The background of the slide features a large, semi-transparent watermark of the University of Pisa crest. The crest is circular and contains a central figure, likely a personification of Justice or Liberty, surrounded by Latin text. The watermark is rendered in a dark blue color that blends with the dark blue background of the slide.

Implicit models – Adversarial Learning

INTELLIGENT SYSTEMS FOR PATTERN RECOGNITION (ISPR)

DAVIDE BACCIU – DIPARTIMENTO DI INFORMATICA - UNIVERSITA' DI PISA

DAVIDE.BACCIU@UNIFI.IT

Generative Learning from a DL Perspective

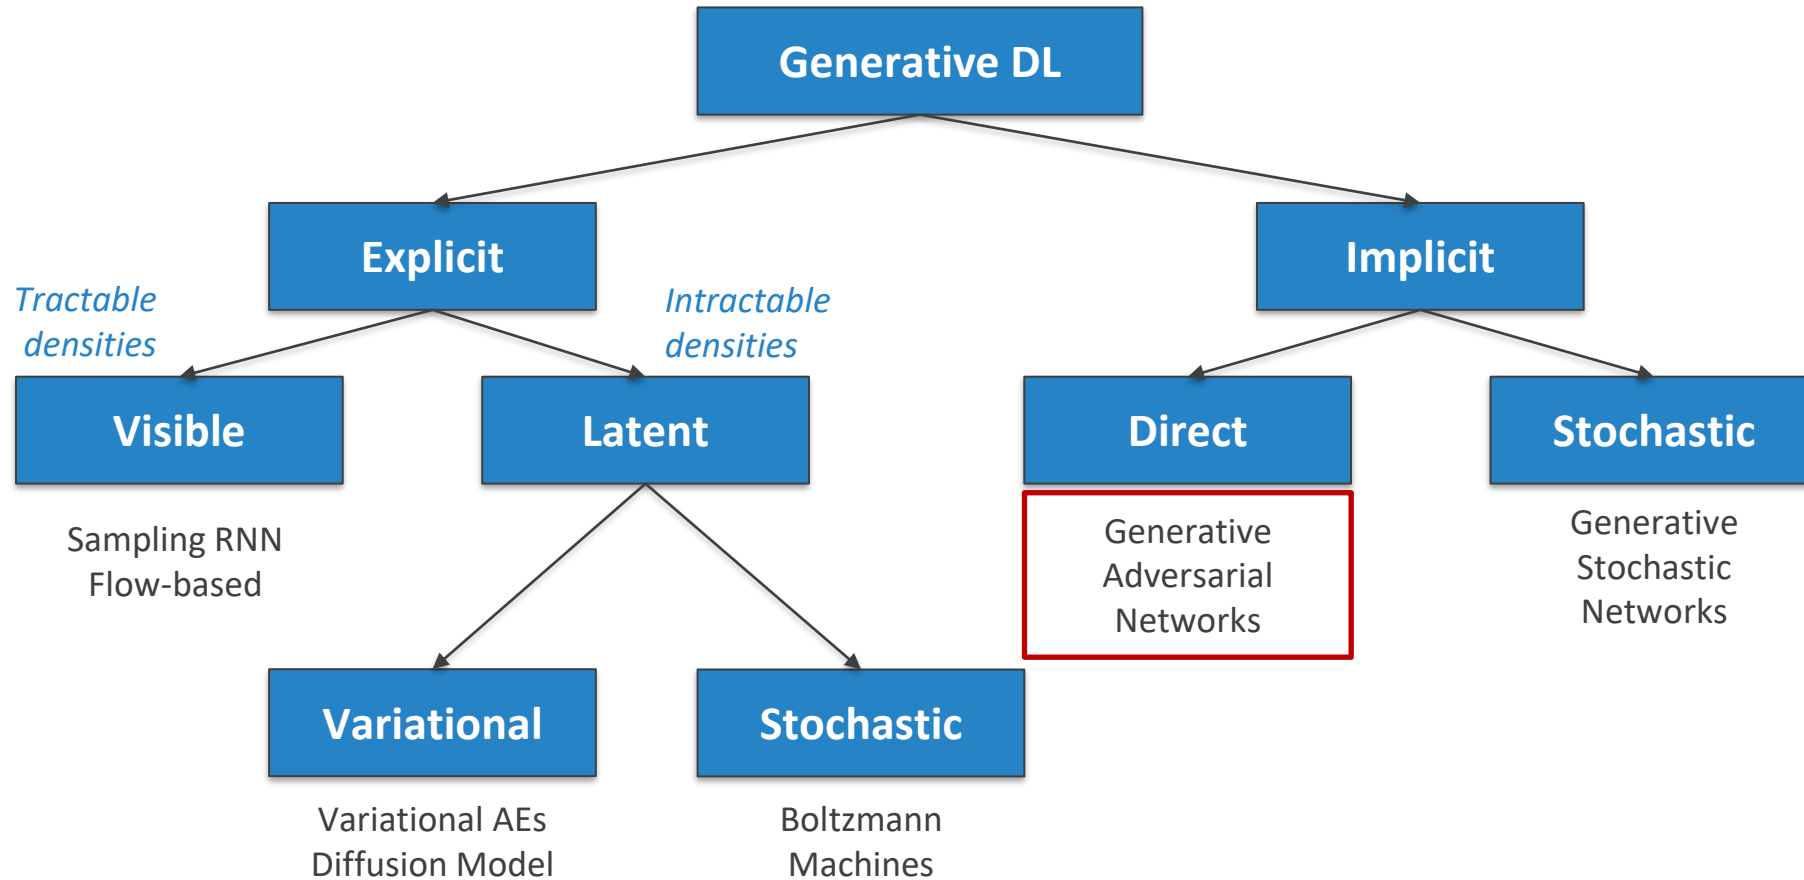
*Given training data, learn a (deep) neural network that **can generate new samples** from (an approximation of) the data distribution*

Two approaches

- **Explicit** \Rightarrow Learn a model density $P_{\theta}(x)$
- **Implicit** \Rightarrow Learn a process that samples data from $P_{\theta}(x) \approx P(x)$



Our Taxonomy



Adapted from I. Goodfellow, Tutorial on Generative Adversarial Networks, 2017

Distribution Learning Vs Learning to Sample

- Variational AEs learn to approximate an intractable distribution

$$P_{\theta}(\mathbf{x}) = \int P_{\theta}(\mathbf{x}|\mathbf{z})P(\mathbf{z})d\mathbf{z}$$

then sample it to generate the output

- What if we learn to generate samples rather than learning the distribution?
 - Generative Adversarial Networks (GAN)
 - Game theoretic approach



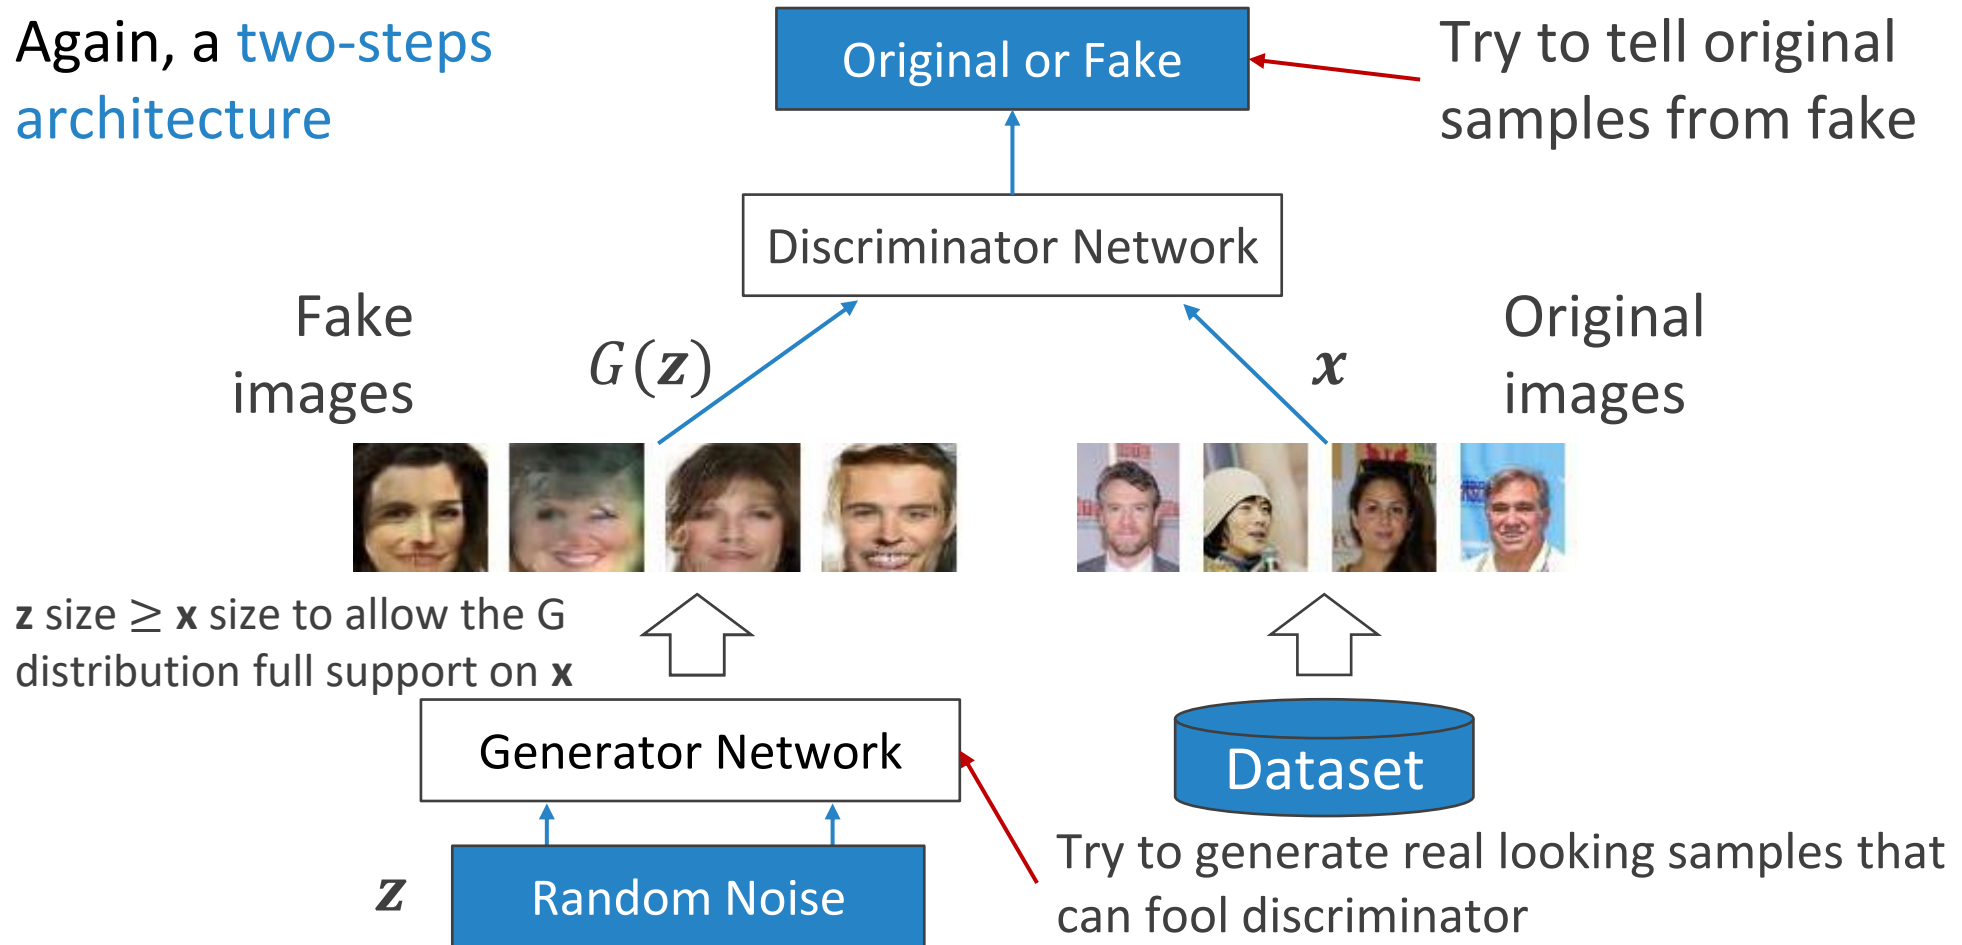
The GAN Catch

- We need to learn to sample from a complex, high-dimensional training distribution
 - No straightforward way to do this
- The **catch**
 - Sample from a simple distribution: **random noise**
 - Train a differentiable deterministic function (neural network) to **transform random noise to the training distribution**



Generative Adversarial Networks

Again, a **two-steps architecture**



Alternate Optimization

$$C = \min_{\theta_G} \max_{\theta_D} \left[\underbrace{\mathbb{E}_x [\log D_{\theta_D}(x)]}_{\text{Discriminator output for real data } x} + \mathbb{E}_z [\log(1 - \underbrace{D_{\theta_D}(G_{\theta_G}(z))}_{\text{Discriminator output for fake data } G(z)})] \right]$$

- Discriminator output is likelihood of input being real
- Discriminator tries to maximize C s.t.
 - $D_{\theta_D}(x) \rightarrow 1$ and $D_{\theta_D}(G_{\theta_G}(z)) \rightarrow 0$
- Generator tries to minimize C s.t.
 - $D_{\theta_D}(G_{\theta_G}(z)) \rightarrow 1$

Alternate Optimization

$$C = \min_{\theta_G} \max_{\theta_D} \left[\mathbb{E}_x [\log D_{\theta_D}(x)] + \mathbb{E}_z [\log(1 - D_{\theta_D}(G_{\theta_G}(z)))] \right]$$

1. Discriminator **gradient ascent**

$$C_D = \max_{\theta_D} \left[\mathbb{E}_x [\log D_{\theta_D}(x)] + \mathbb{E}_z [\log(1 - D_{\theta_D}(G_{\theta_G}(z)))] \right]$$

2. Generator **gradient descent**

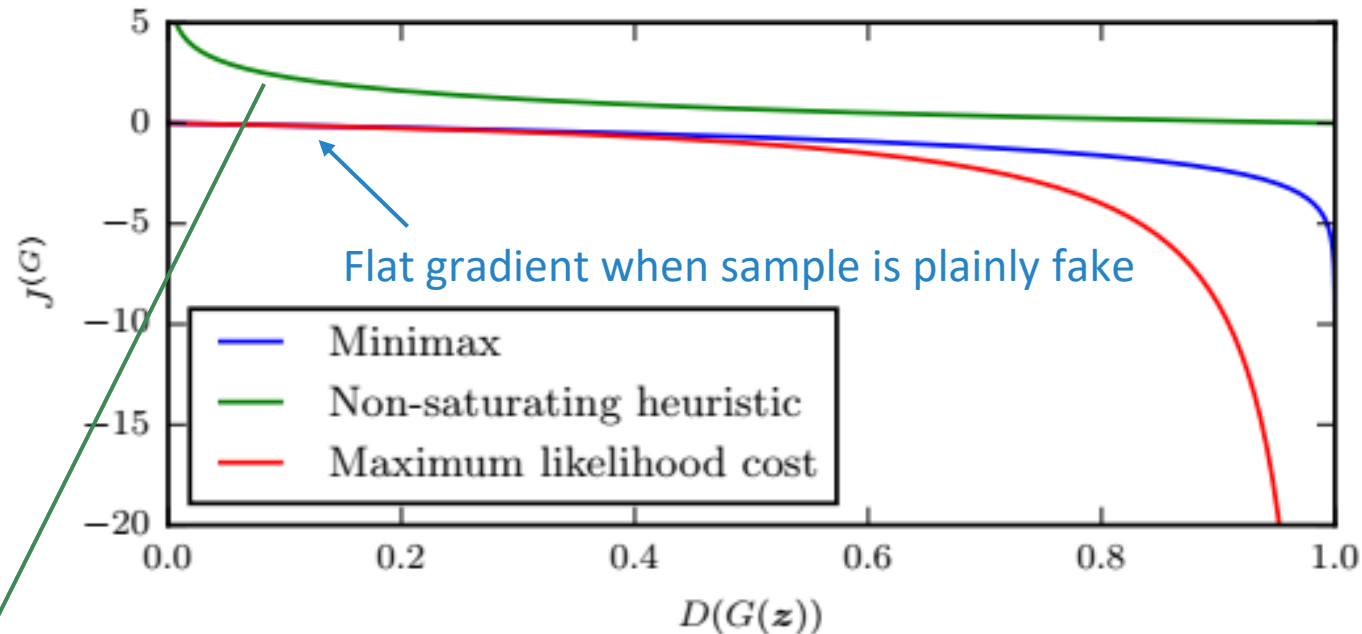
$$C_G = \min_{\theta_G} \left[\mathbb{E}_z [\log(1 - D_{\theta_D}(G_{\theta_G}(z)))] \right]$$

Optimizing this doesn't really work



The Issue and a Solution

The **cost** that the Generator receives in response to generate $G(\mathbf{z})$ depends only on the **Discriminator** response



$$C_G = \max_{\theta_G} \left[\mathbb{E}_z \left[\log(D_{\theta_D}(G_{\theta_G}(z))) \right] \right] \quad \text{maximize likelihood of discriminator being wrong}$$



GAN Training Pseudo-Algorithm

for number of training iterations **do**

for k steps **do**

Stability trick

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

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m [\log D_{\theta_d}(x^{(i)}) + \log(1 - D_{\theta_d}(G_{\theta_g}(z^{(i)})))]$$

end for

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

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log(D_{\theta_d}(G_{\theta_g}(z^{(i)})))$$


end for

Expectation



A Hard Two-Player Game

- The optimal solution of the min-max problem is a saddle point
- Little stability
 - Initially lot of heuristic work
 - Now converged to more principled solutions



Wasserstein Distance Models

Attempts to solve the hardness of training adversarial generators by **optimizing the Wasserstein distance** (EMD) between the generator and empirical distribution filtered through the discriminator function D

$$\begin{aligned} G^* &= \operatorname{argmin}_G W(\mu, \mu_G) \\ &= \operatorname{argmin}_G \sup_{\|D\|_L \leq 1} \left[\mathbb{E}_{x \sim \mu} [D(x)] - \mathbb{E}_{x \sim \mu_G} [D(x)] \right] \end{aligned}$$

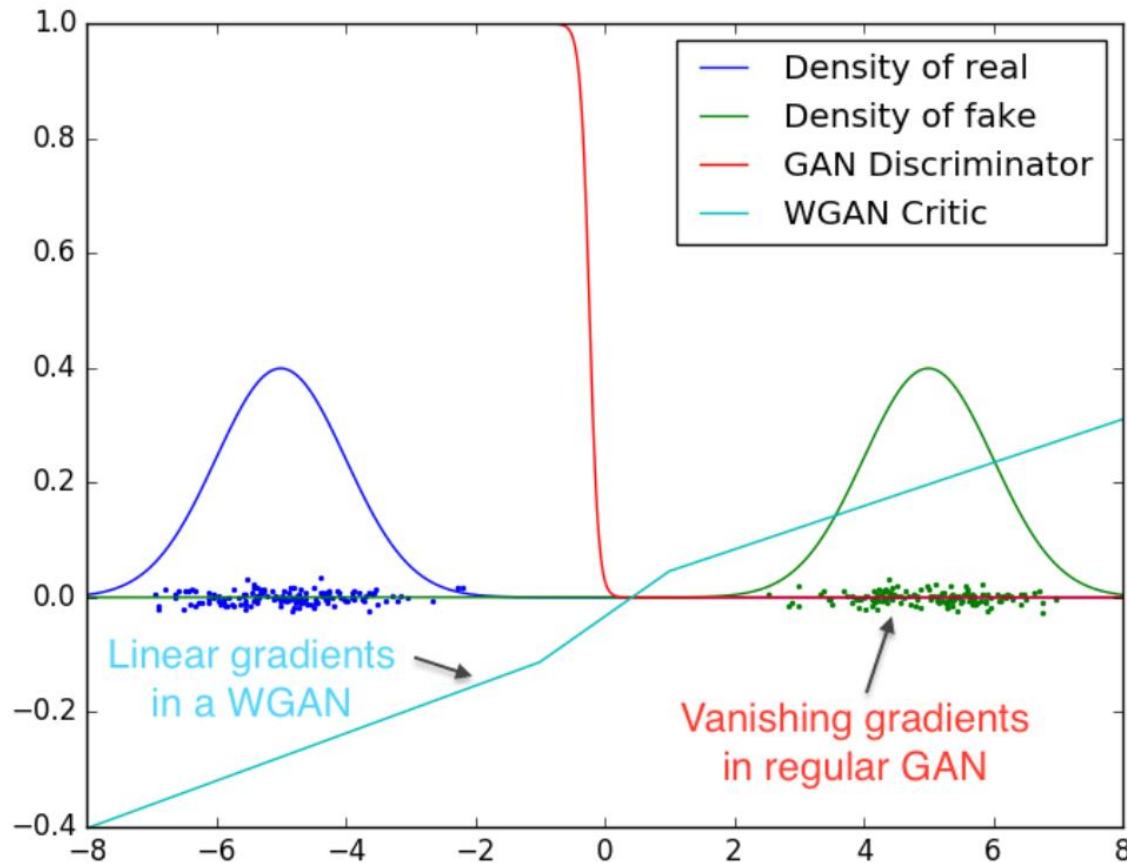
$$\|D\|_L \leq 1$$

Requires **optimizing D under a constraint** on Lipschitz seminorm

- Clipping D weights (slow to converge)



Effect of the Wasserstein Loss



- Classical GAN loss results in saturation (discriminator with zero loss)
- WGAN provide gradients across all the range of training conditions

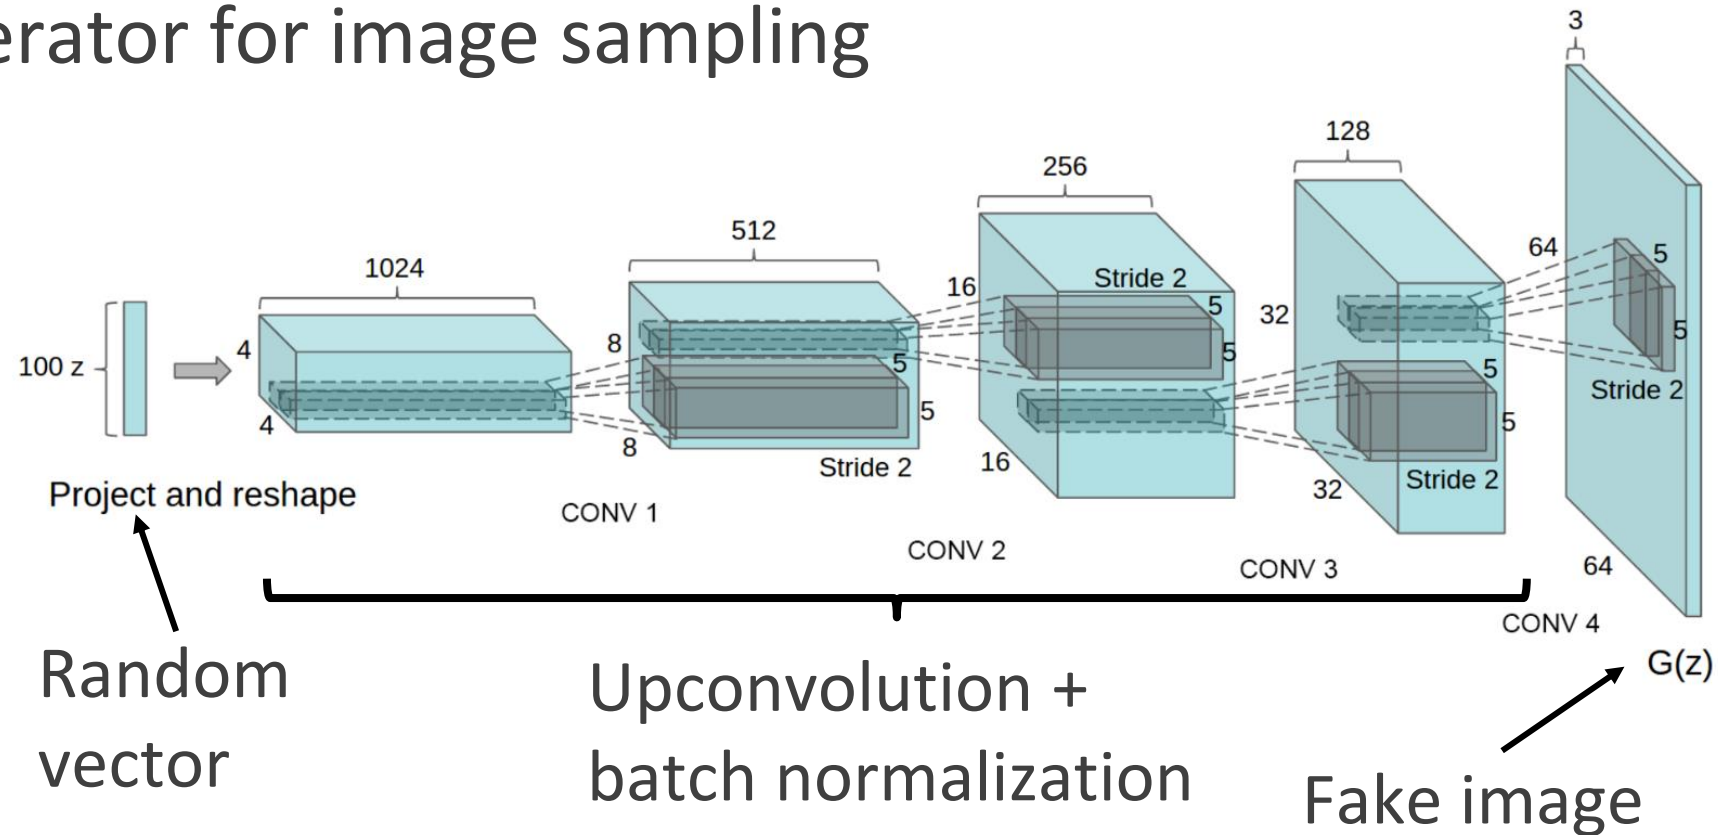
<https://arxiv.org/pdf/1701.07875.pdf>



UNIVERSITÀ DI PISA

The DCGAN Architecture

Generator for image sampling



Radford et al, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", ICLR 2016



UNIVERSITÀ DI PISA

GAN and Images



Radford et al, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", ICLR 2016

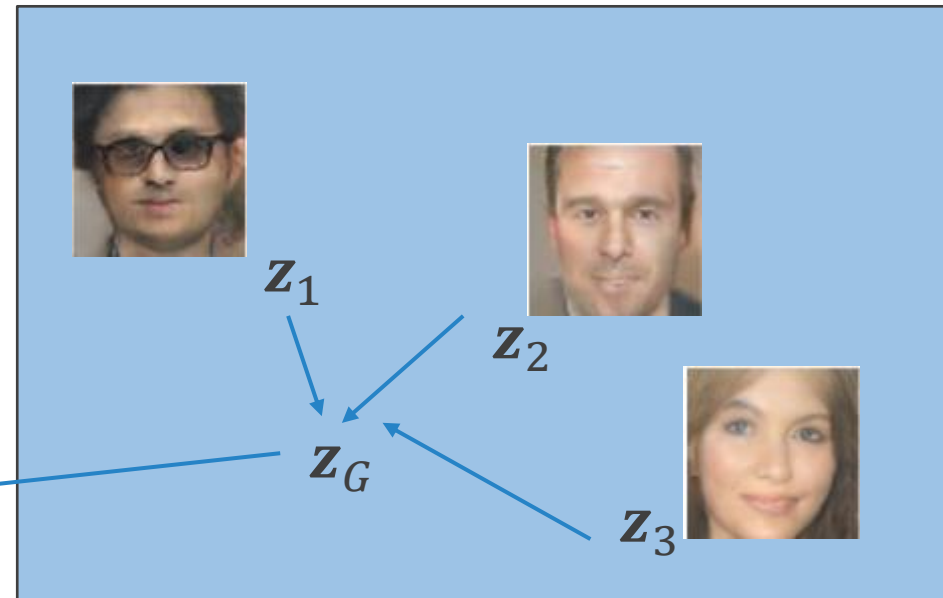


UNIVERSITÀ DI PISA

Latent Space Arithmetic

Can do sensible linear operations on noise vectors (arithmetic, interpolation)

$$z_G = z_1 - z_2 + z_3$$



Radford et al, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", ICLR 2016



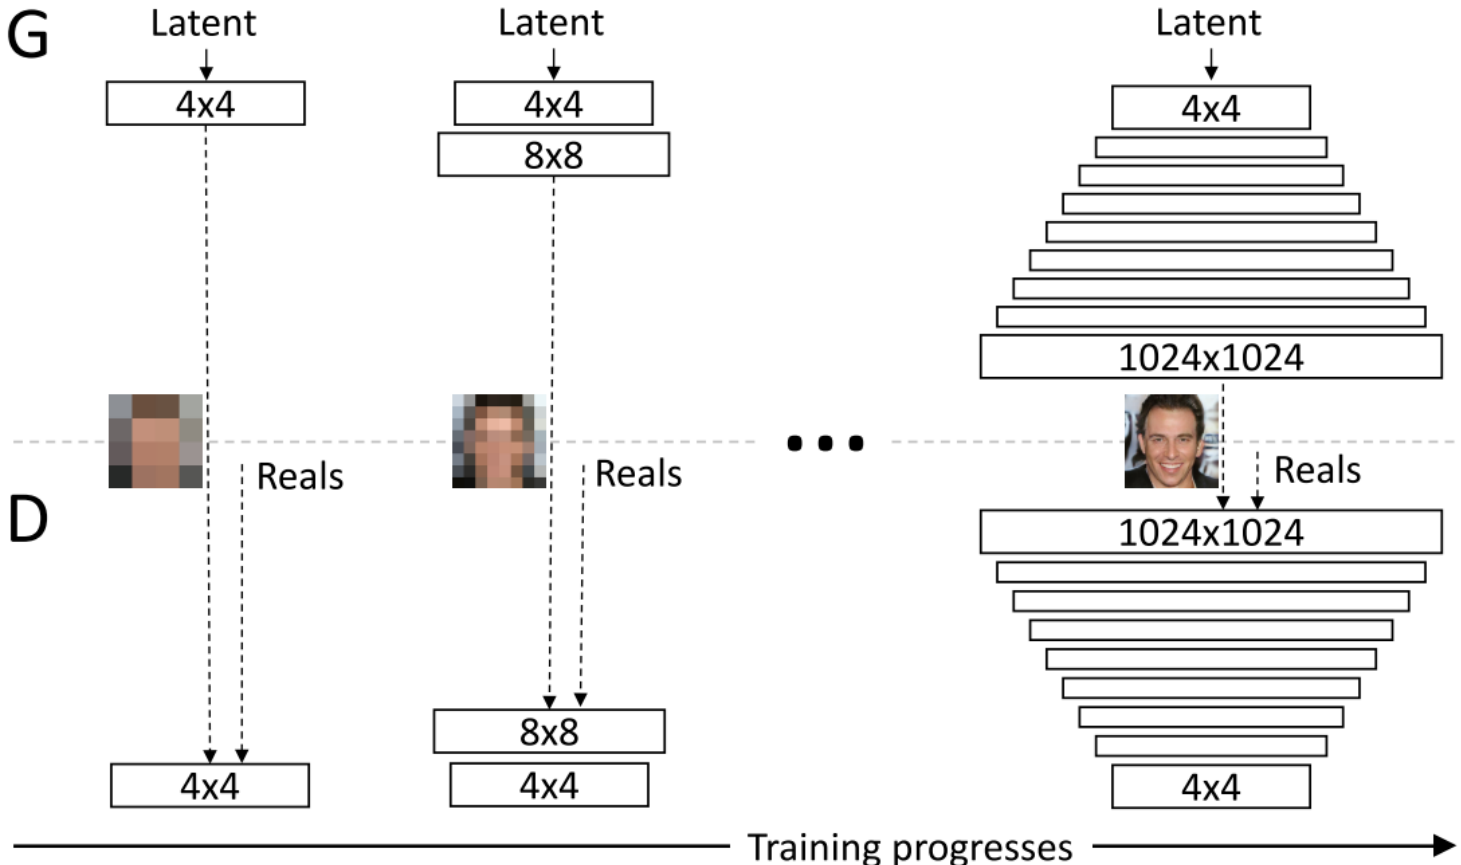
UNIVERSITÀ DI PISA

How to Get High Resolution Samples?



<https://arxiv.org/abs/1710.10196>

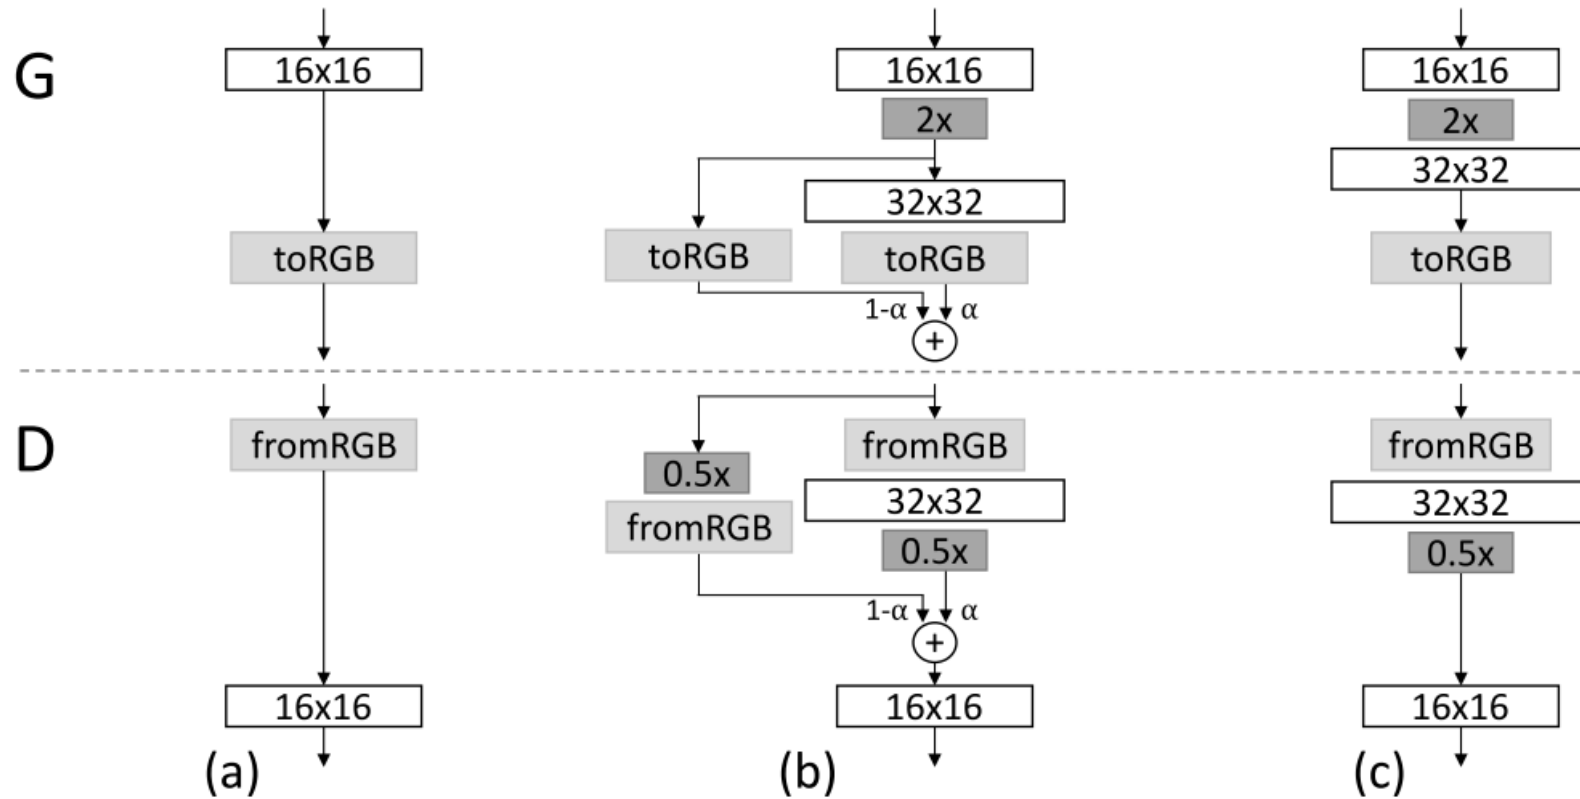
Progressive GAN



<https://arxiv.org/abs/1710.10196>



Progressive GAN – Smooth Transition



<https://arxiv.org/abs/1710.10196>




UNIVERSITÀ DI PISA

Conditional Generation

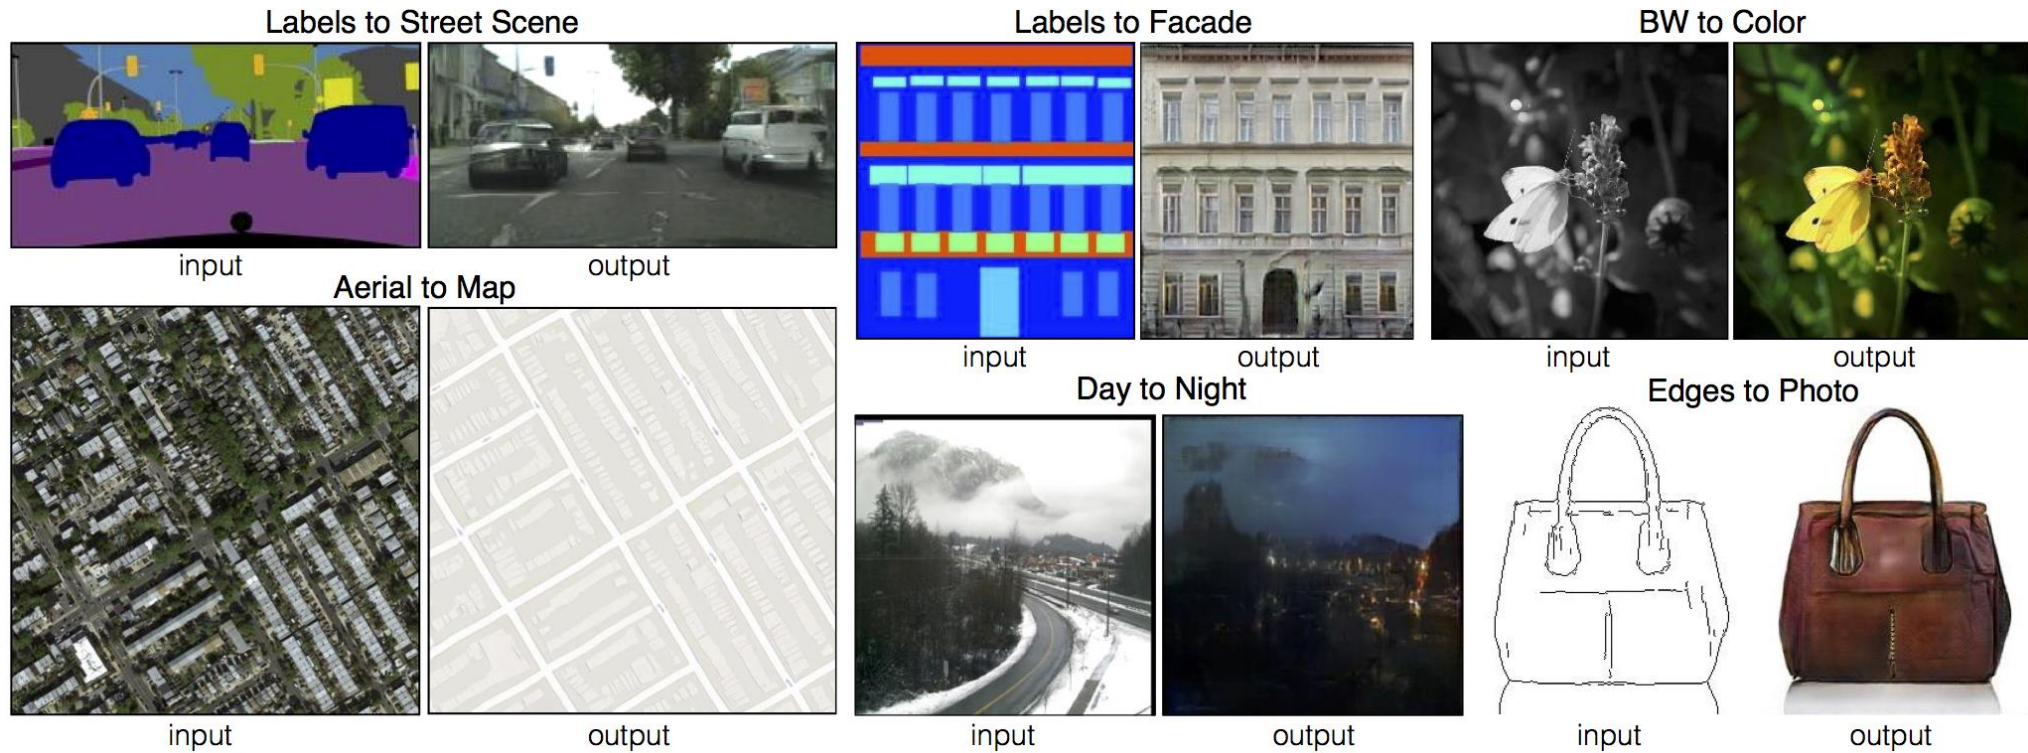
Learn a mapping from an observed side information \mathbf{x} and a random noise vector \mathbf{z} to the fooling samples \mathbf{y}

$$G: \{x, z\} \rightarrow y$$



Antipov et al, "Face Aging With Conditional Generative Adversarial Networks", ICIP 2017

Conditional Generation – Image2Image

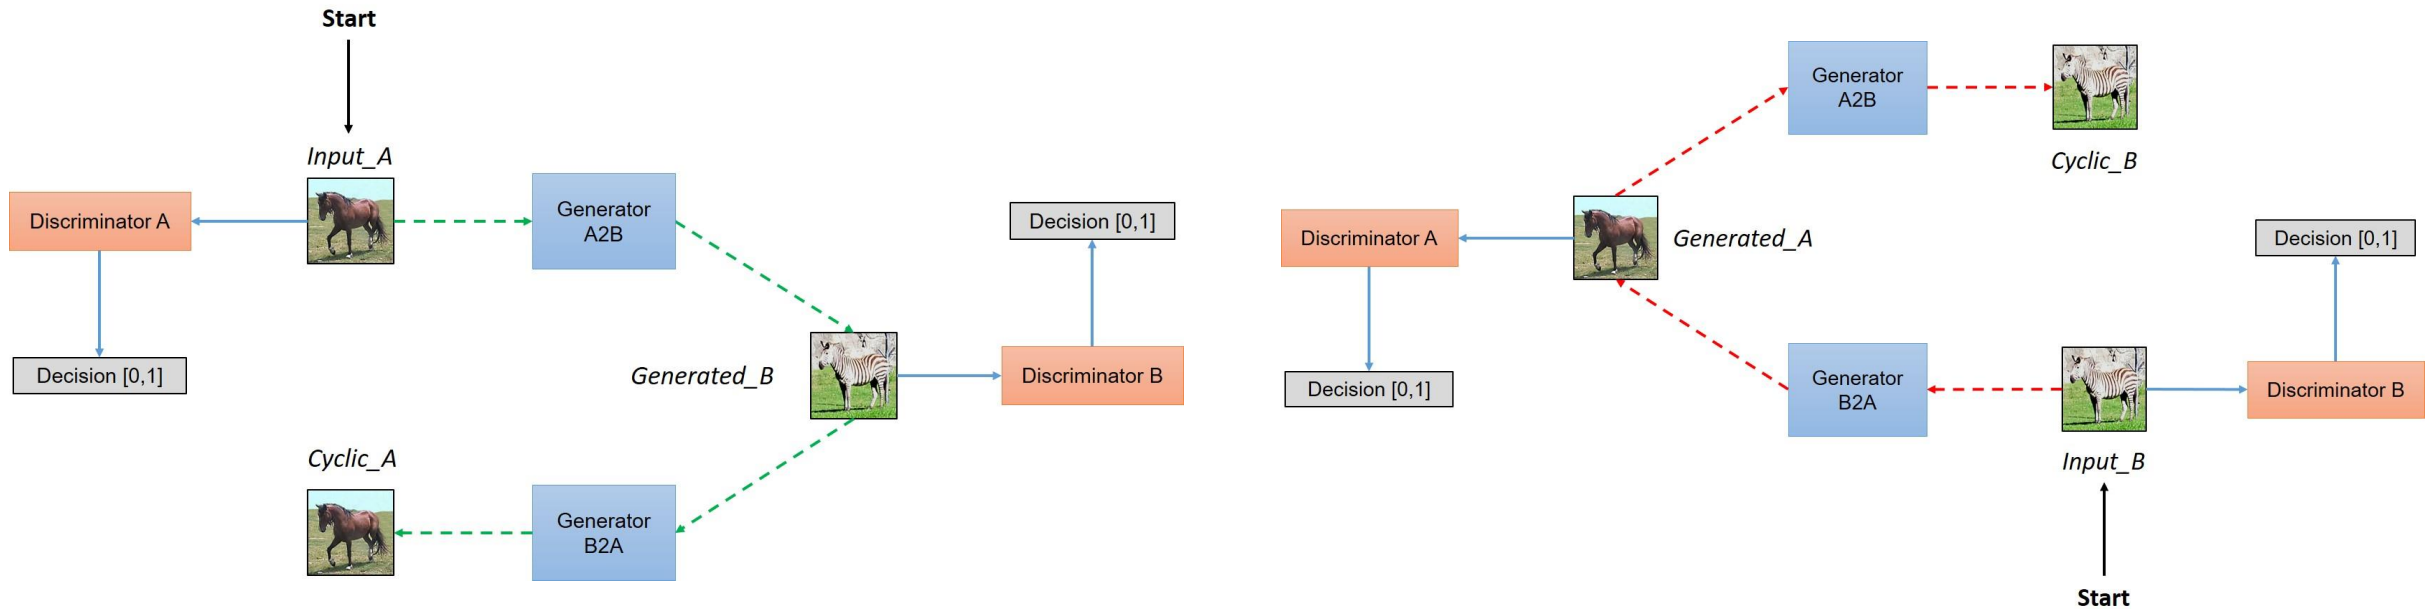


Isola et al, "Image-to-Image Translation with Conditional Adversarial Networks", 2016



UNIVERSITÀ DI PISA

CycleGAN – Style transfer without pairing



Enforce alignment by ensuring that generated images in domain B can lead to good fakes in domain A and vice versa

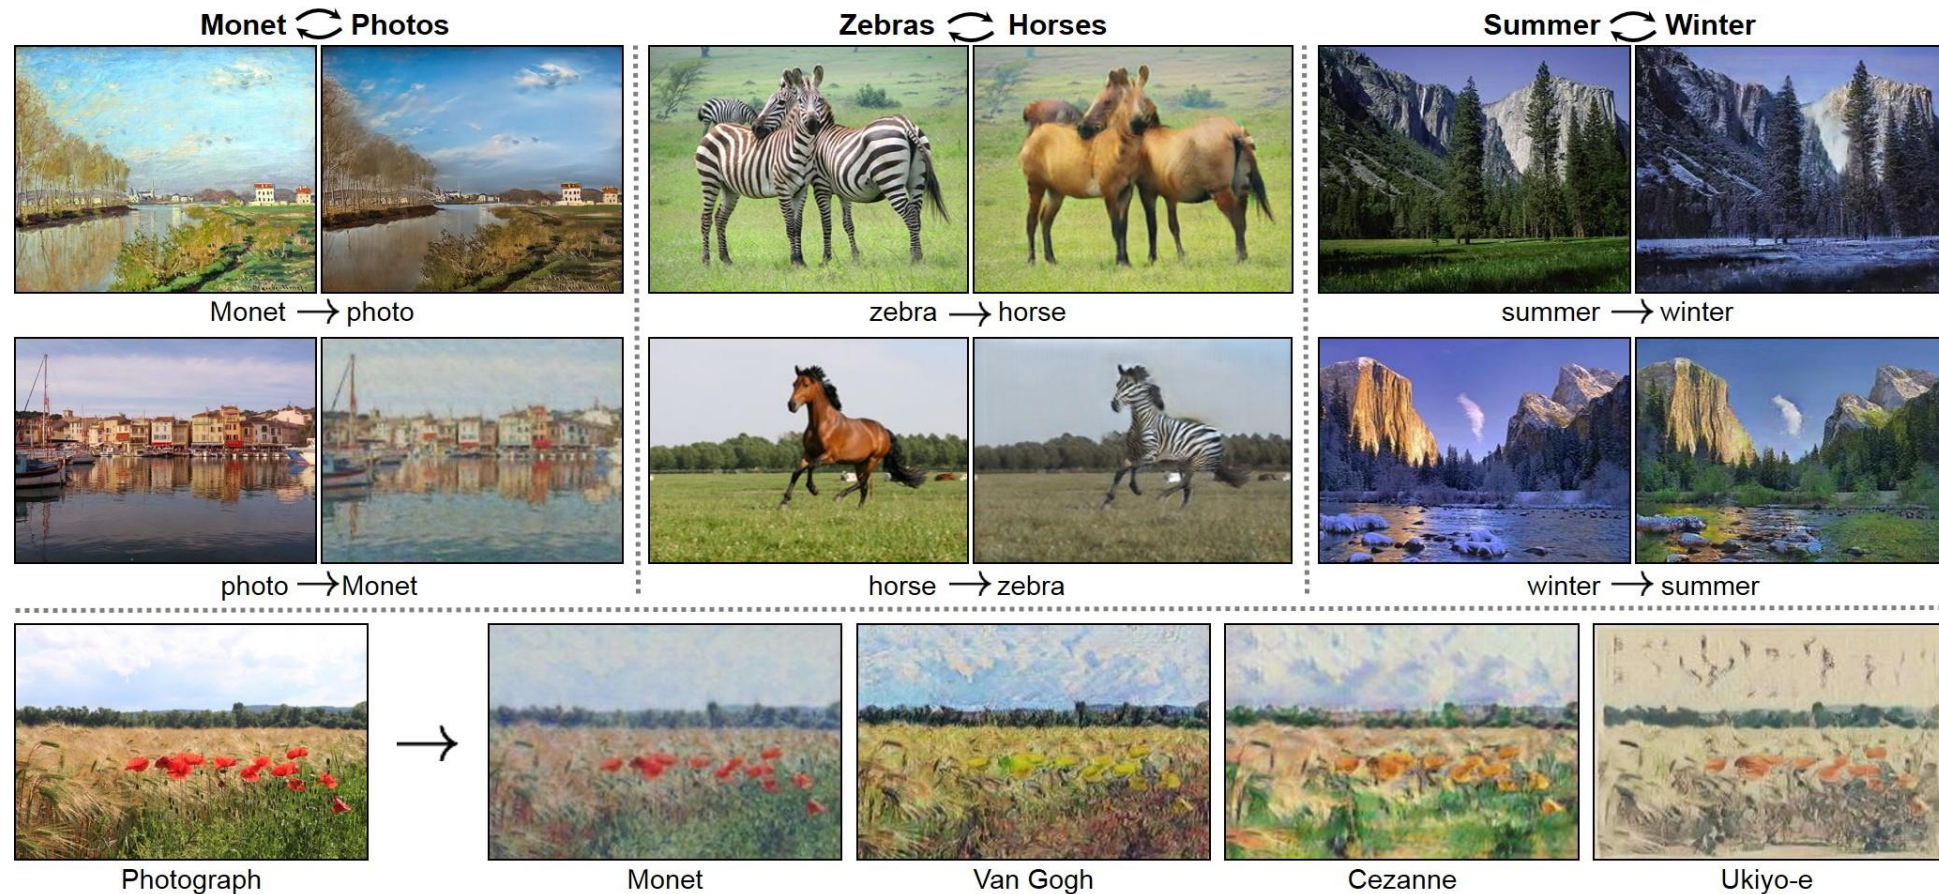
<https://junyanz.github.io/CycleGAN/>



UNIVERSITÀ DI PISA

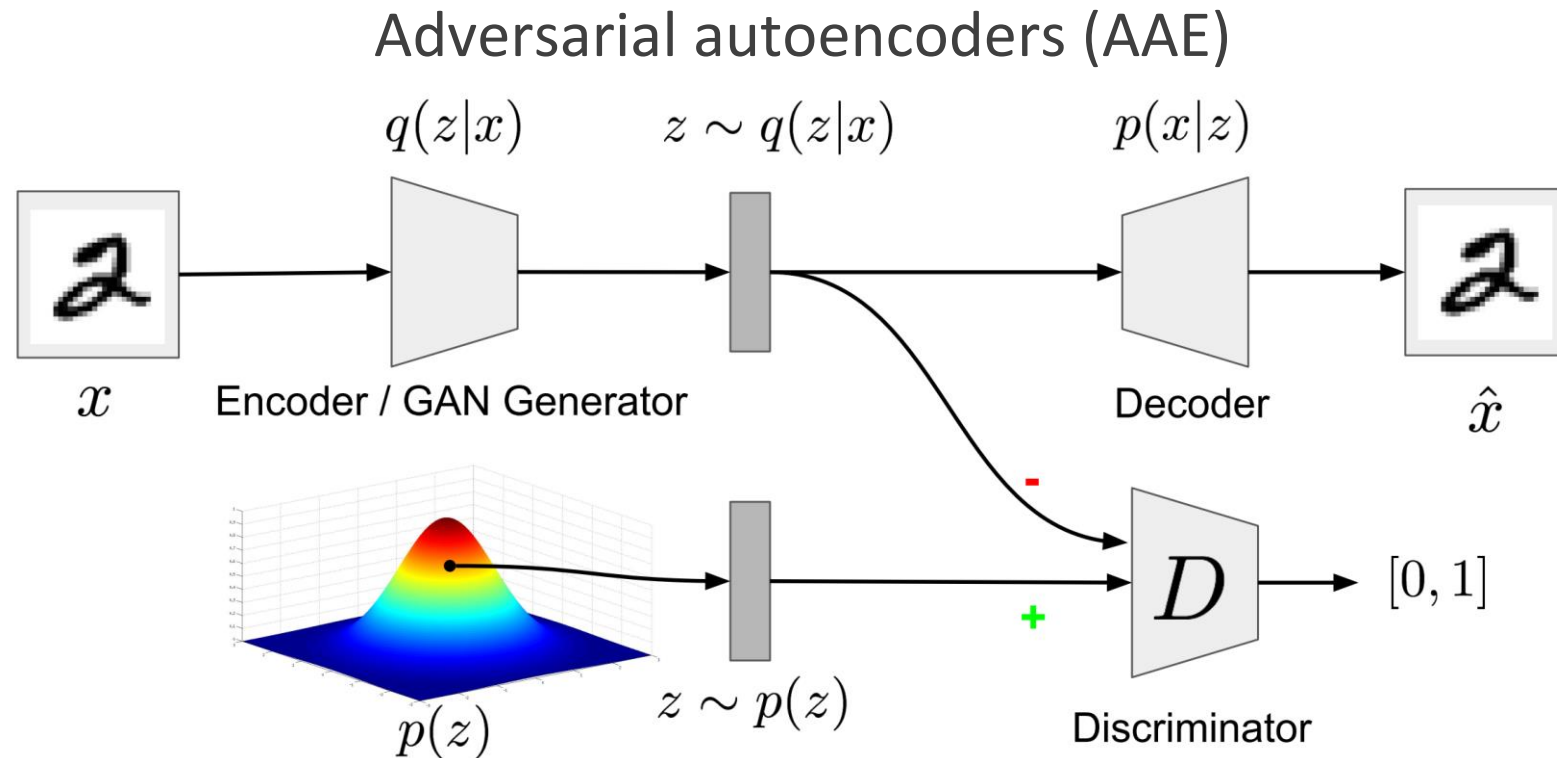
CycleGAN Examples

<https://junyanz.github.io/CycleGAN/>



UNIVERSITÀ DI PISA

Best of 2 worlds?



Force the latent codes to be **indistinguishable from samples of a priori distribution**



Training AAE

$$\mathcal{L}(x) = \mathbb{E}_Q[\log P(x|z)] - \underbrace{KL(Q(z|x) || P(z))}$$

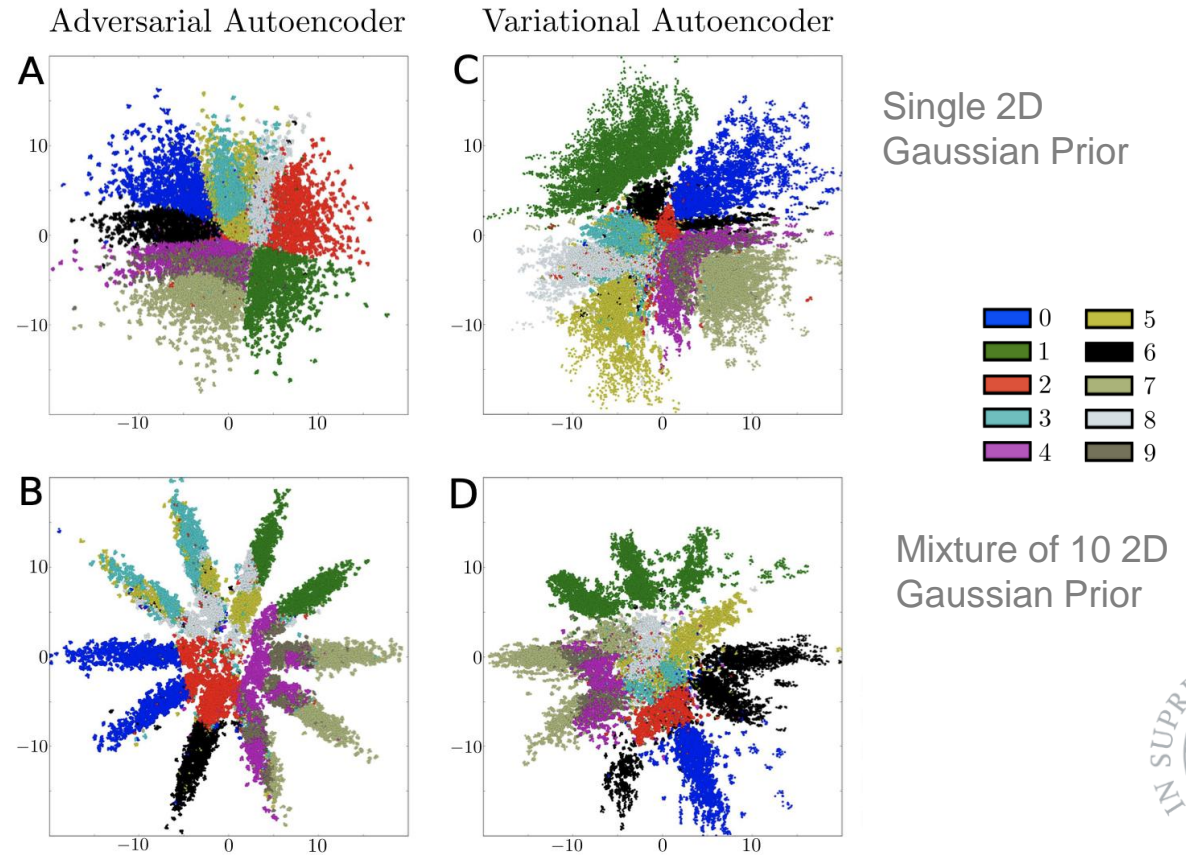
Replaced by an adversarial loss

- **Reconstruction phase** - Update the encoder and decoder to minimize reconstruction error
- **Regularization phase** - Update discriminator to distinguish true prior samples from generated samples; update generator to fool the discriminator
- Adversarial regularization allows to impose priors for which we cannot compute the KL divergence

AAE Vs VAE

AAE yields a smoother coverage of the latent space

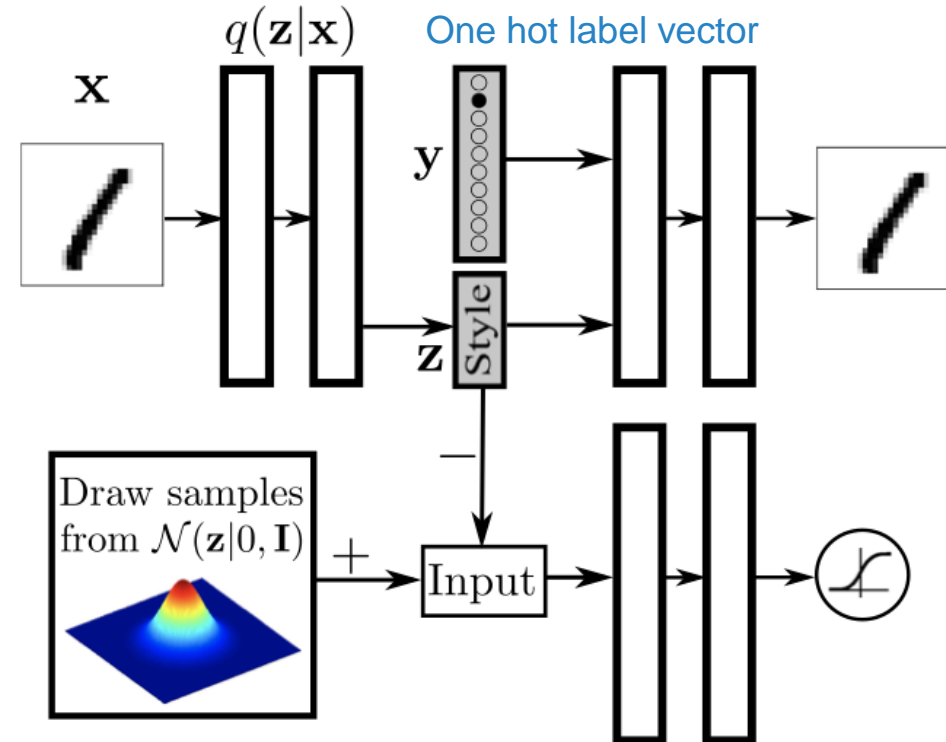

<https://arxiv.org/abs/1511.05644>



UNIVERSITÀ DI PISA


AAE – Style transfer (supervised)

Incorporate label information explicitly to force \mathbf{z} to capture class-independent information (e.g. style)



<https://arxiv.org/abs/1511.05644>

AAE – Semi-supervised learning



Factorize latent code in

- One hot encoding vector \mathbf{y}
- Continuous code \mathbf{z}

Distribution of \mathbf{y} made little distinguishable from a multinomial (induced from data)

<https://arxiv.org/abs/1511.05644>



UNIVERSITÀ DI PISA

Software

- A [TF implementation](#) of PixelCNN
- A list of acknowledged VAE implementations is kept by Kingma [here](#)
- Plenty of DCGAN implementations
 - [Torch](#)
 - [Tensorflow](#)
- Conditional GAN for image-to-image
 - [Pytorch](#) code
 - [Demo](#)
- So many GANs: check out the [GAN-Zoo](#)

Take Home Messages

- GAN – **Learn to sample** rather than learn the distribution
 - Sample quality only recently surpassed by diffusion models (tomorrow)
 - Unstable/difficult to train
 - Cannot perform inference (no distribution learning)
 - Needs differentiable generator
- Adversarial Autoencoders
 - Leveraging **adversarial penalties** in place of KL regularization
 - Useful to impose “**complex**” or **empirical priors**



Next Lecture

Back to explicit approaches

- Diffusion models
- Generative learning as a noising-denoising process
- Guided diffusion
- Conditional diffusion
- Latent space diffusion

