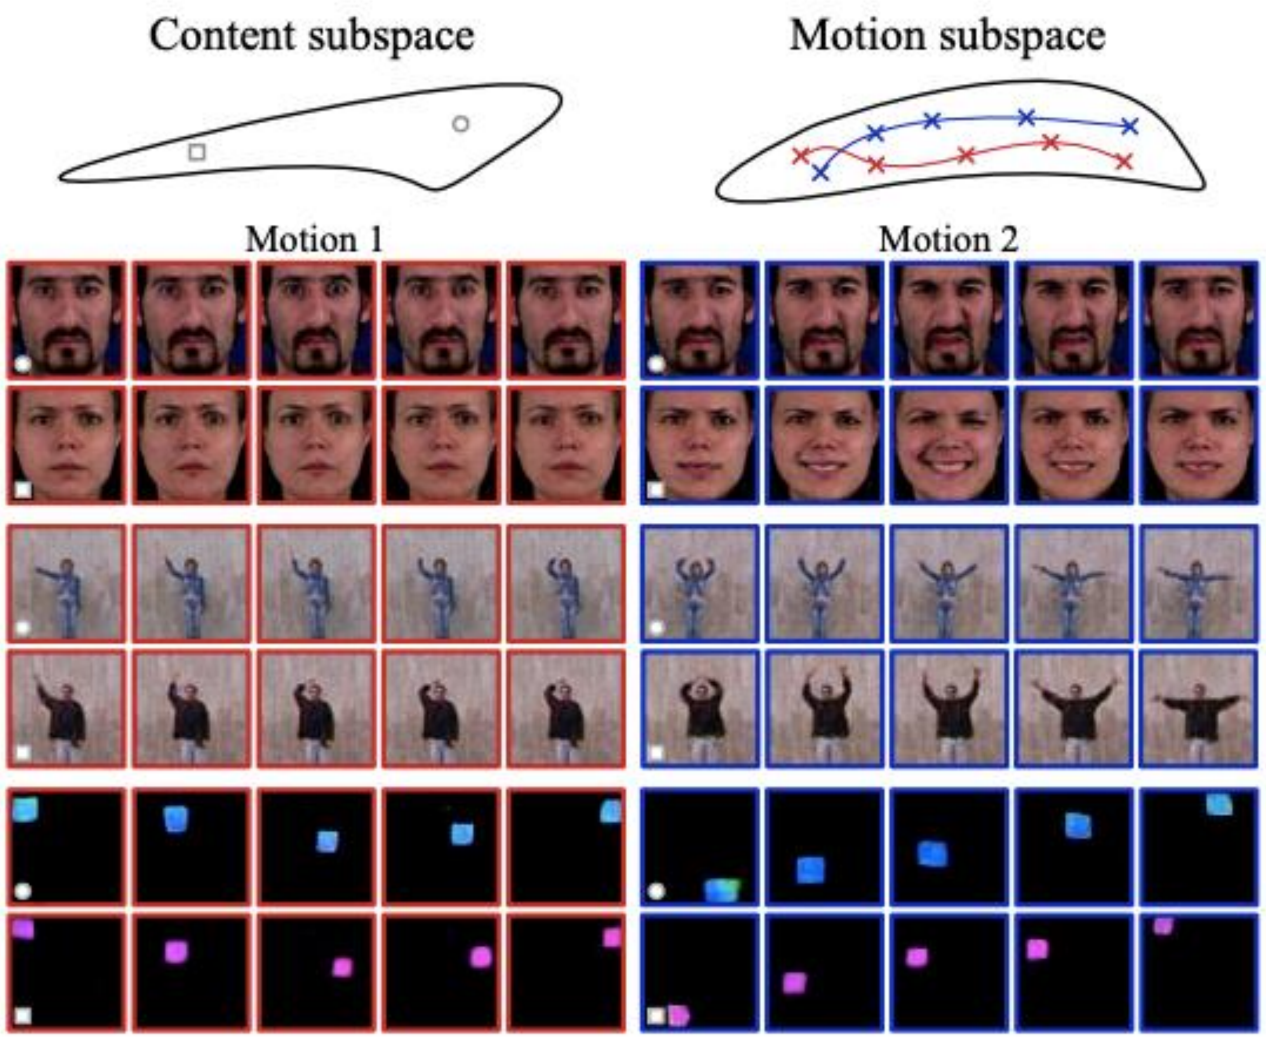

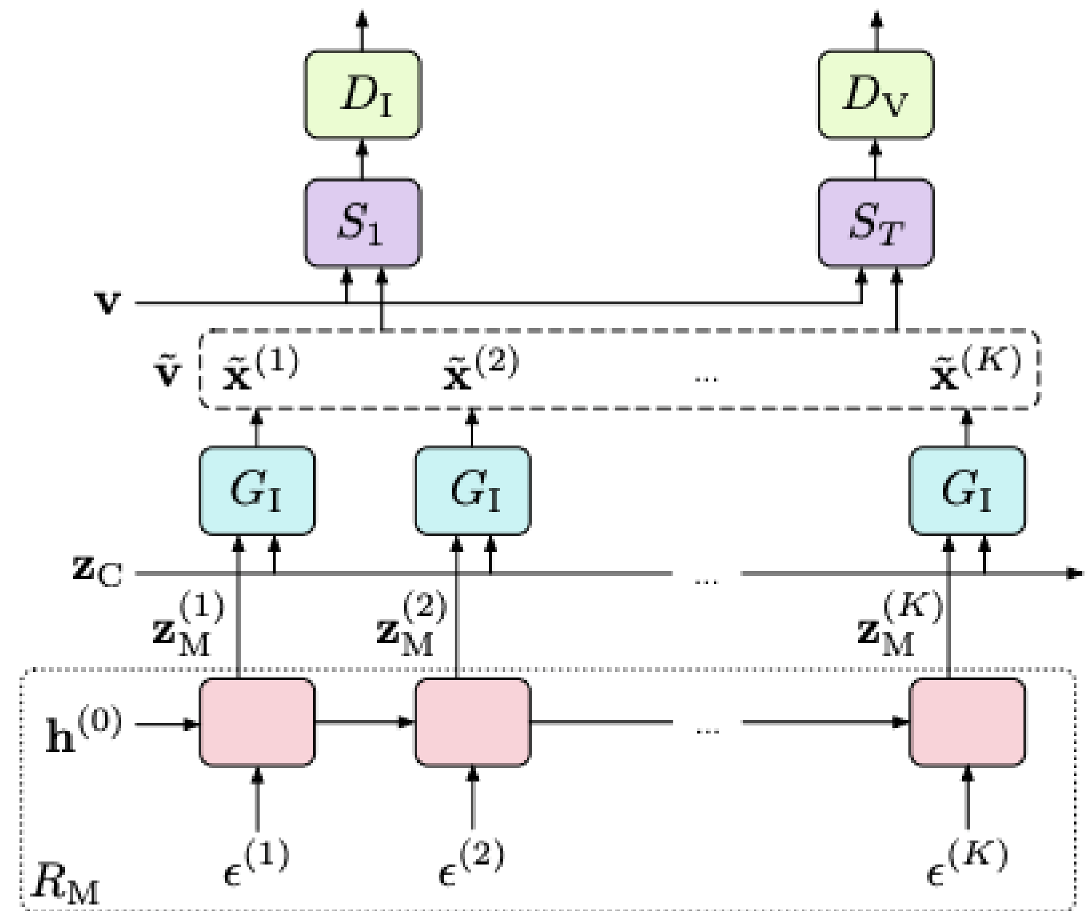

MoCoGAN [CVPR’19]

Video Generation

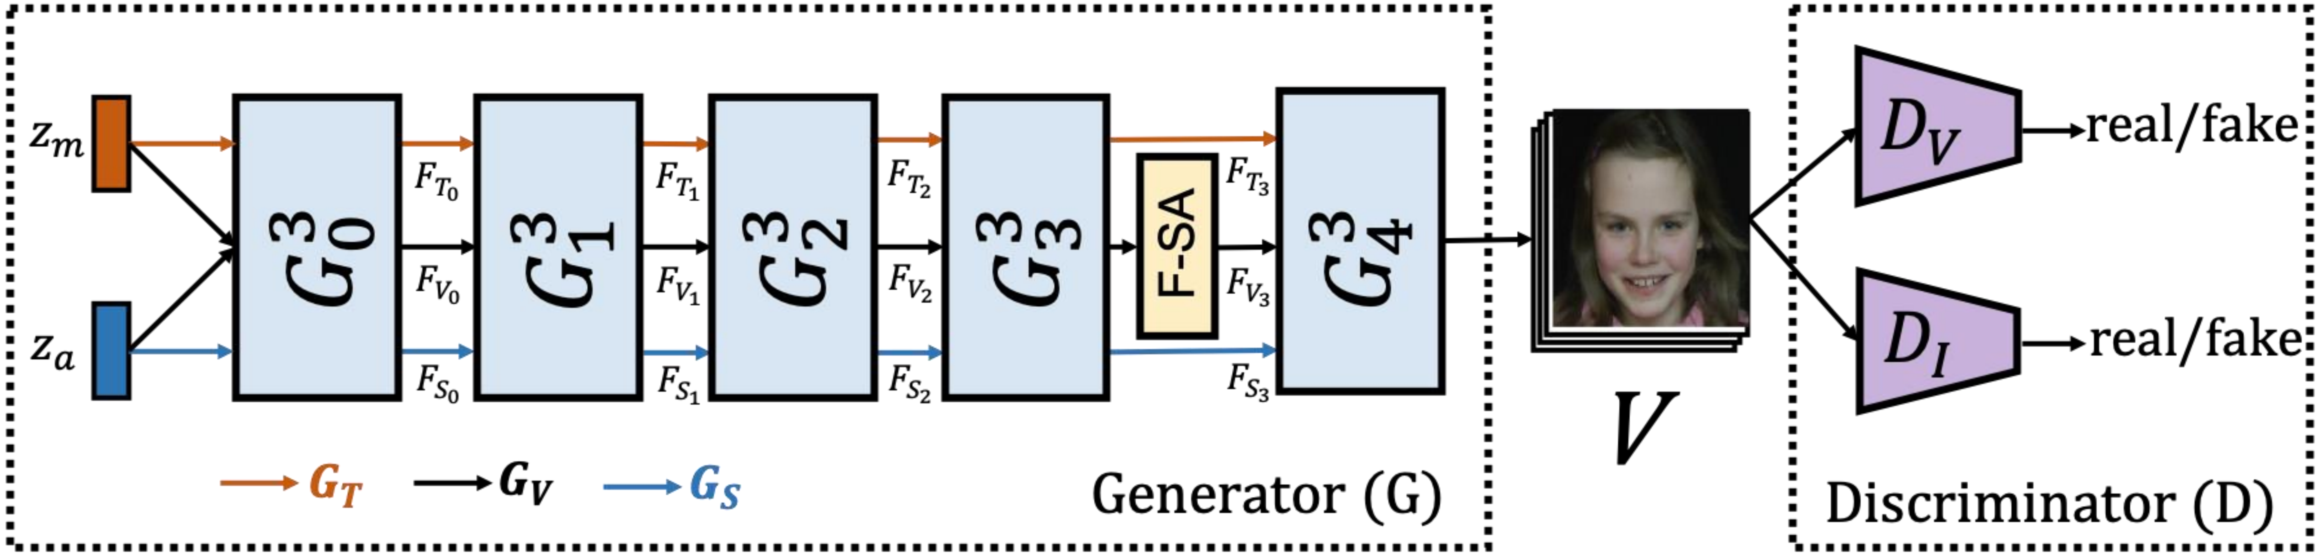

G3AN [CVPR'20]

What I can not create, I do not understand

- R. Feynman

Thank You !