Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. **Conditional GANs**

Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. **Conditional GANs**
    1. **Principle**

Motivation

c: train

Prior distribution z

Generator

Text

Generator

Image

this white and yellow flower have thin white petals and a round yellow stamen

Generator

Image

Generator

Image

Generator

Conditional GAN

c^1: a dog is running $\hat{x}^1$: 

c^2: a bird is flying $\hat{x}^2$: 

- ***Text to image*** by traditional supervised learning

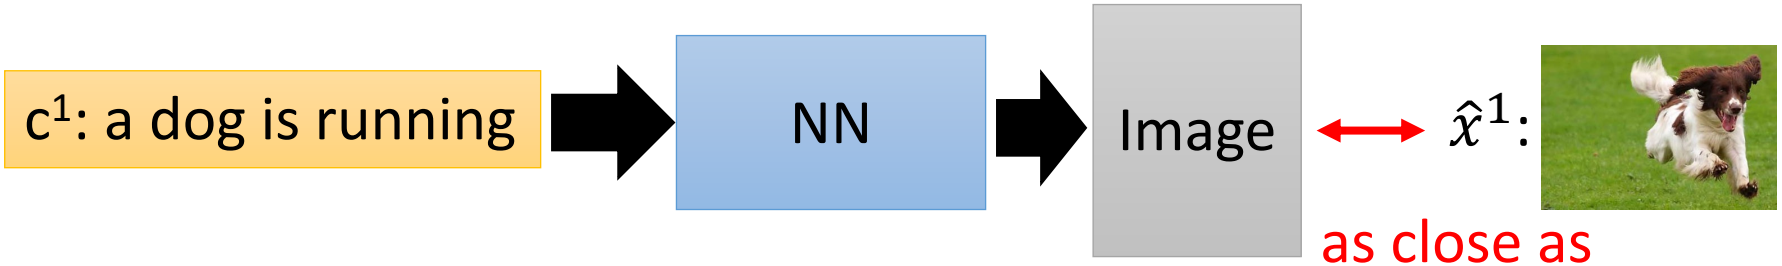

Text: "train"

Target of NN output

A blurry image!

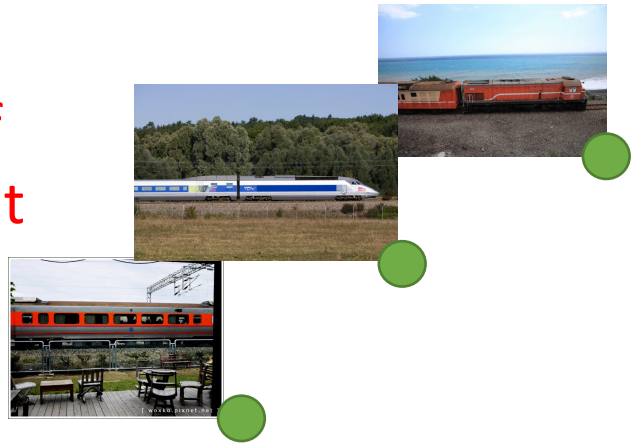

Conditional GAN

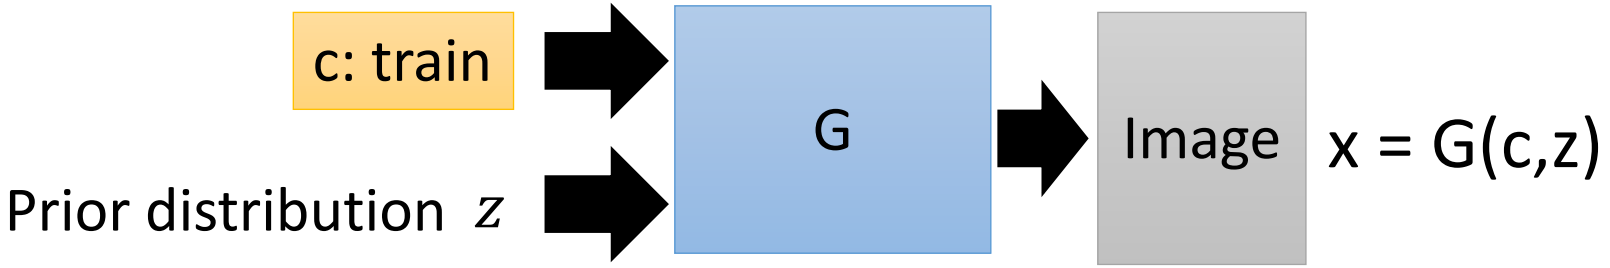

It is a distribution

Approximate the distribution of real data

Text: “train”

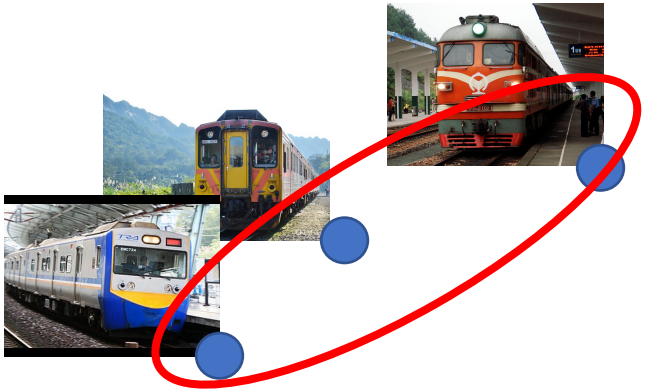

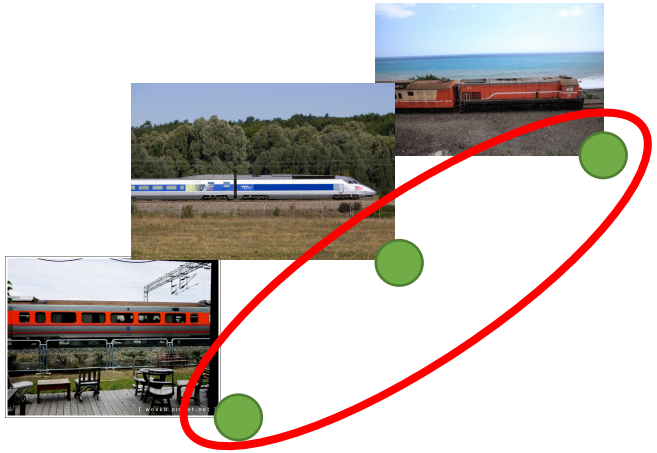

Conditional GAN

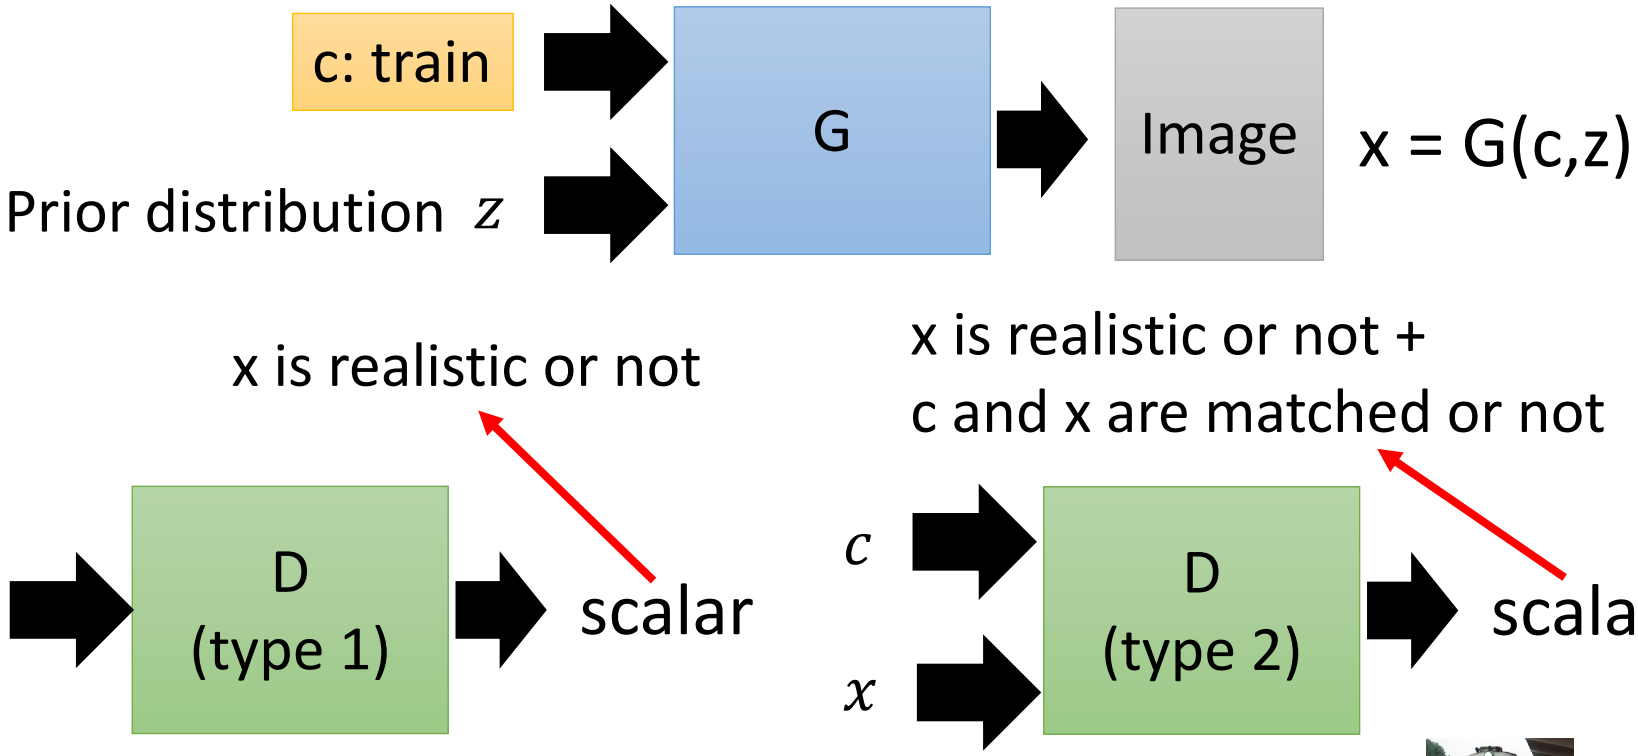

Conditional GAN (cGAN model)

GAN

$$V(G, D) = \mathbb{E}_{x \sim P_{data}}[logD(x)] + \mathbb{E}_{x \sim P_G}\left[log\left(1 - D(x)\right)\right]$$

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

cGAN

$$\min_G \max_D \left(\mathbb{E}_{\mathbf{x},\mathbf{y} \sim p_{\text{data}}(\mathbf{x},\mathbf{y})} [\log D(\mathbf{x}, \mathbf{y})] + \mathbb{E}_{\mathbf{y} \sim p_{\mathbf{y}},\, \mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z}, \mathbf{y}), \mathbf{y}))]\right)$$

Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. Conditional GANs
    1. Principle
    2. **Text2Image**

Text2Image: architecture example

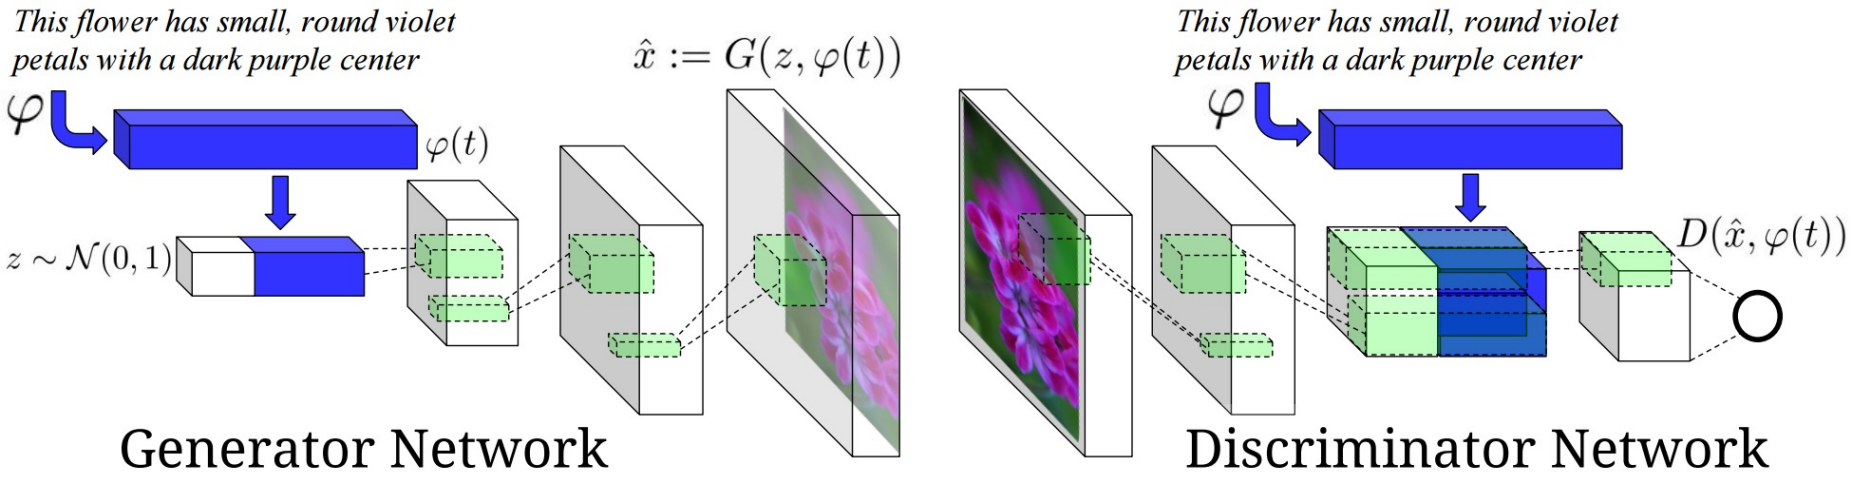

- Positive samples:
  - real image + right texts
- Negative samples:
  - fake image + right texts
  - Real image + wrong texts

Text2Image results

this small bird has a pink breast and crown, and black primaries and secondaries.

this magnificent fellow is almost all black with a red crest, and white cheek patch.

the flower has petals that are bright pinkish purple with white stigma

this white and yellow flower have thin white petals and a round yellow stamen

Text2Image results

Caption	Image
this flower has white petals and a yellow stamen	
the center is yellow surrounded by wavy dark purple petals	
this flower has lots of small round pink petals	

Text2Image: architecture example (2)

Generating higher resolution images (from 64 to 256)

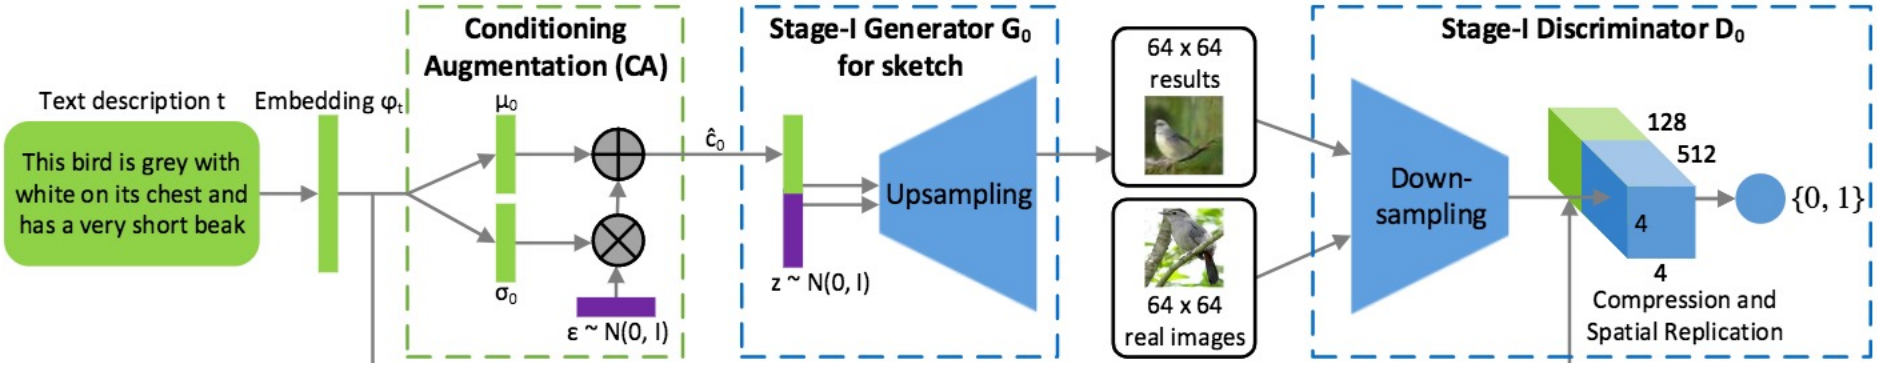

Text2Image: architecture example (2)

Generating higher resolution images (from 64 to 256)

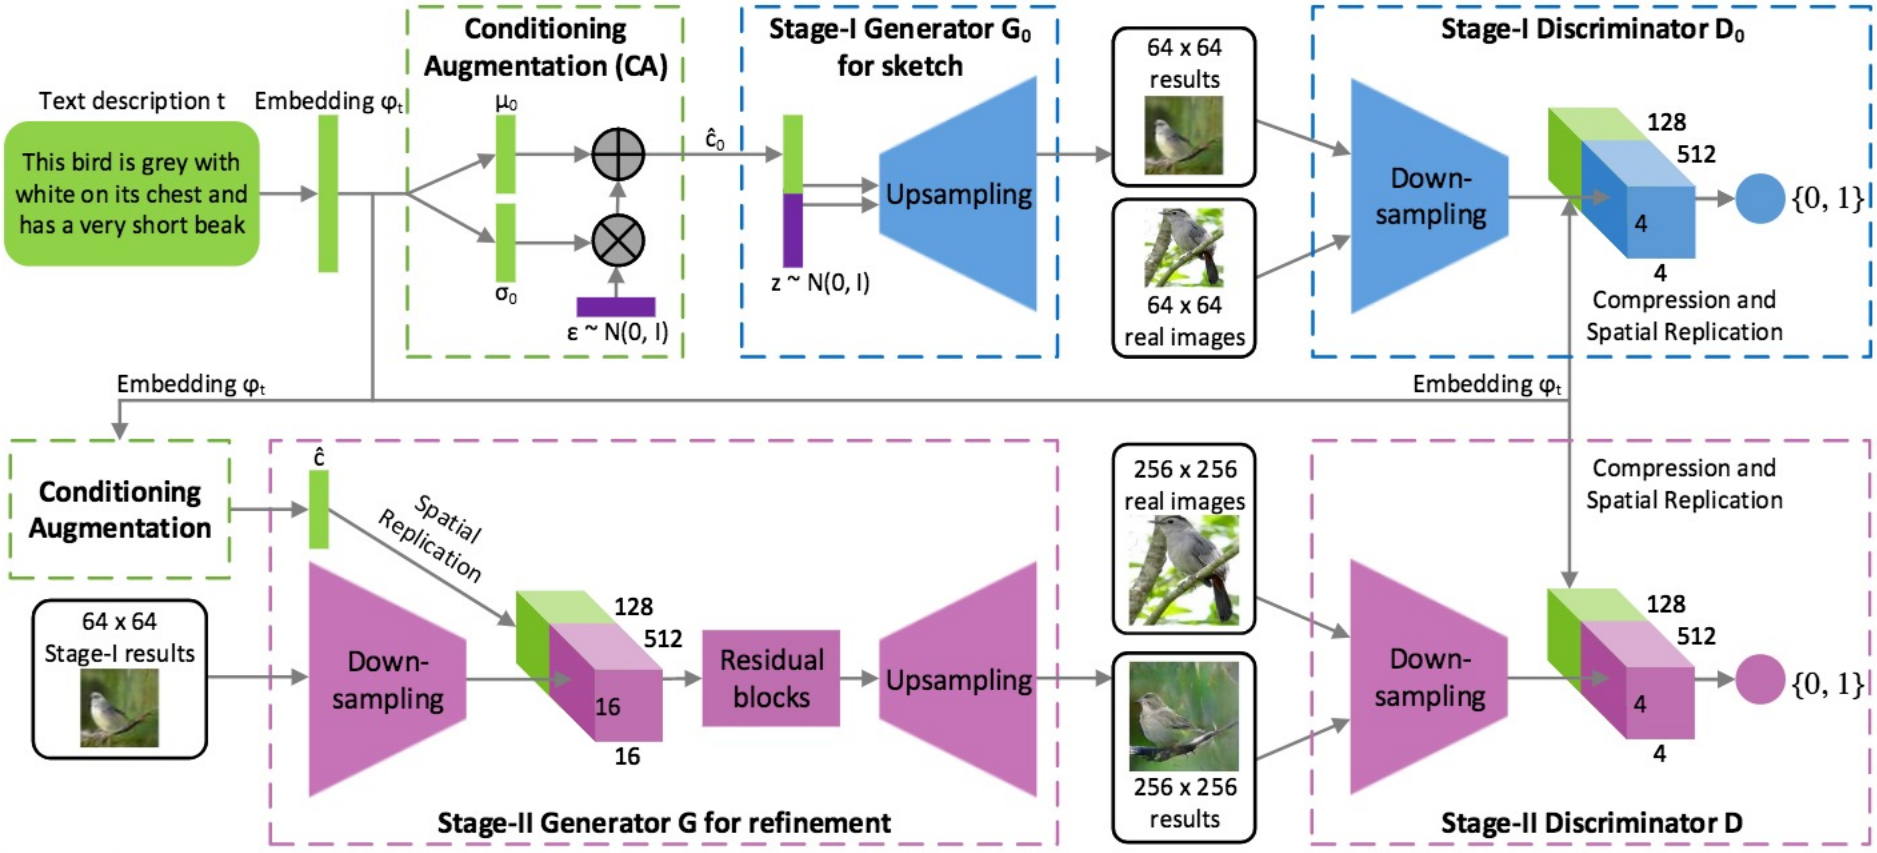

StackGAN results

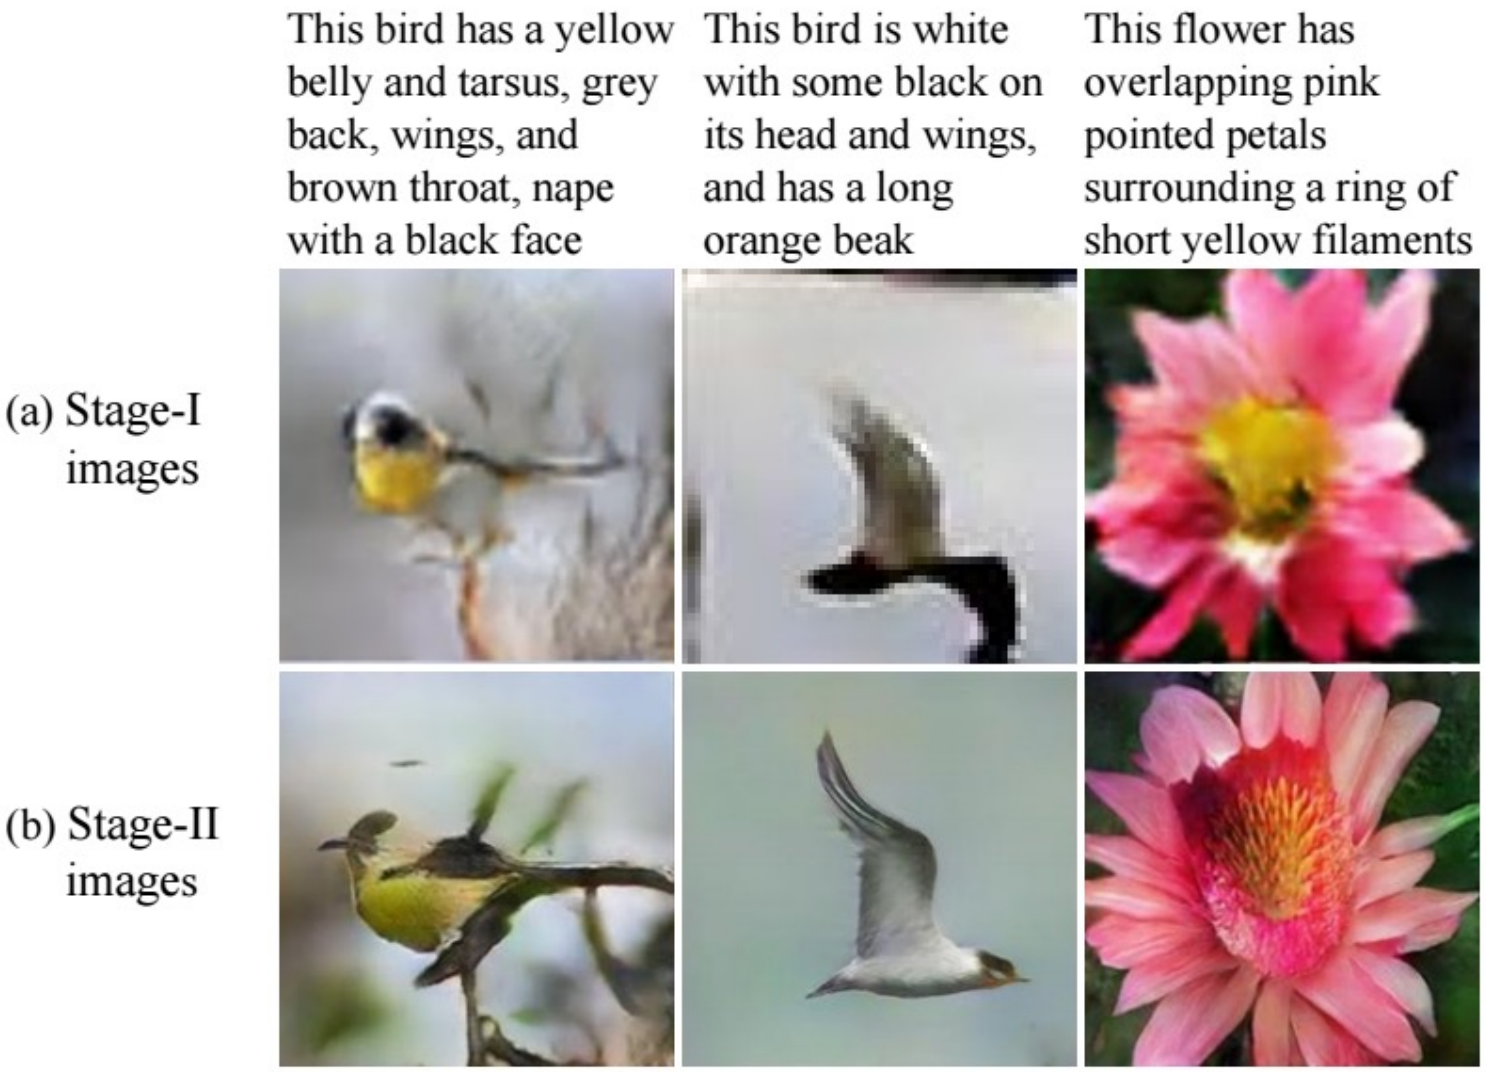

[Zhang et al. 2016]

Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. Conditional GANs
    1. Principle
    2. Text2Image
    3. **Image2Image**

Image-based Conditional GAN

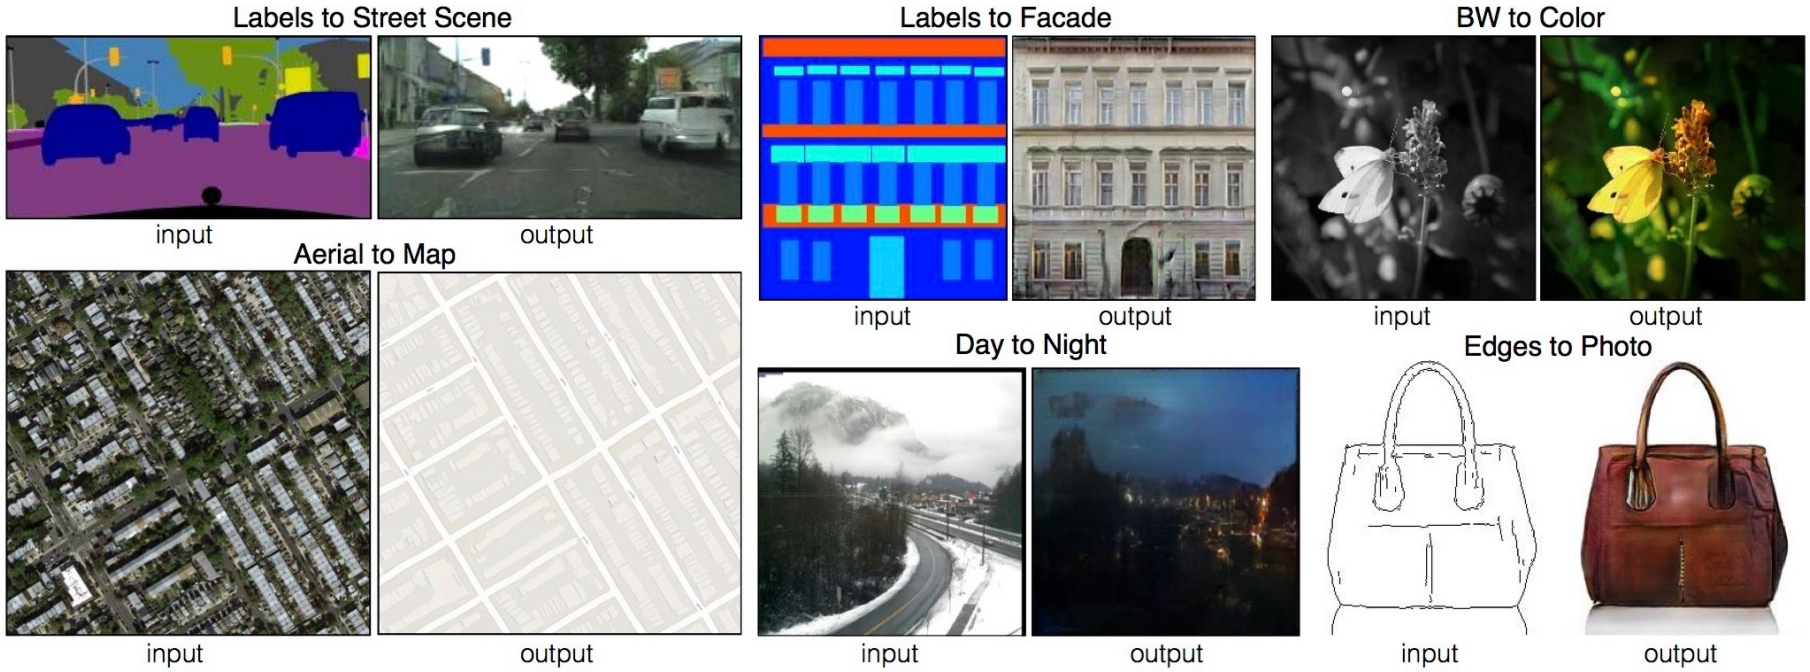

- Conditioned on an image of different modality
- Image-to-Image Translation => pix2pix

Image-to-image pix2pix

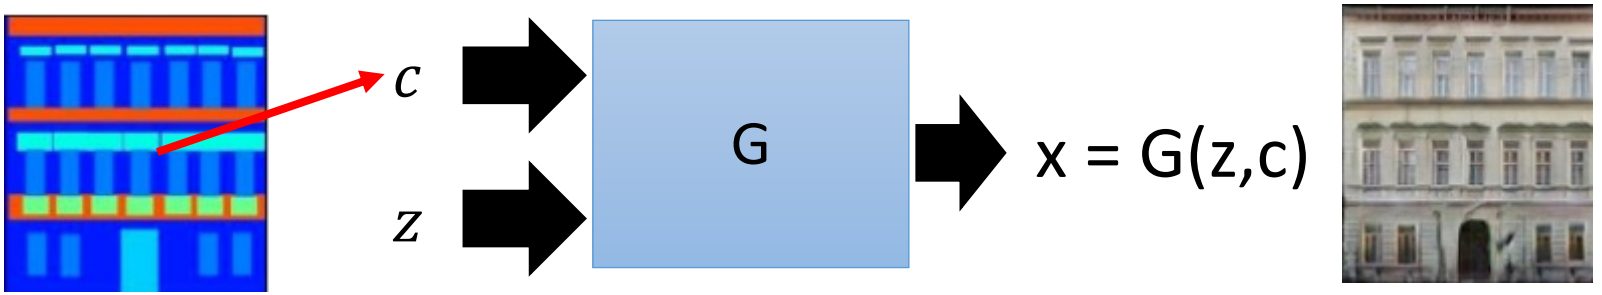

Image-to-image pix2pix

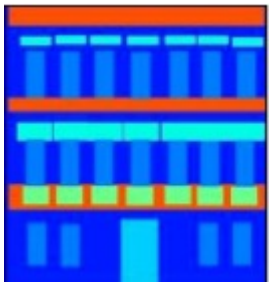


- Traditional supervised approach

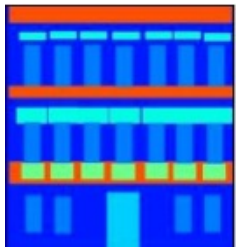

Testing:

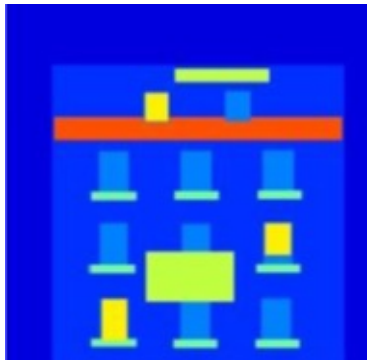
input

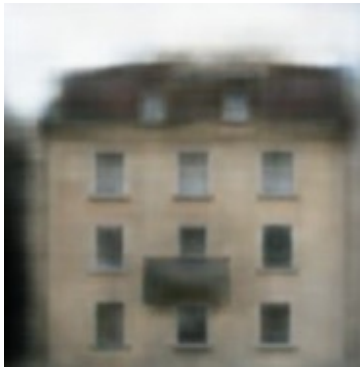
close

It is blurry because it is the average of several images.

Image-to-image

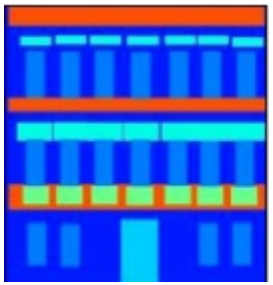

- Conditional GAN

Supplementary loss (close)

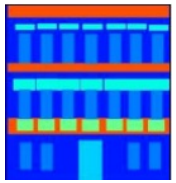
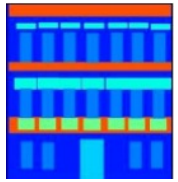

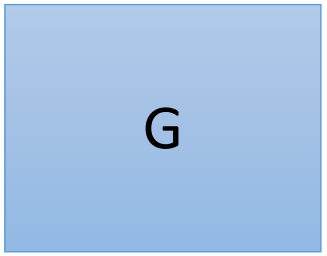

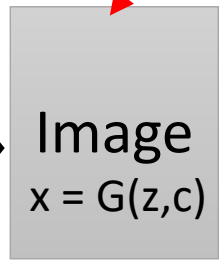

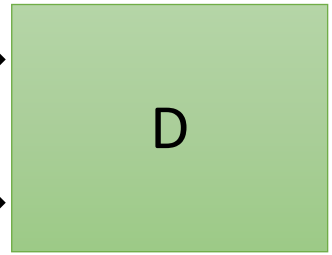

scalar

z

Testing:

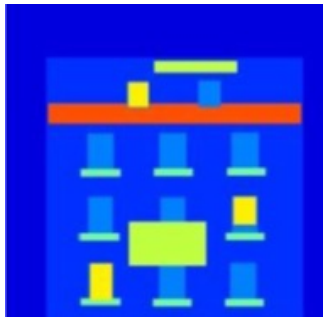
input

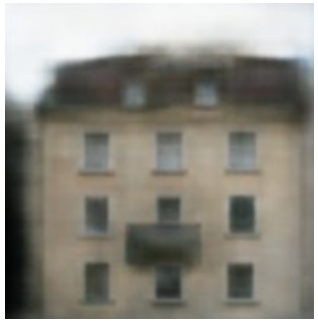
close


GAN


GAN + close


GT

Positive examples

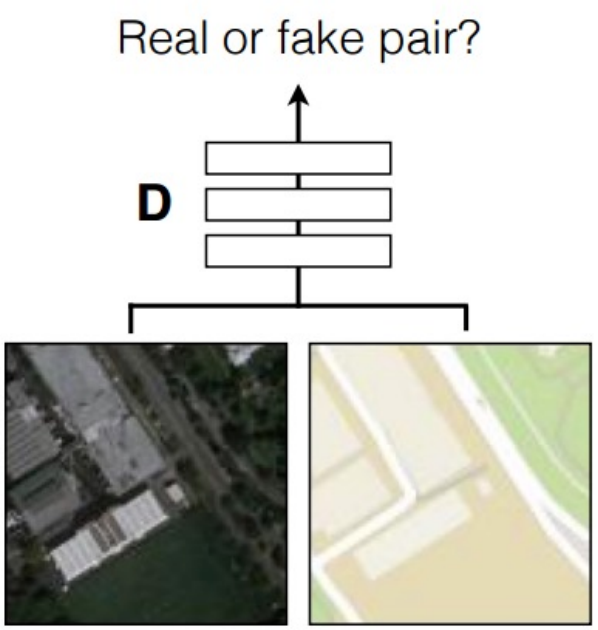

Negative examples

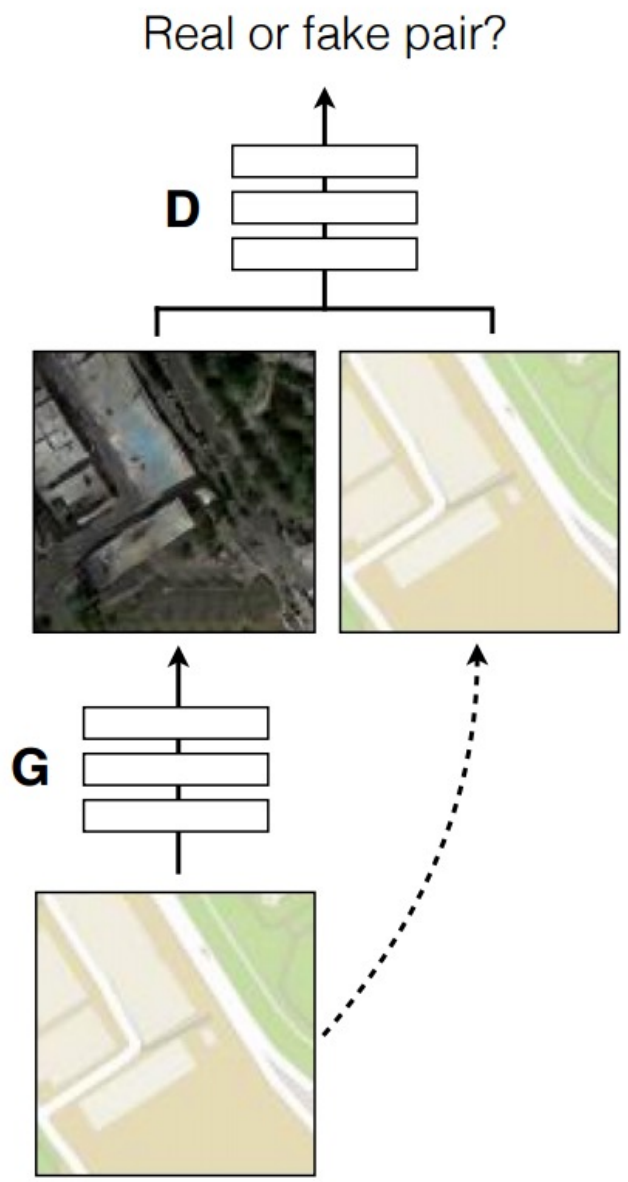

G tries to synthesize fake images that fool **D**

D tries to identify the fakes

Label2Image

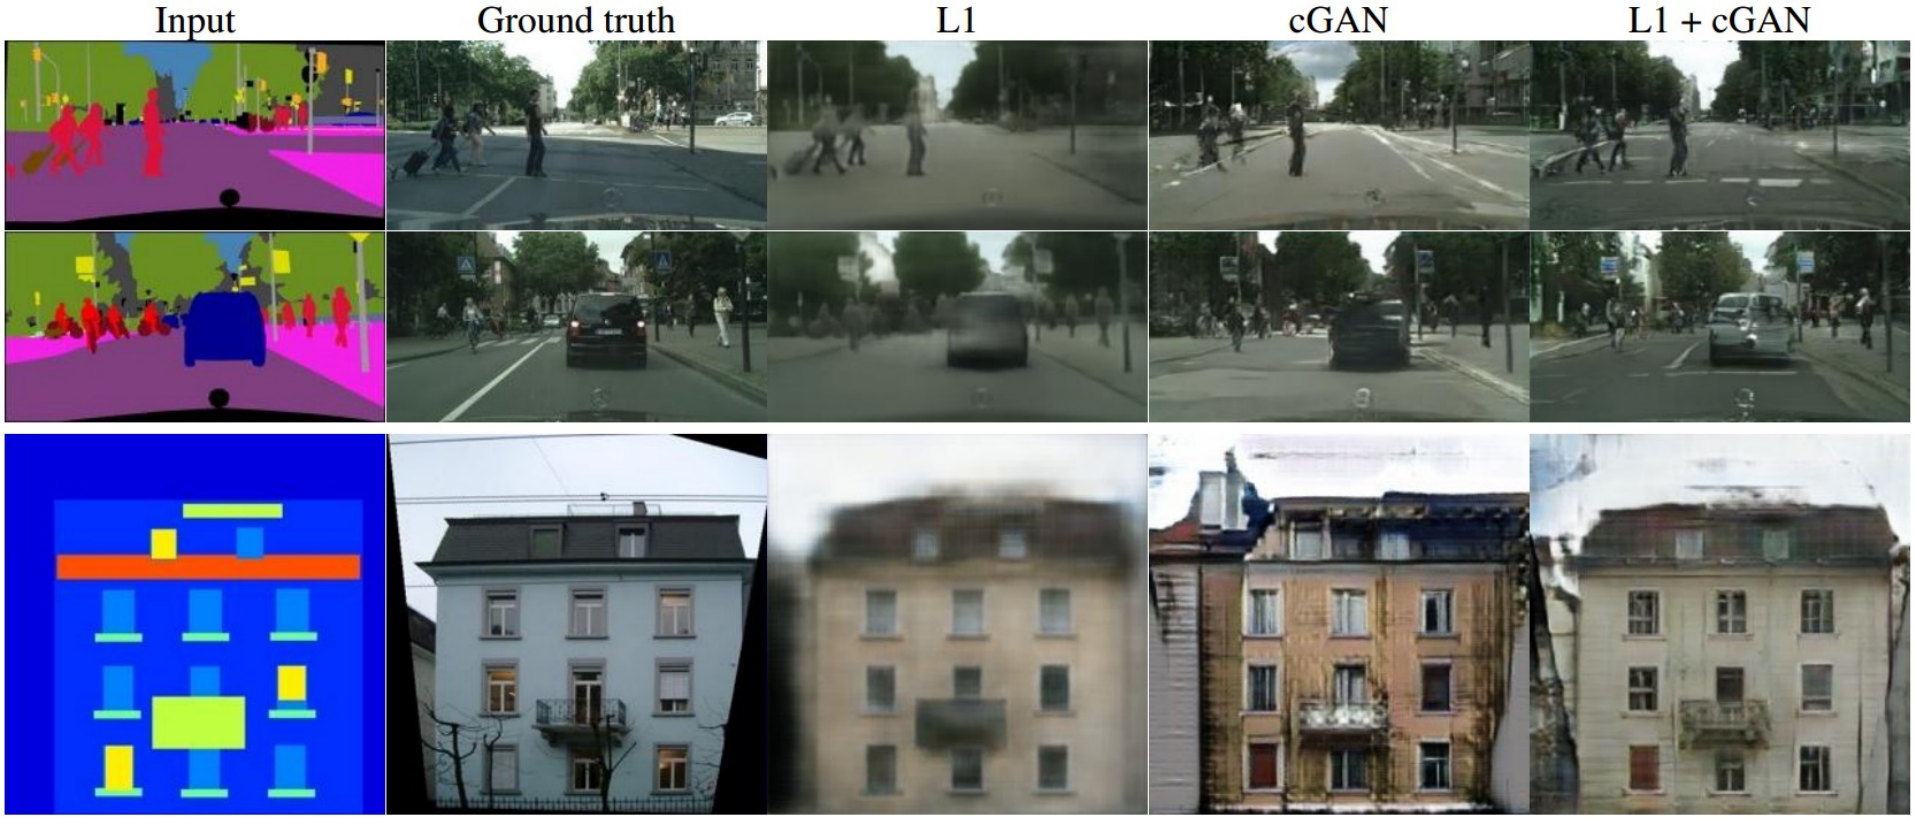

Edges2Image

Pix2pixHD

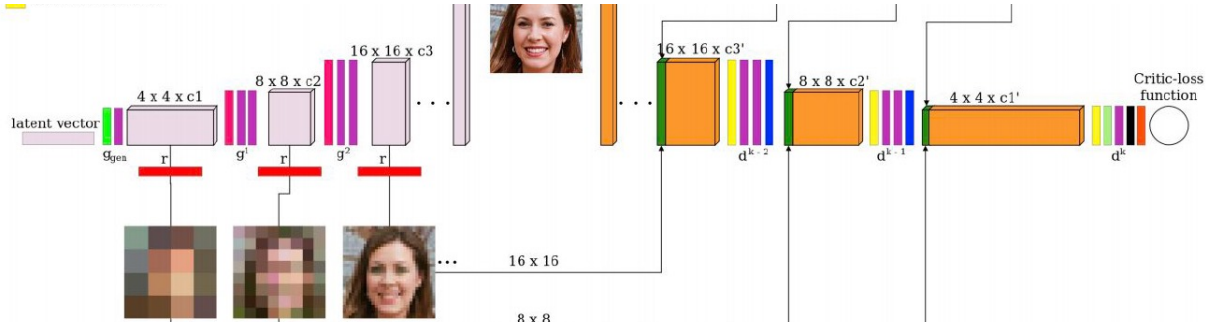

Coarse-to-fine Generator

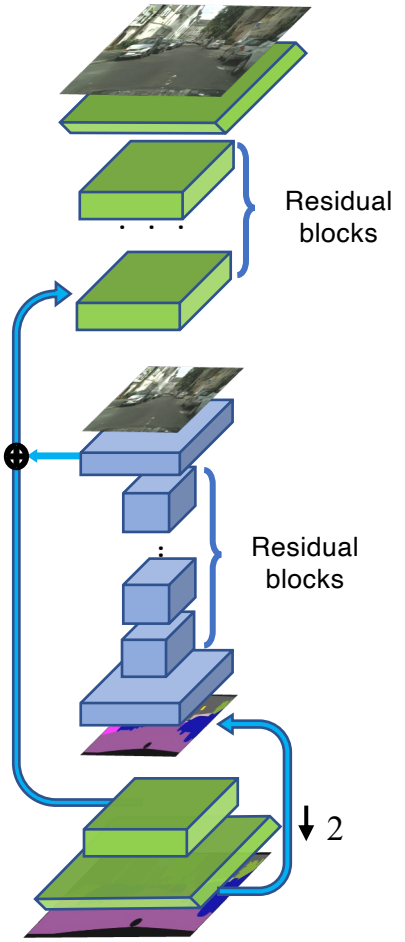

Multi-scale Discriminators

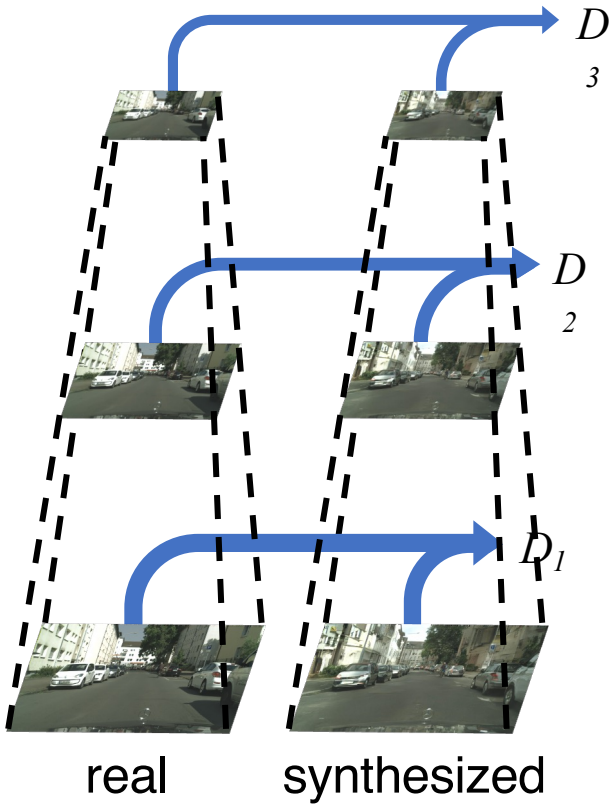

Robust Objective

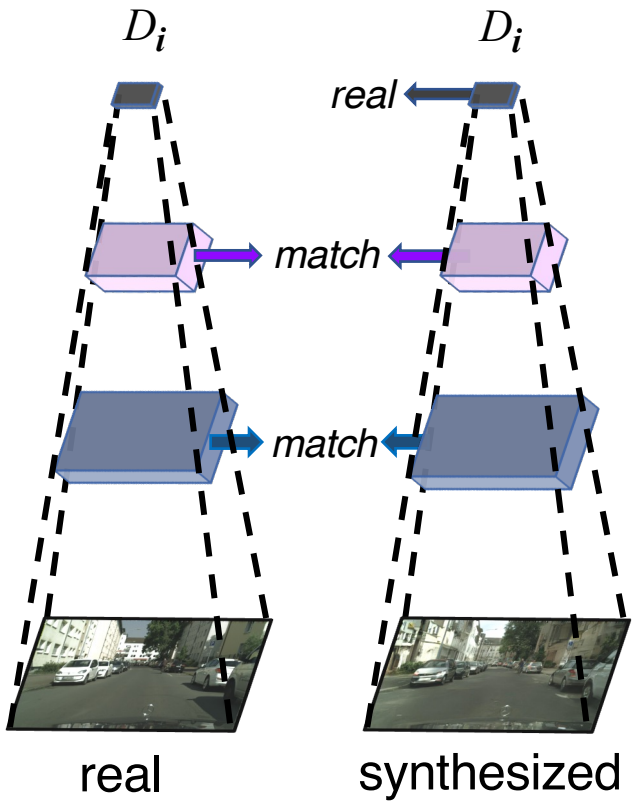

Semantic Map

pix2pix

CRN

Ours

Improving Segmentation2Image strategy?

Limitation of the approach:

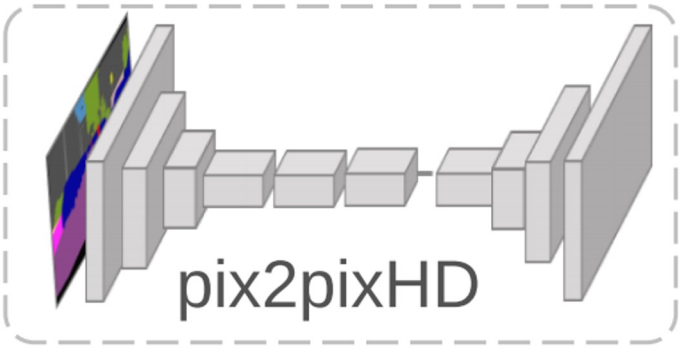

Directly feed the semantic layout as input to the deep network, which is processed through stacks of convolution, normalization, and nonlinearity layers.

However, this is suboptimal as the normalization layers tend to “wash away” semantic information in input semantic segmentation masks.

Improving Segmentation2Image strategy

Proven effective for recent generative adversarial networks such as StyleGAN

Can we do the same for conditional GAN?
Conditional Normalization Layers?

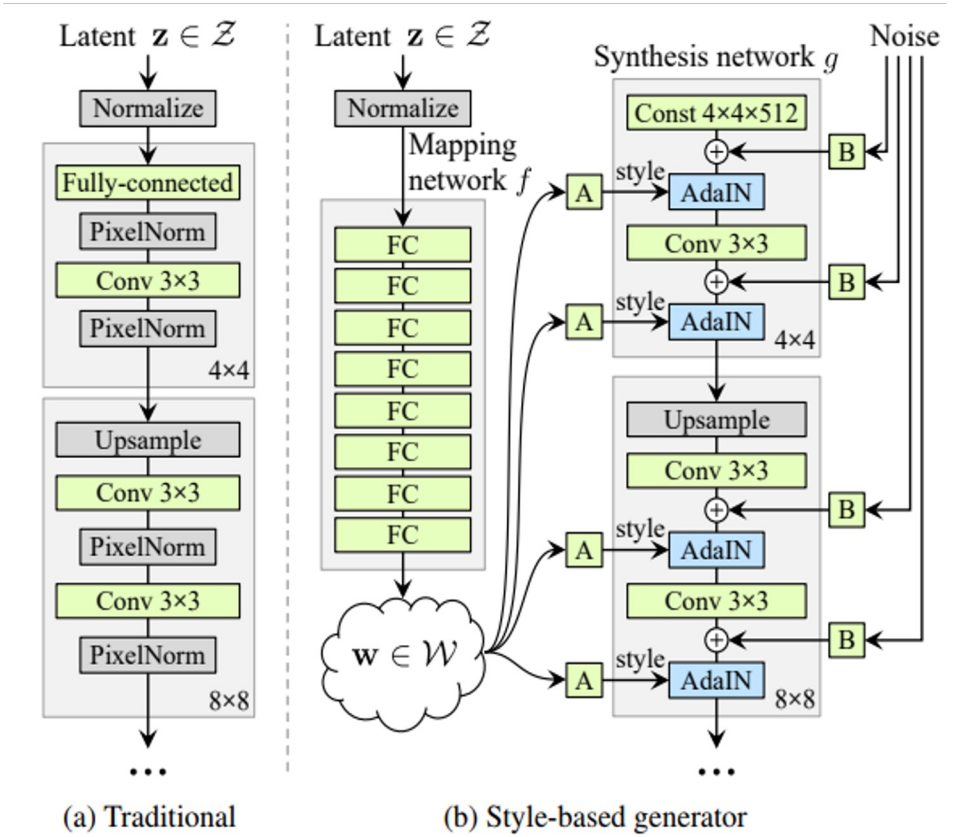

Improving Segmentation2Image strategy

Recall: Adaptive instance normalization

$$\text{AdaIN}(x, y) = \sigma(y)\left(\frac{x - \mu(x)}{\sigma(x)}\right) + \mu(y)$$

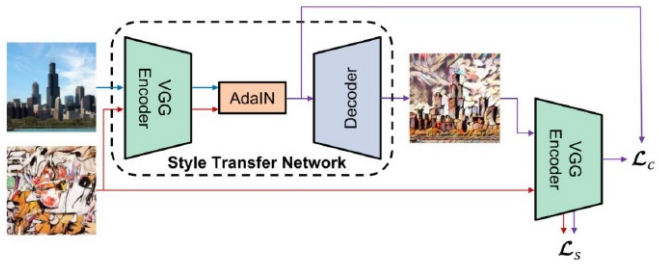

SPADE block= spatially-adaptive denormalization:

Same idea but per class c over each channel i (N=batch size)

$$\gamma^i_{c,y,x}(\mathbf{m})\frac{h^i_{n,c,y,x} - \mu^i_c}{\sigma^i_c} + \beta^i_{c,y,x}(\mathbf{m})$$

$$\mu^i_c = \frac{1}{NH^iW^i}\sum_{n,y,x} h^i_{n,c,y,x}$$

$$\sigma^i_c = \sqrt{\frac{1}{NH^iW^i}\sum_{n,y,x}(h^i_{n,c,y,x})^2 - (\mu^i_c)^2}$$

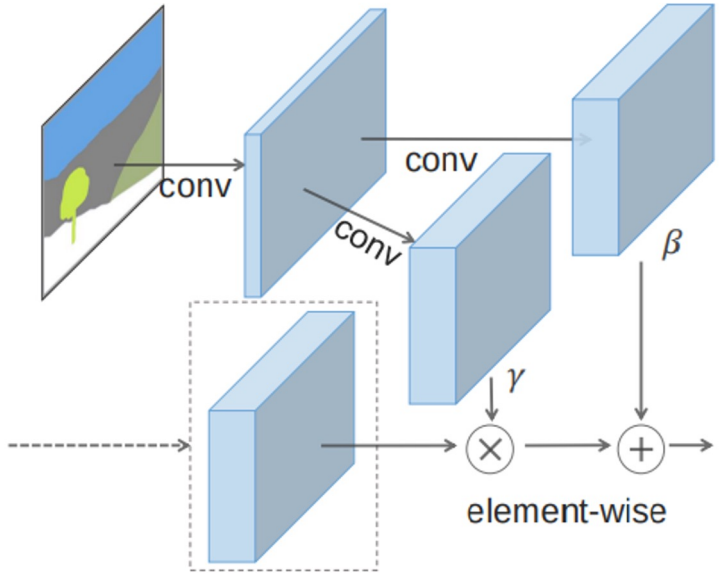

SPADE paper = [Semantic Image Synthesis with Spatially-Adaptive Normalization CVPR 2019]

SPADE Generator

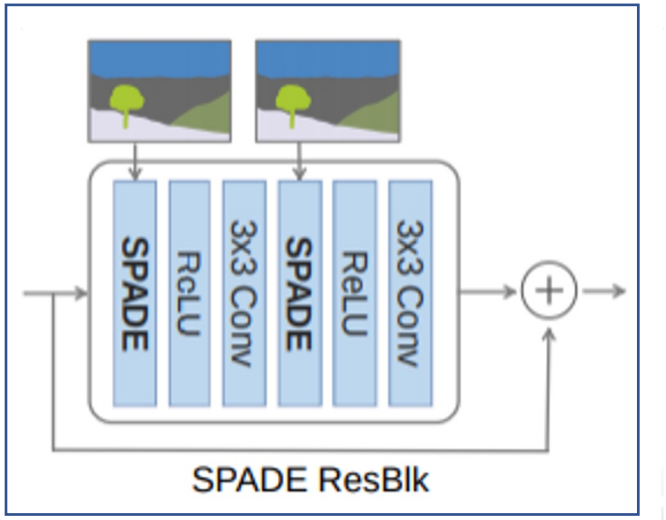

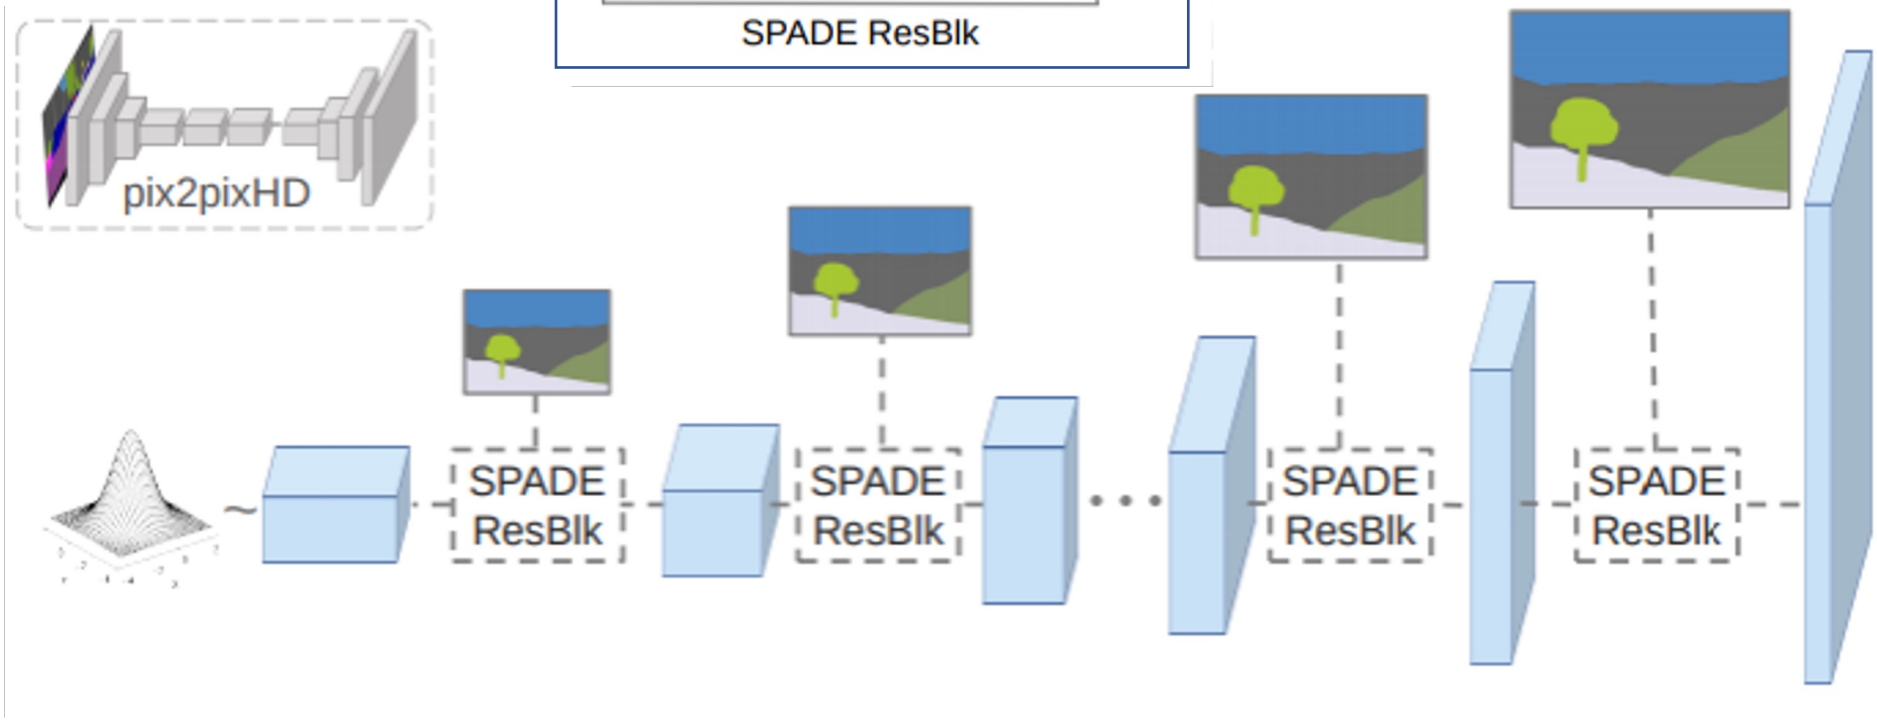

SPADE Generator

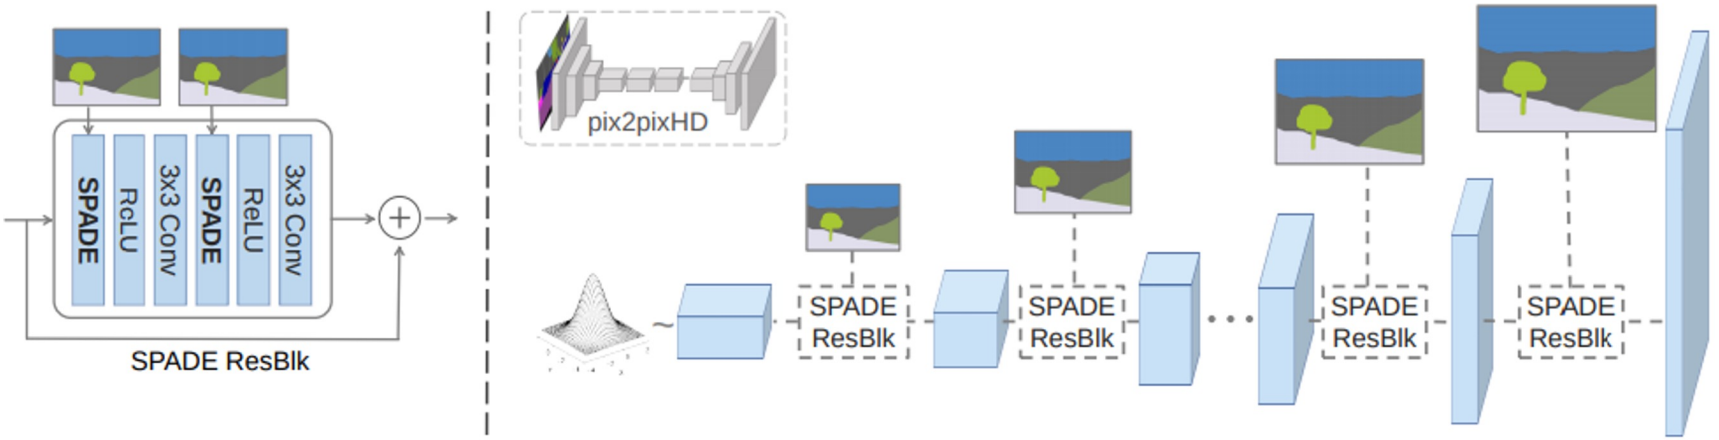

Better preserve semantic information against common normalization layers

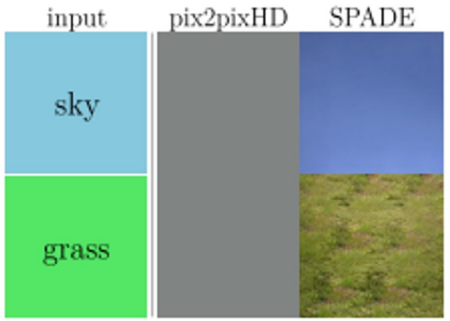

SPADE results

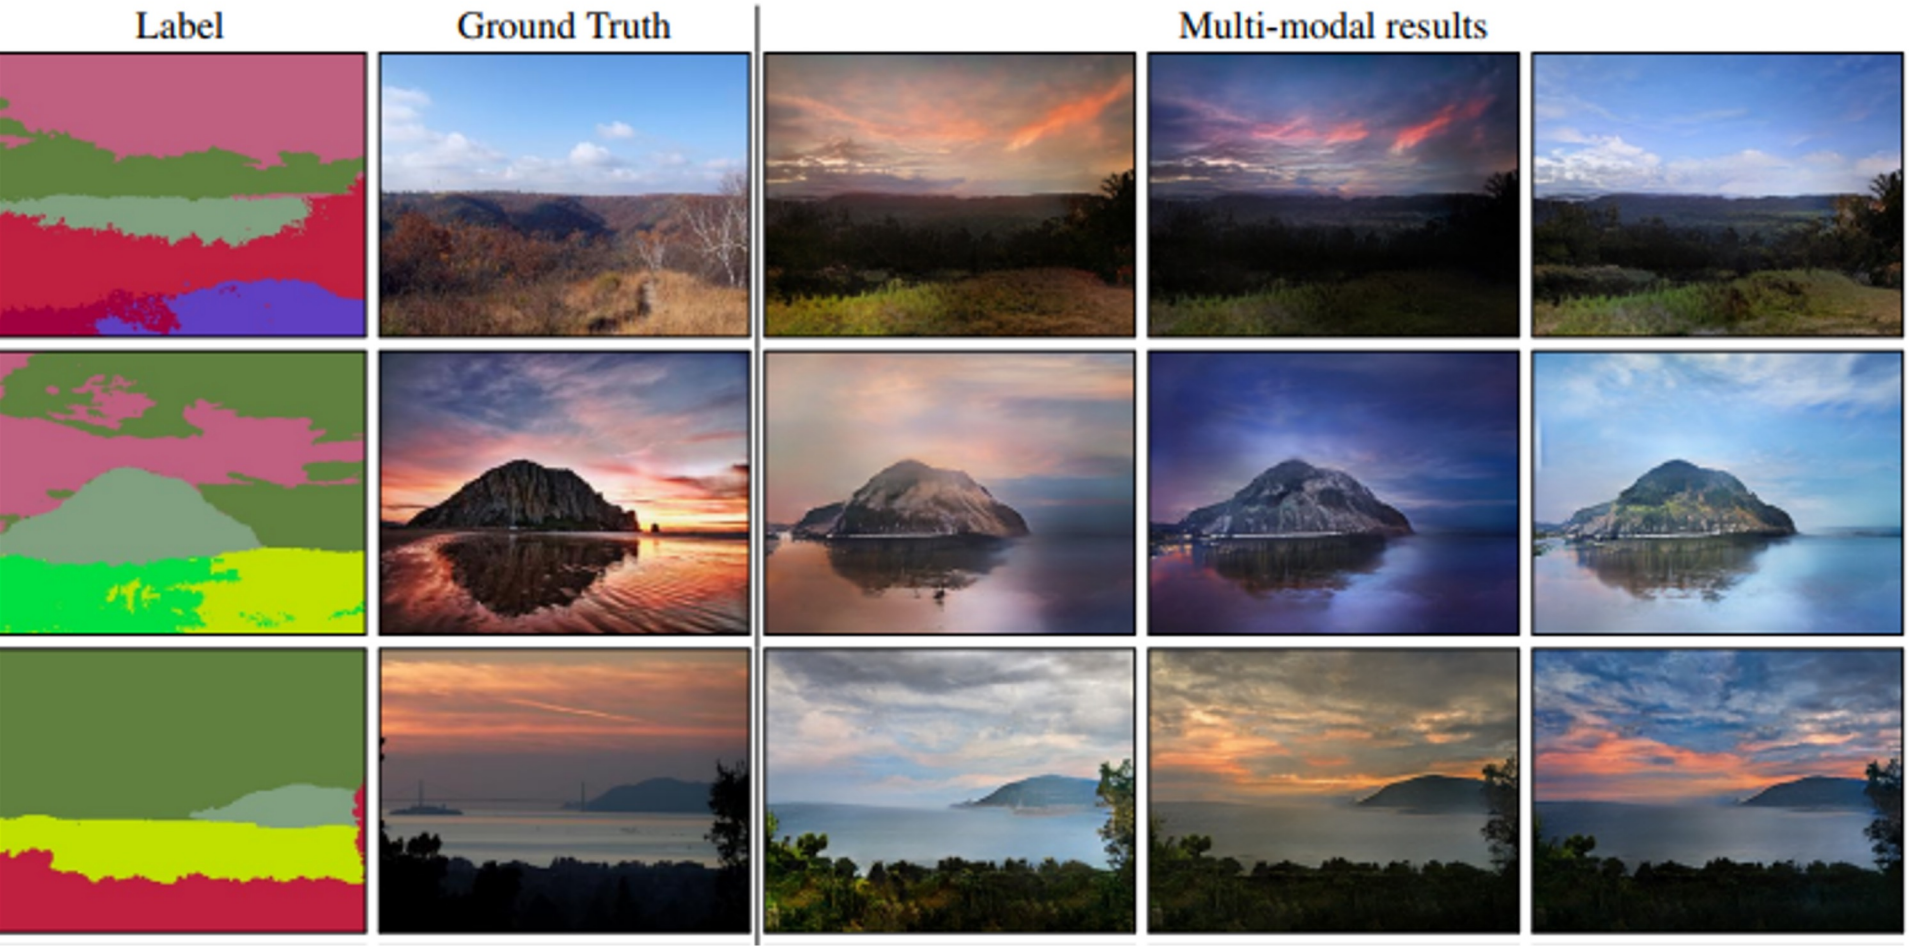

Spade and follow-up approaches

Editing with *conditional-or-structured-latent* GANs

BlobGan [Epstein *et al.* ECCV'22]

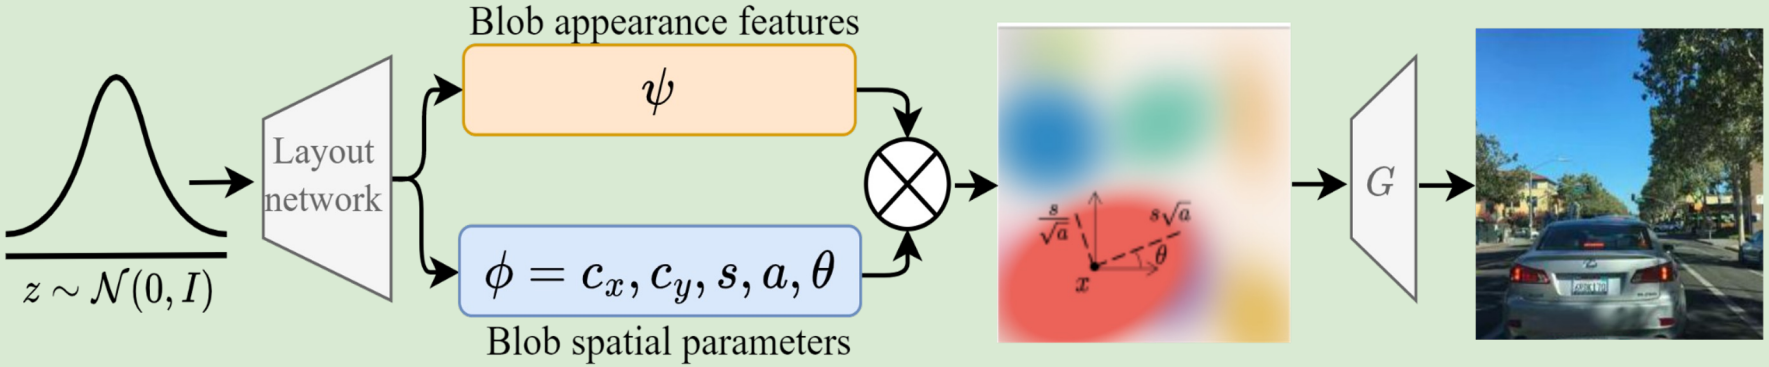

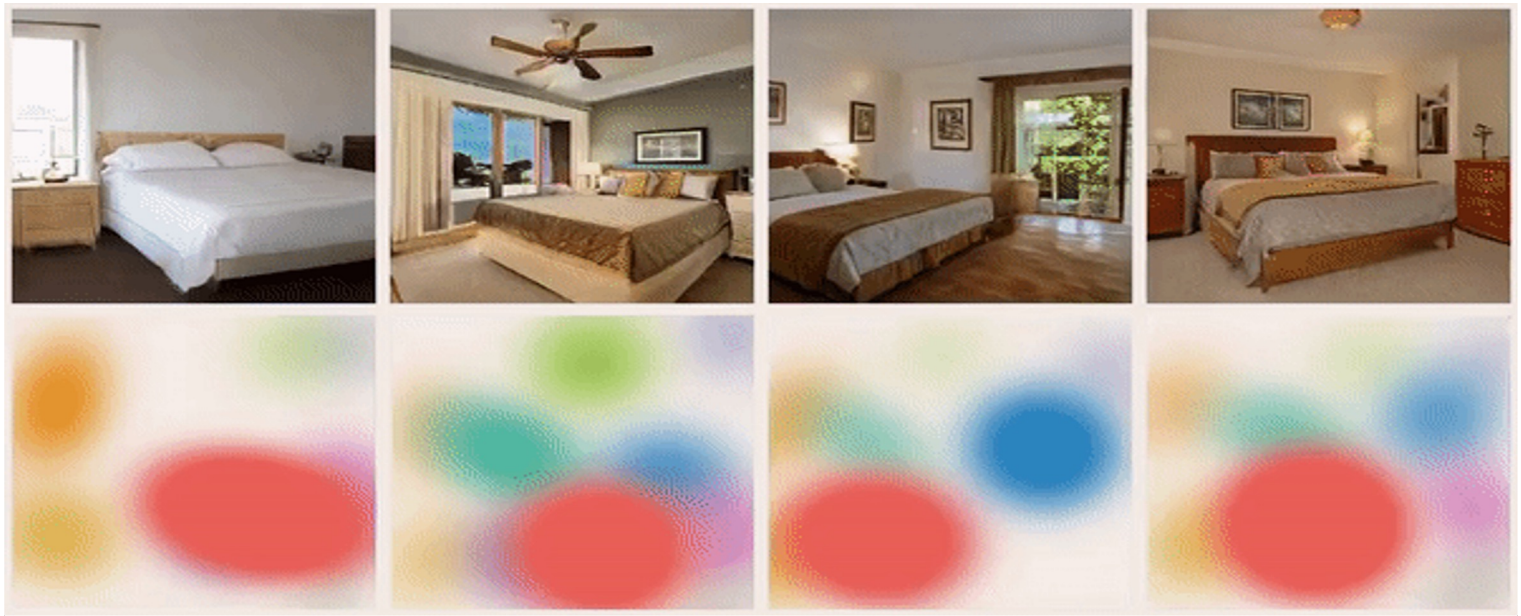

Editing with *conditional-or-structured-latent* GANs

Editing with *conditional-or-structured-latent* GANs

Example of Counterfactual optimization for editing

Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. Conditional GANs
    1. Principle
    2. Text2Image
    3. Image2Image
    4. **Inpainting and general missing data encoder**

Inpainting task

- Complete the missing part

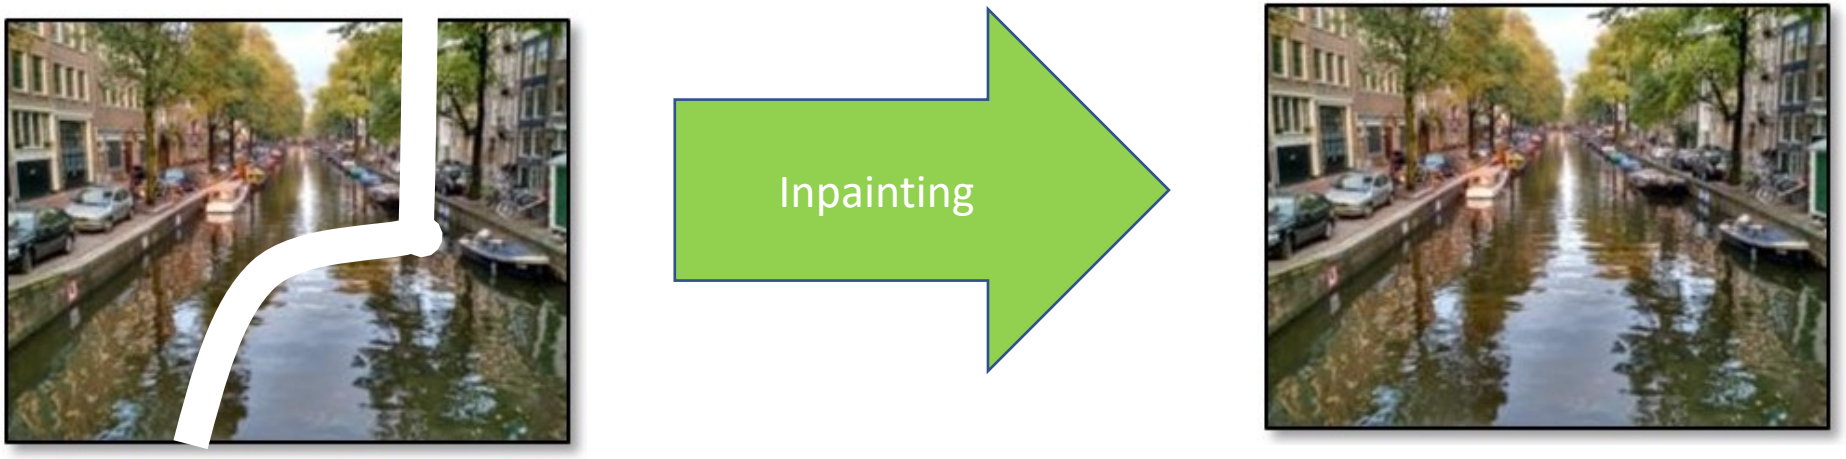

Inpainting as unsupervised learning with GAN loss

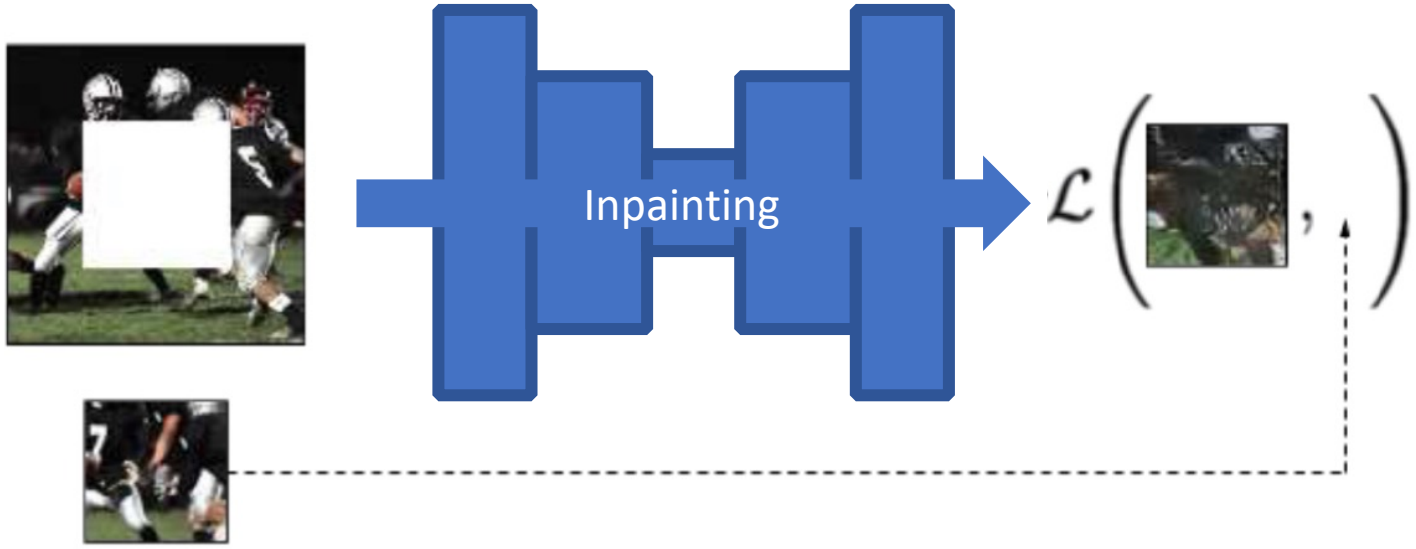

Reconstruct missing pixels by decoding using context

Loss defined on the predicted patch and the real one (known at training time)

First proposition -- Architecture

- Architecture: Encoder/Fully connected/Decoder

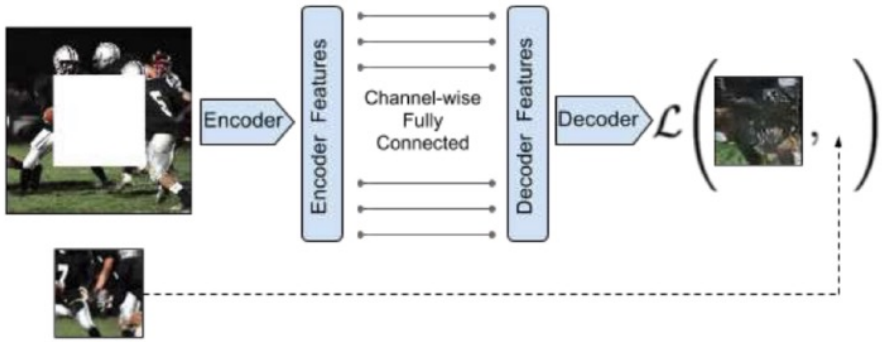

- DC-GAN for inpainting task
- **Input:** $227 \times 227 \times 3$ image
- **Output:** encoder context features ($6 \times 6 \times 256$)

Channel-wise fully-connected layer

- **Input / output:** $6 \times 6 \times 256$ channels
- **First layer:** Channel-wise fully-connected
  (each 6×6 input connected to the corresponding 6×6 output)
- **Second layer:** Stride 1 convolution to mix channels

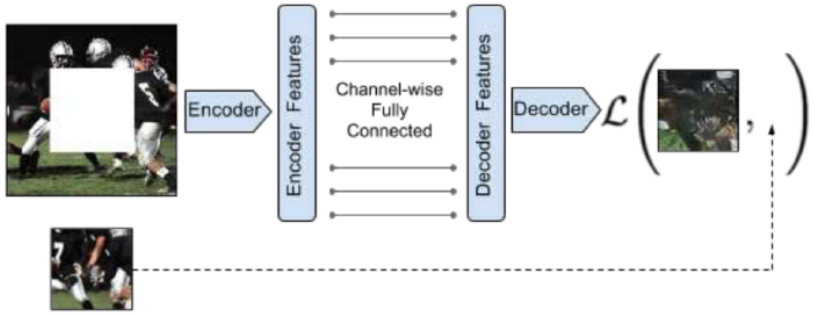

Decoder

- **Architecture:** Same as DC-GAN: 5 up-convolutional layers
  ("deconv" + ReLU)
- **Input:** decoder context features $6 \times 6 \times 256$
- **Output:** $227 \times 227 \times 3$ image

Training: Masking the images

- **How to define the mask ?**
  - Center region of the image
  - Random regions (chosen solution)
  - Random segmentation mask from VOC (said to be equivalent to random regions)
- **Formal definition:** Defined by a mask $\hat{M} \in \{0,1\}^{227 \times 227}$ with 1 if the pixel should be masked

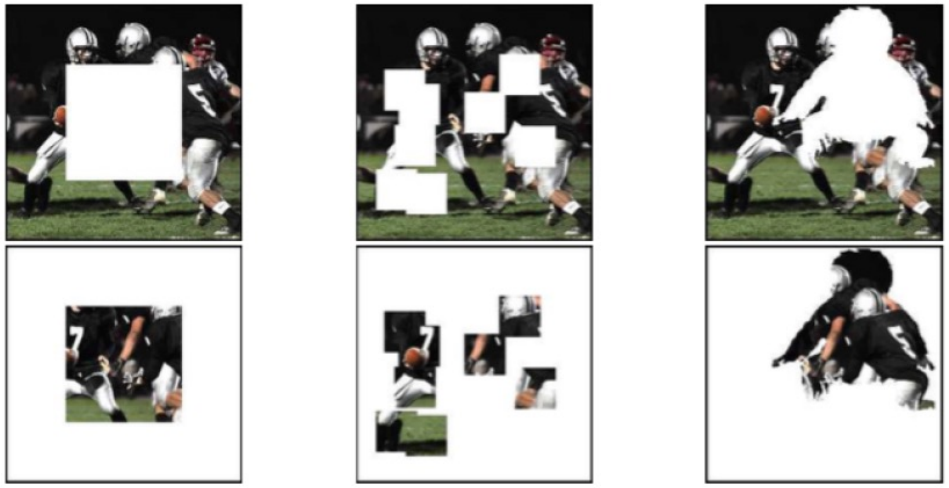

Training: Loss - Overview

- Trained completely from scratch to fill-up the masked areas
- **Problem:** multiple plausible solutions
- **Solution:** combining 2 losses:
  - $\mathcal{L}_{rec}$ **L2 reconstruction loss:** learn the structure of the missing region (average multiple modes in prediction)
  - $\mathcal{L}_{adv}$ **Adversarial loss:** make it look real (pick a mode from the distribution)

$$\min_F \mathcal{L} = \lambda_{rec}\mathcal{L}_{rec} + \lambda_{adv}\mathcal{L}_{adv}$$

$$\mathcal{L}_{rec}(x) = \left\| \hat{M} \odot \Big(x - F\big((1 - \hat{M}) \odot x\big) \Big) \right\|_2$$

$$\mathcal{L}_{adv} = \max_D \mathbb{E}_{x \in \mathcal{X}} \Big[\log(D(x)) + \log \Big(1 - D\big(F((1 - \hat{M}) \odot x)\big) \Big) \Big]$$

- Rq: The encoder-decoder is the generator, D is a CNN

Results

Dataset: StreetView Paris and ImageNet

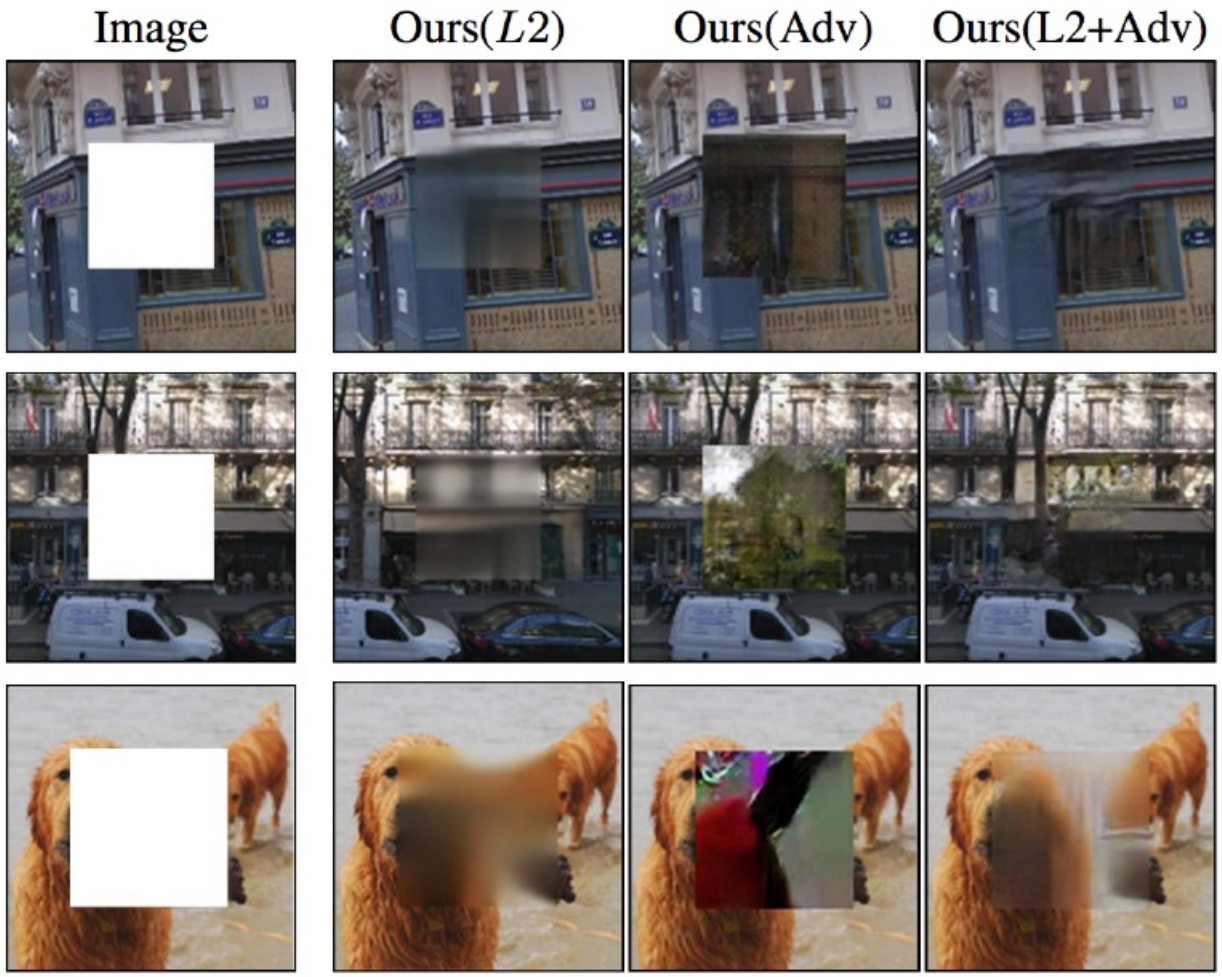

Semantic inpainting - Qualitative results

Generalizing inpainting: missing data encoder

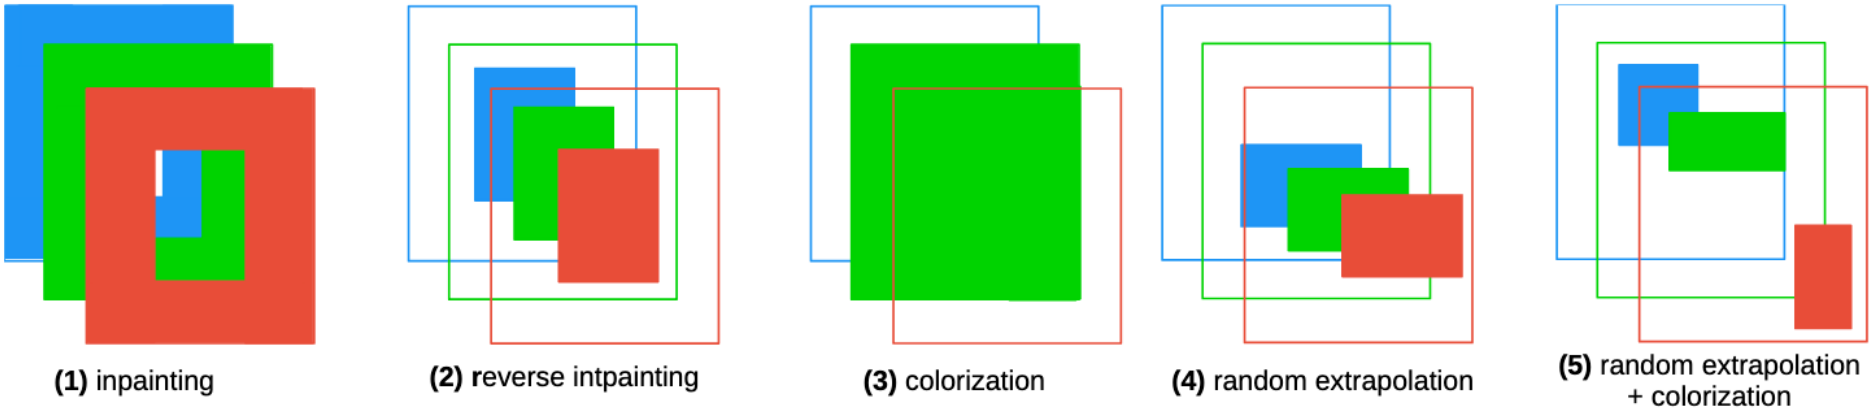

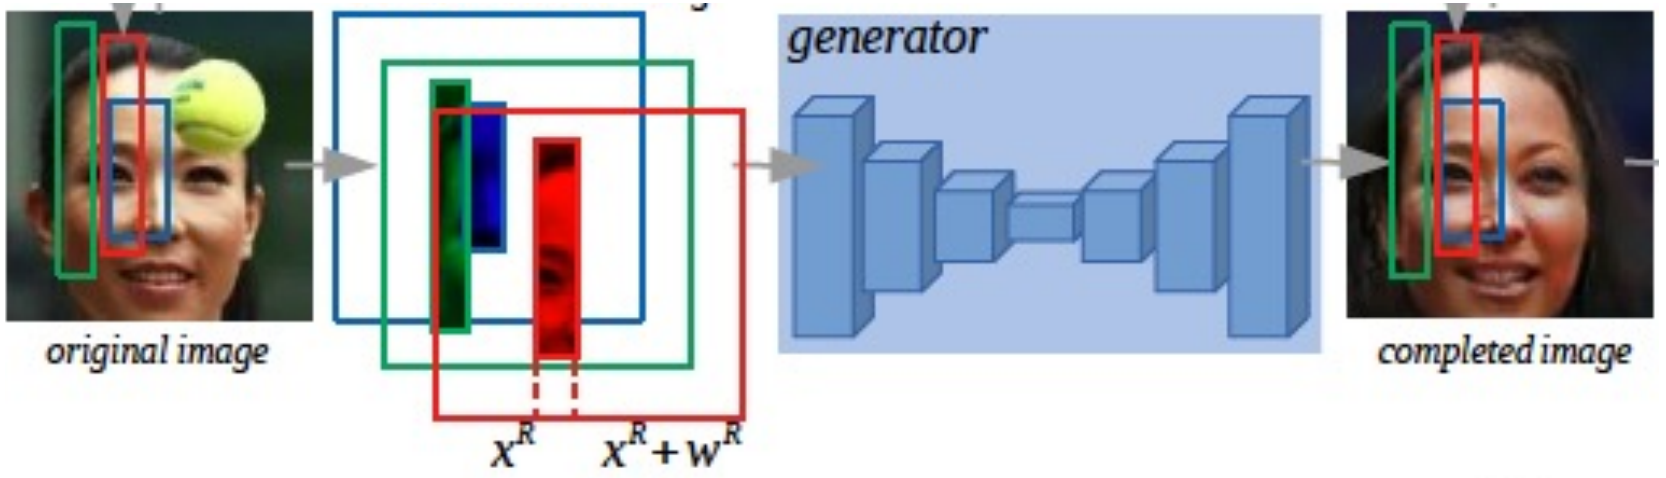

Results

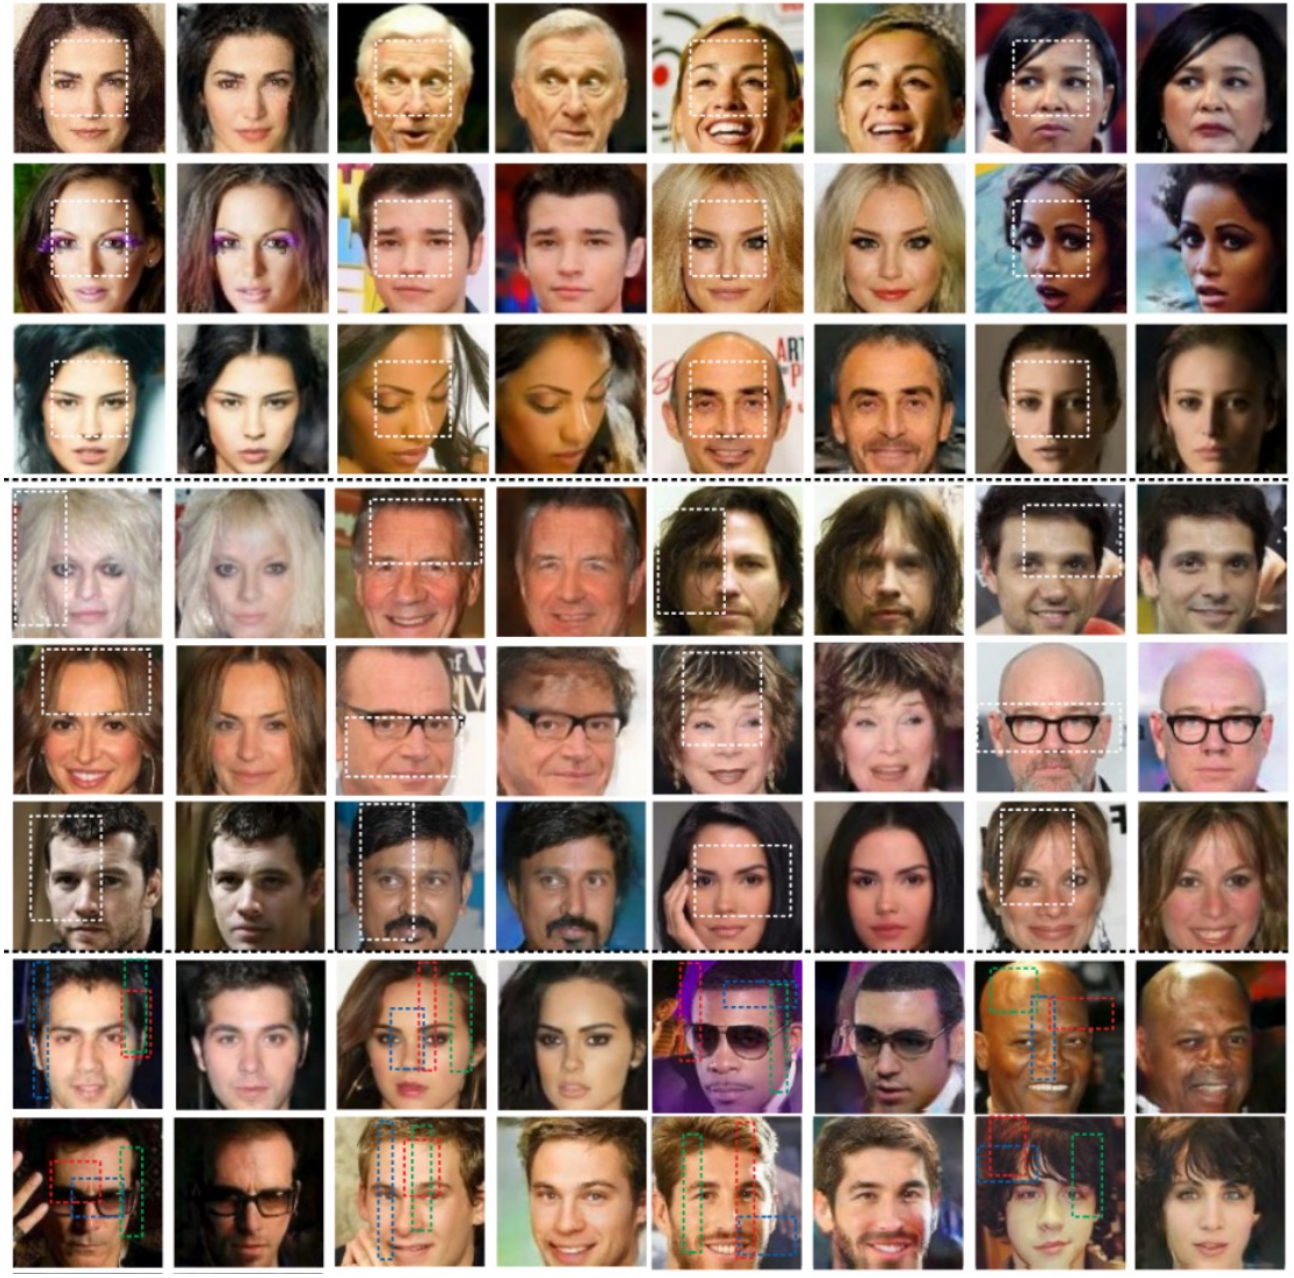

Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. Conditional GANs
    1. Principle
    2. Text2Image
    3. Image2Image
    4. Inpainting and general missing data encoder
    5. **Learning unpaired Transformation**

Unpaired Transformation

paired data

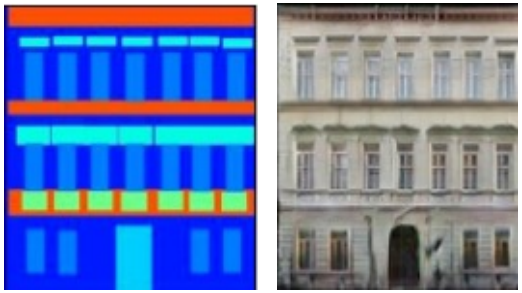

Transform an object from one domain to another ***without paired data***

photo van Gogh

Domain X

Domain Y

Cycle GAN

https://arxiv.org/abs/1703.10593
https://junyanz.github.io/CycleGAN/

Domain X

Domain Y

Domain X

$G_{X \to Y}$

Become similar to domain Y

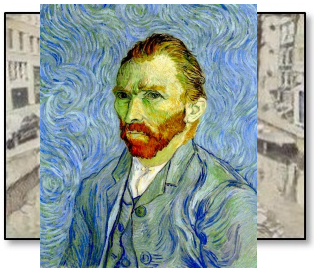

Not what we want

ignore input

D_Y → scalar

Input image belongs to domain Y or not

Domain Y

Cycle GAN

Domain X

Domain Y

as close as possible

$G_{X \to Y}$

$G_{Y \to X}$

Lack of information for reconstruction

D_Y

scalar

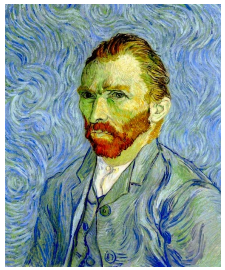

Input image belongs to domain Y or not

Domain Y

Cycle GAN

as close as possible

$G_{X \to Y}$

$G_{\mathrm{Y} \to \mathrm{X}}$

scalar: belongs to domain X or not

D_X

D_Y

scalar: belongs to domain Y or not

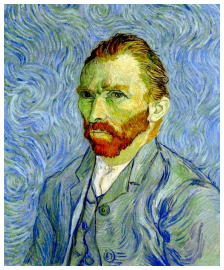

$G_{\mathrm{Y} \to \mathrm{X}}$

$G_{X \to Y}$

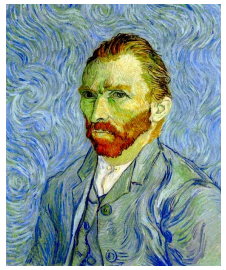

as close as possible

Results -- Cycle GAN

photo van Gogh

Domain X

Domain Y

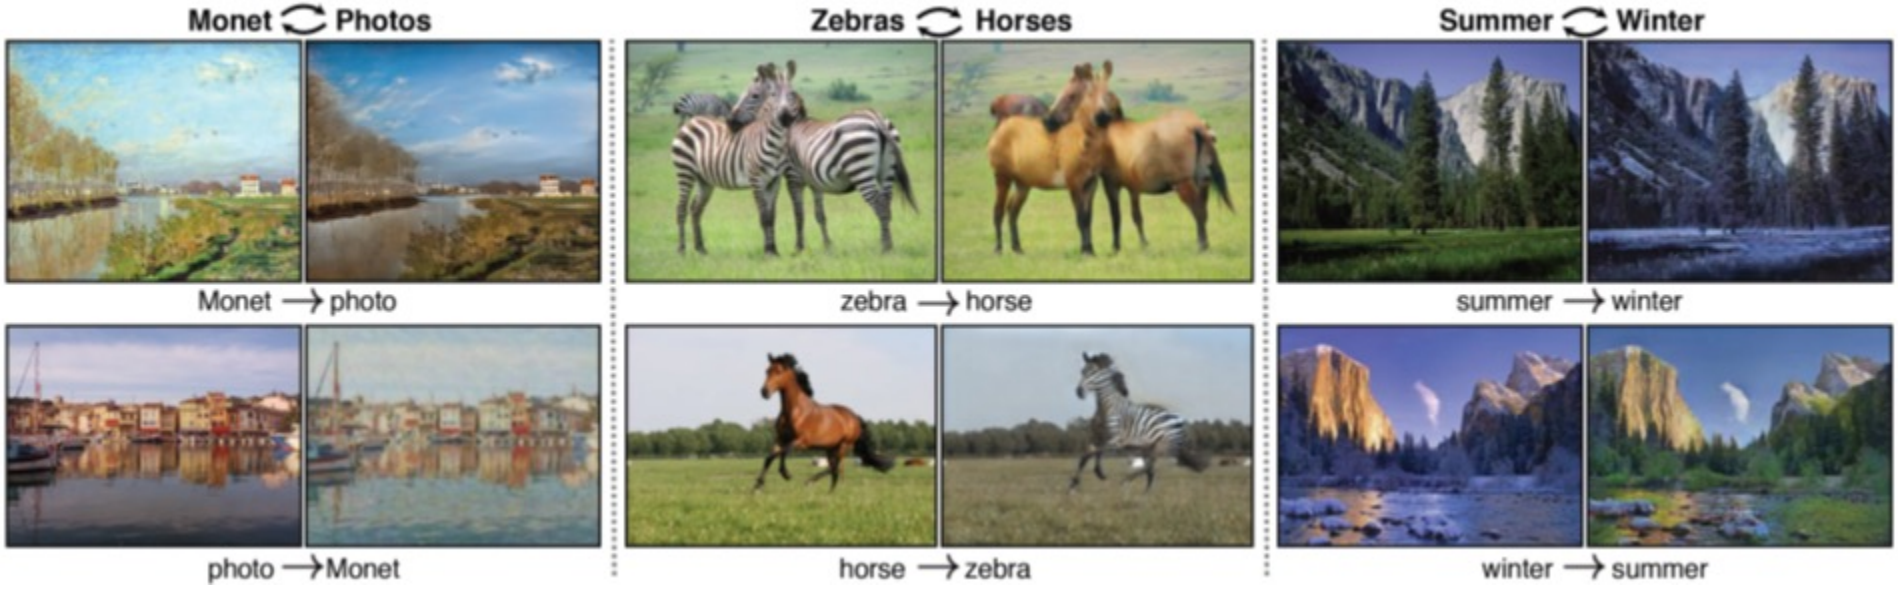

GANs: works in progress

A lot of things to better understand, to use, adapt, …

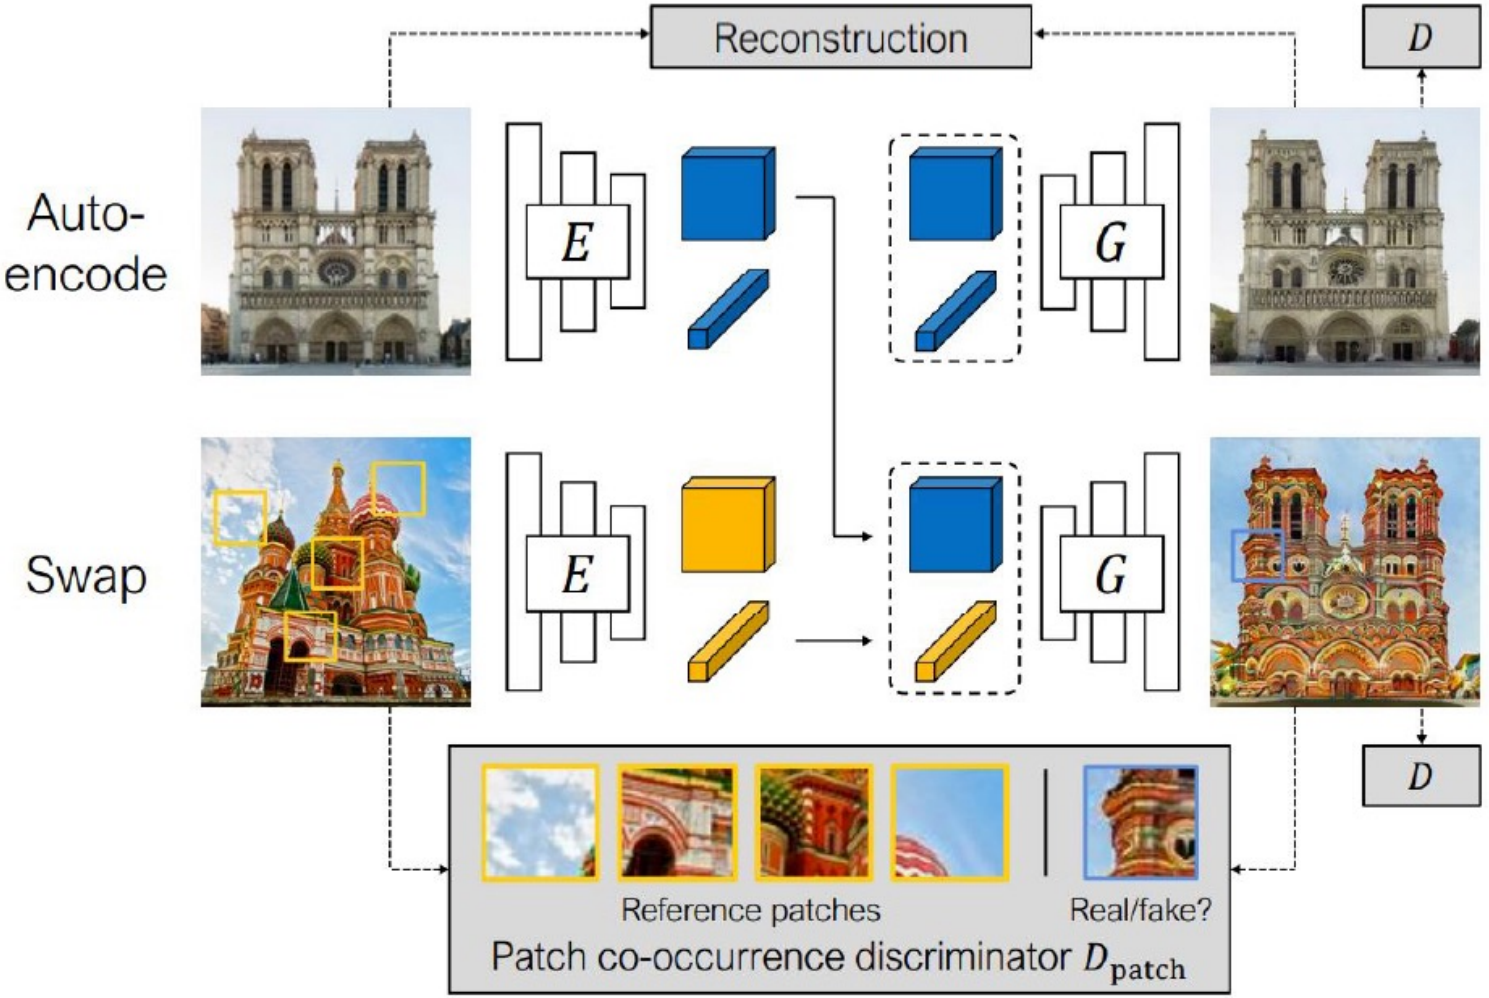

Generative models

Outline

1. Preview: Auto-Encoders, VAE
2. Generative models with GAN
3. GAN architectures
4. Editing
5. Conditional GANs
6. **Diffusion models**

Generative Modelling with Diffusion models

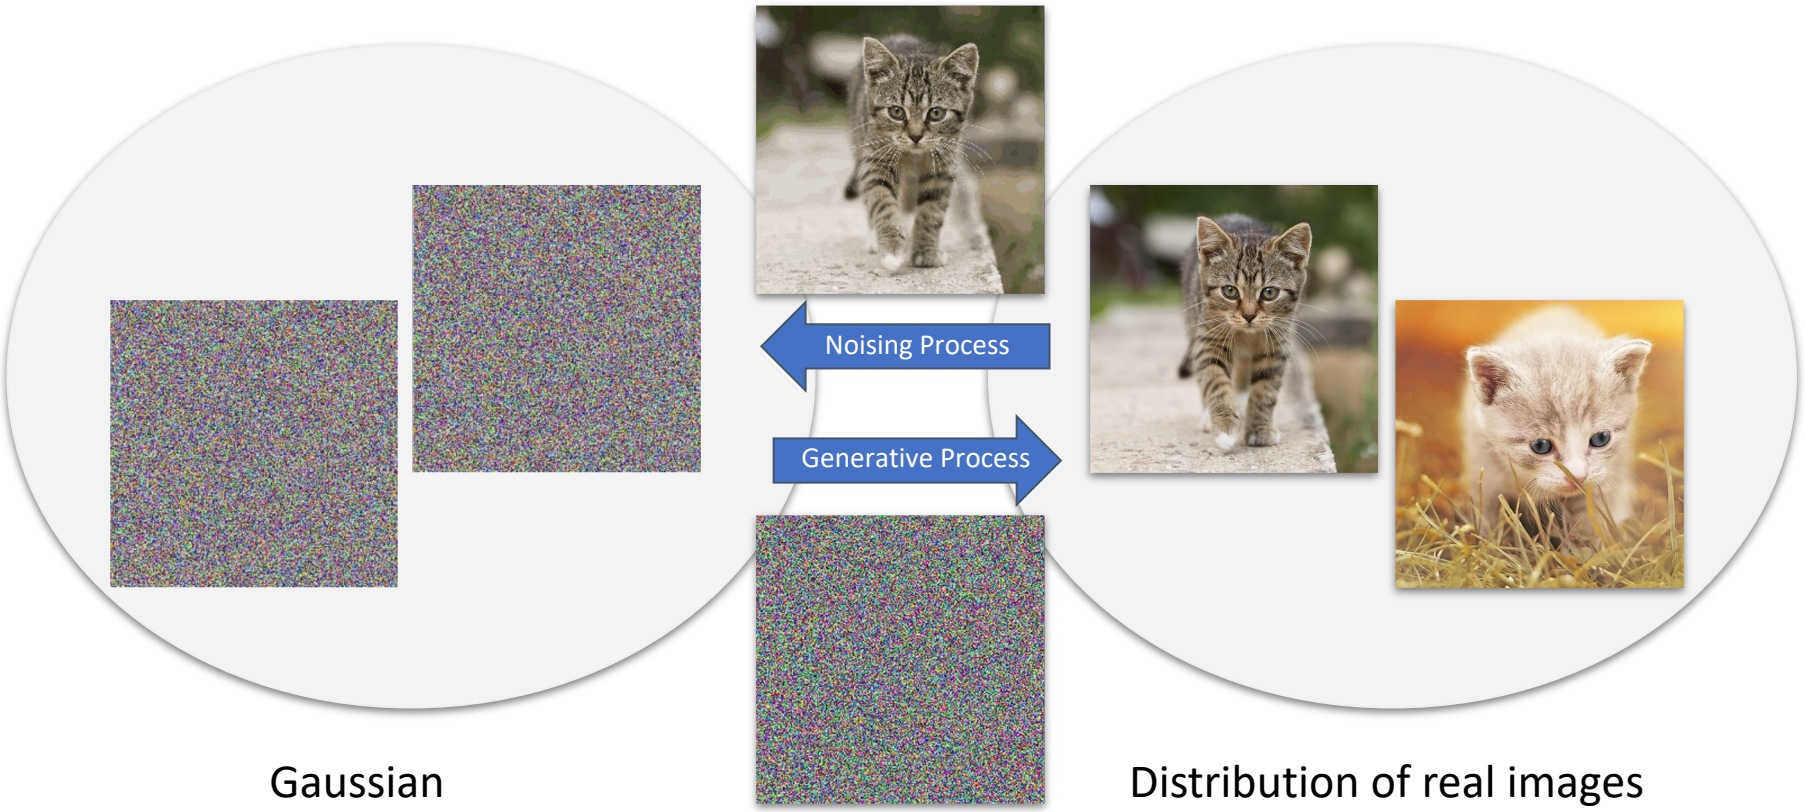

Generative Modelling with Diffusion models

DDPM: Denoising Diffusion Probabilistic Models
In context with other genenrative Models: