



# SC395:

# Image Generative Models in Computer Vision

## Viraj Shah

Lecture 2  
Jan 7<sup>th</sup>, 2026

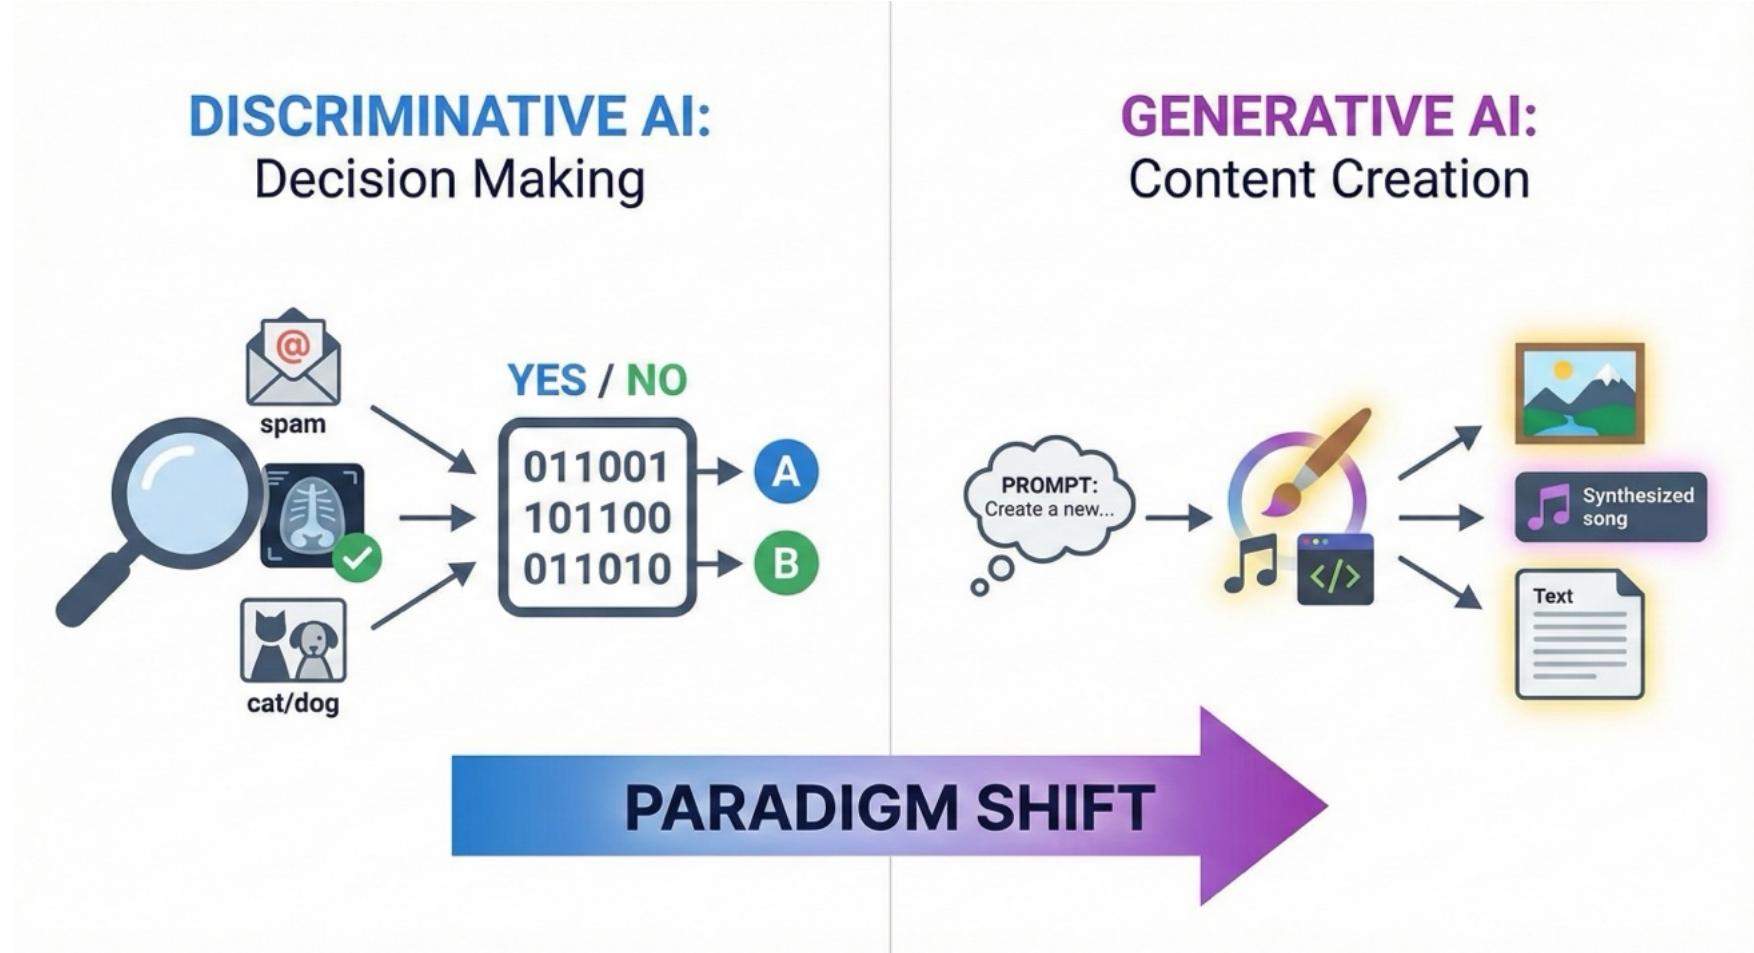
[sc395.virajshah.com](http://sc395.virajshah.com)



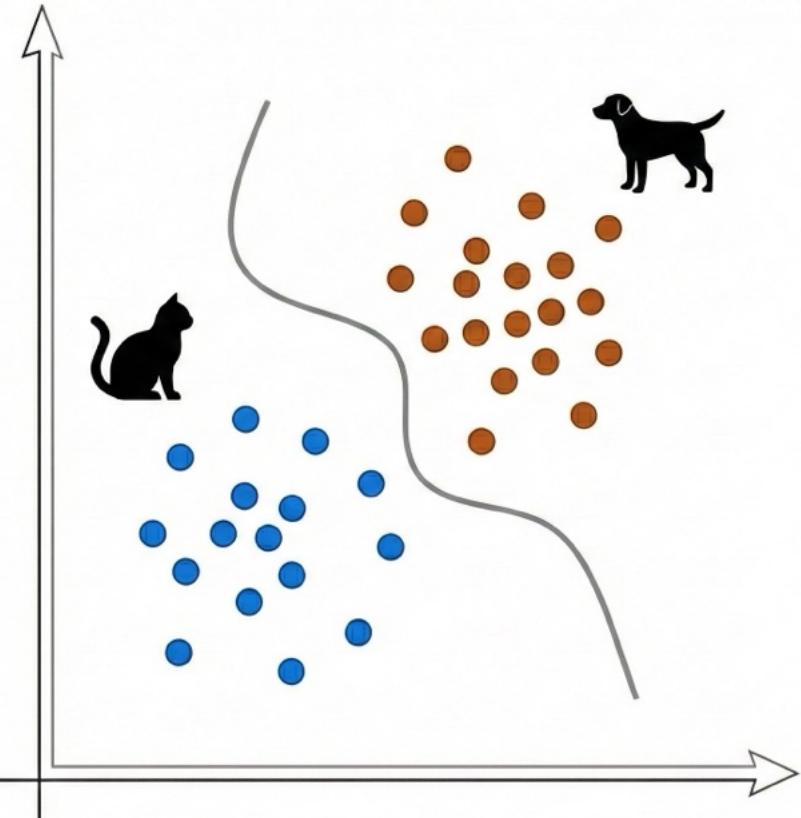


# Lecture 2: GANs and Applications

# Paradigm Shift: Discriminative Models → Generative Models

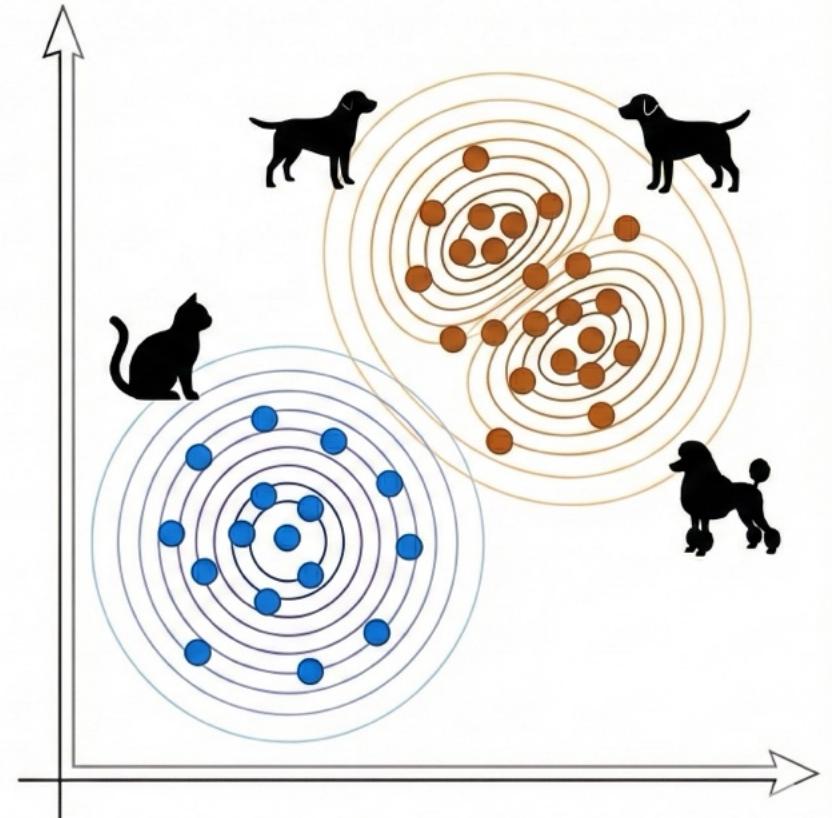


# Discriminative Models vs. Generative Models



Predict the label given the Input

Learns  $P(Y/X)$



Generate new content by understanding **the abstract patterns of the data**

Learns  $P(X, Y)$



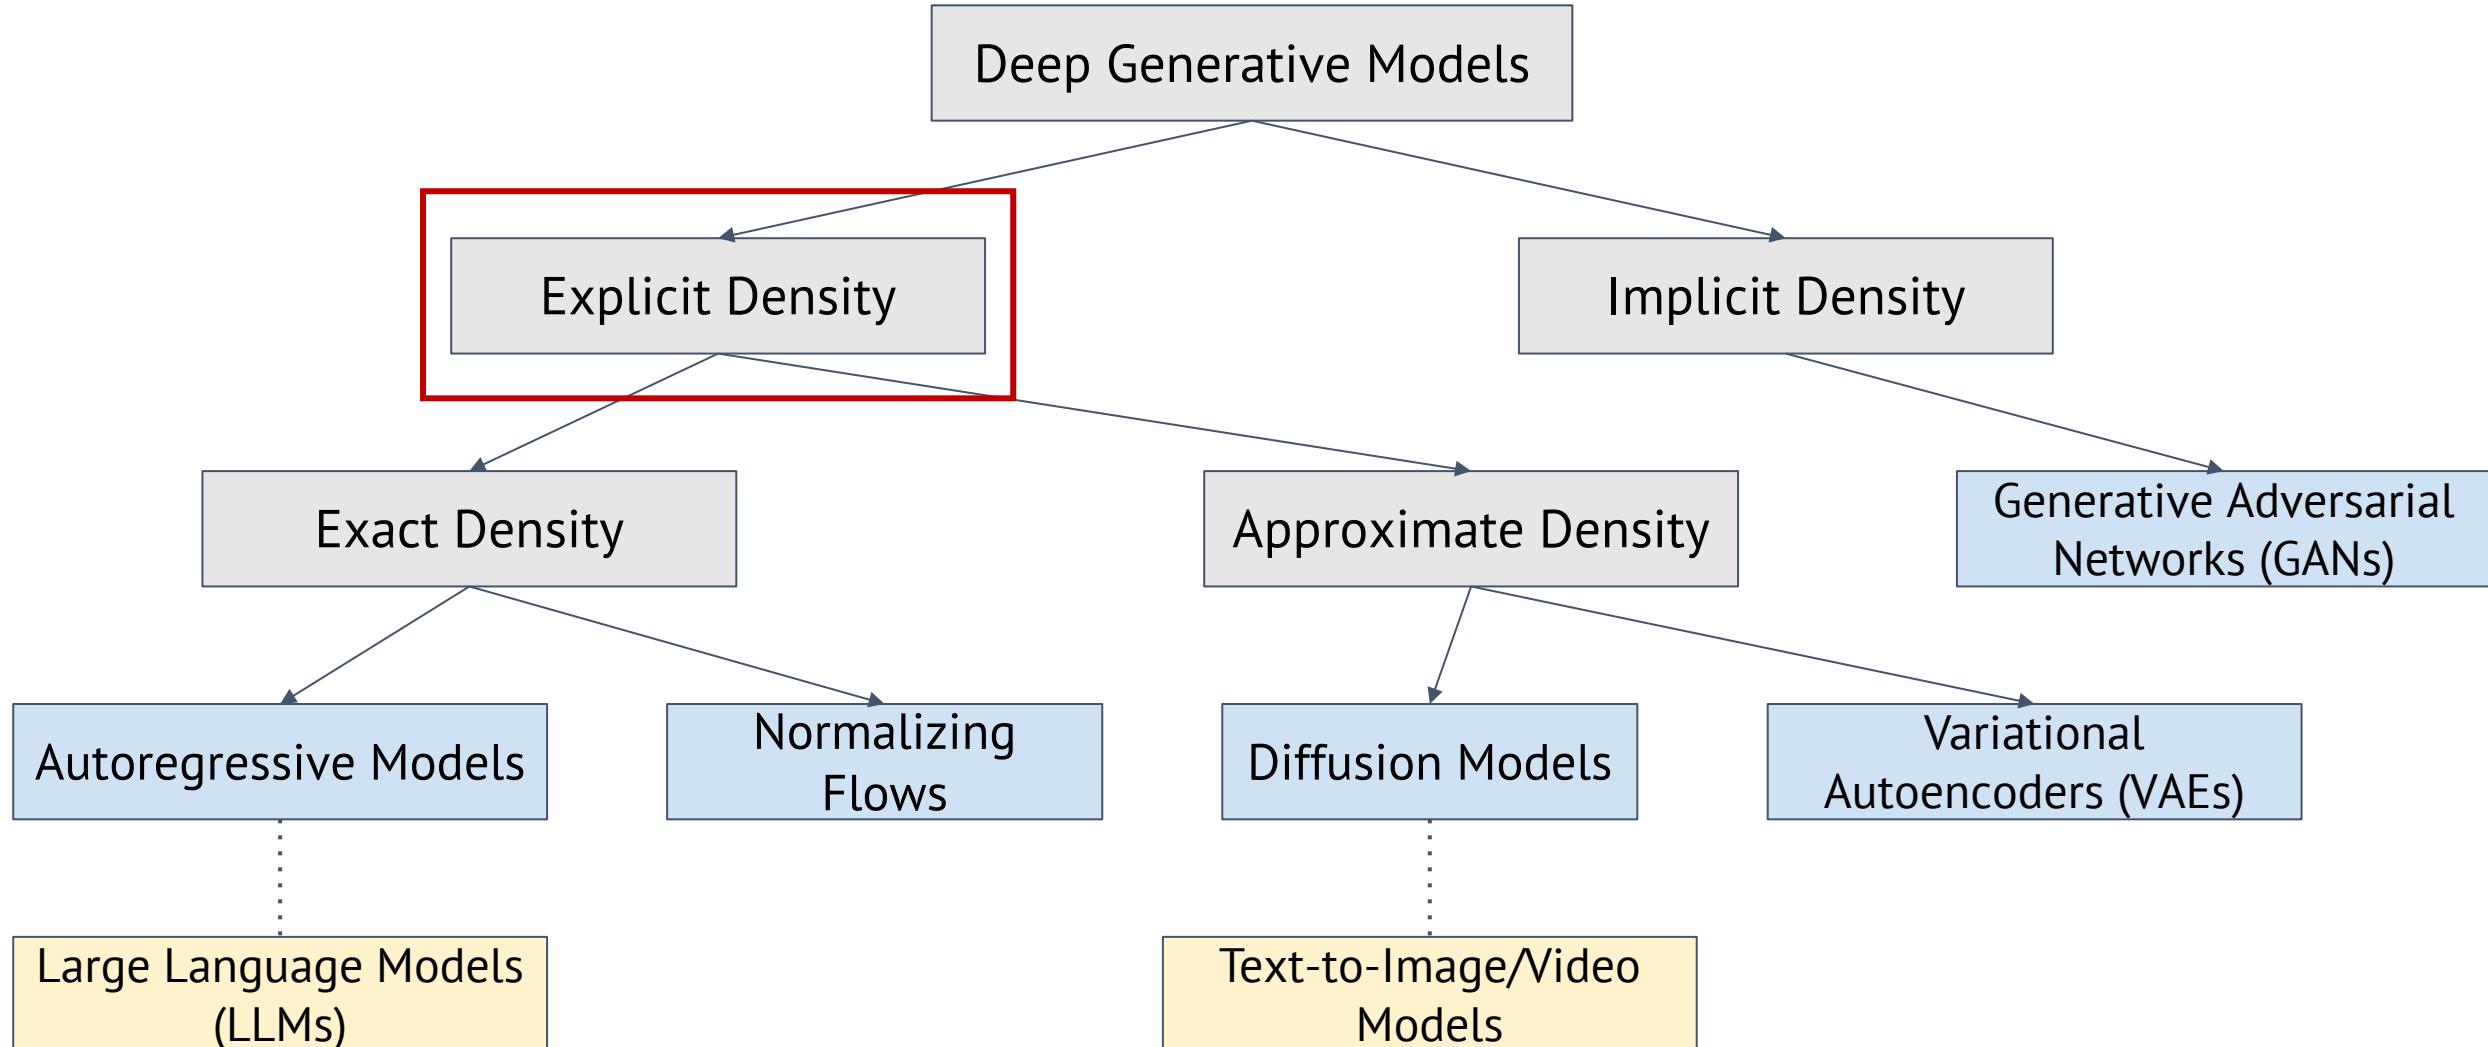
# Naïve Attempt: Maximum Likelihood Estimation

- Our data distribution is  $p_{data}(x)$  : we have  $m$  samples:  $(x^{(1)}, x^{(2)}, \dots, x^{(m)})$
- Our generative model is parameterized by  $\theta$ ,
  - Likelihood of each sample is given by:  $p_{model}(x^{(i)}; \theta)$
- MLE: choose the parameters for the model that maximize the likelihood of the training data:

$$\theta^* = \arg \max_{\theta} \prod_{i=1}^m p_{model}(x^{(i)}; \theta)$$

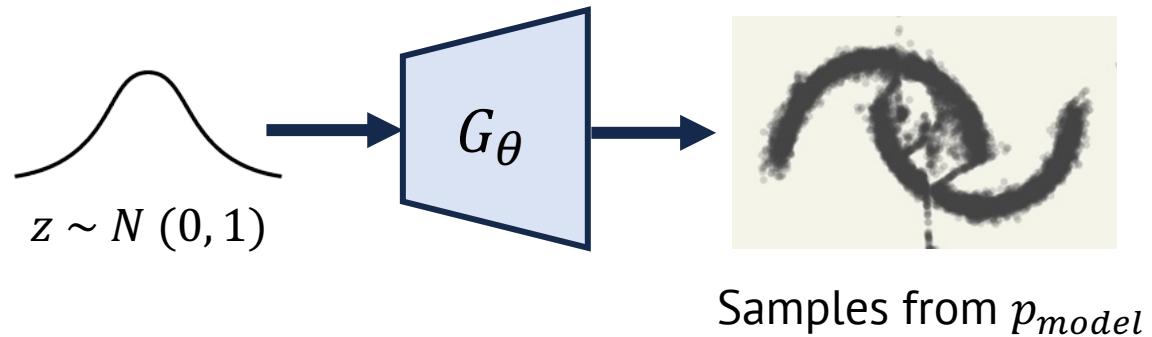
**Minimize KL Divergence**

# Deep Generative Models



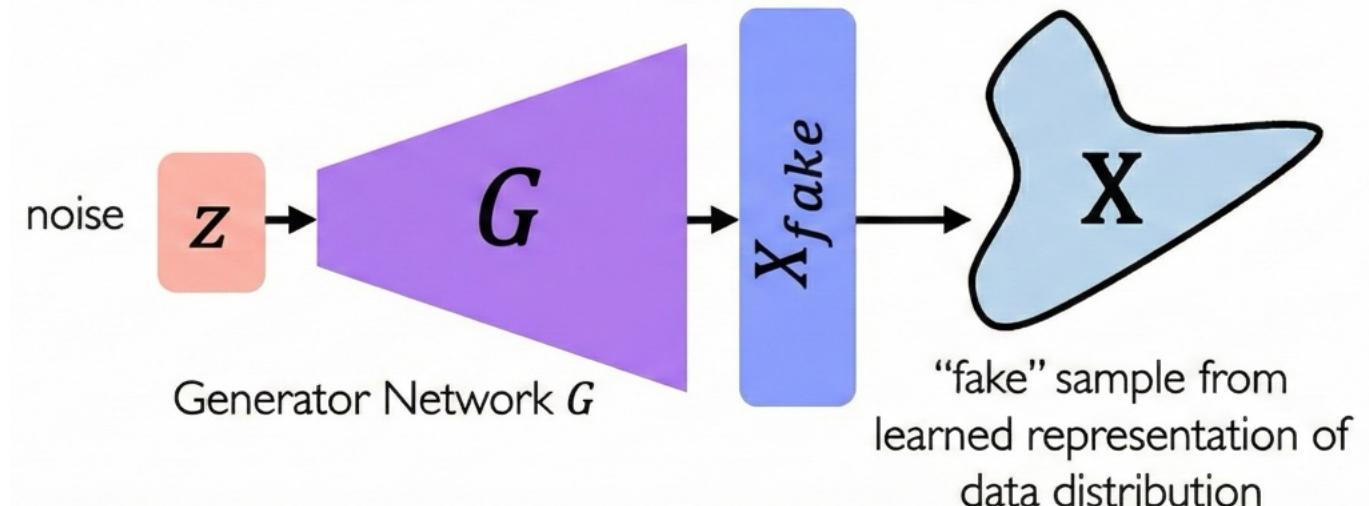
# Implicit Generative Models

- No explicit definition of the density function  $p_{model}$
- Interact only indirectly with  $p_{model}$ :
  - via **sampling** from it
  - No need to parameterize  $p_{model}$
  - Instead, use **latent variable**  $z$  and generator  $G_\theta$  such that  $G_\theta(z) \sim p_{model}$



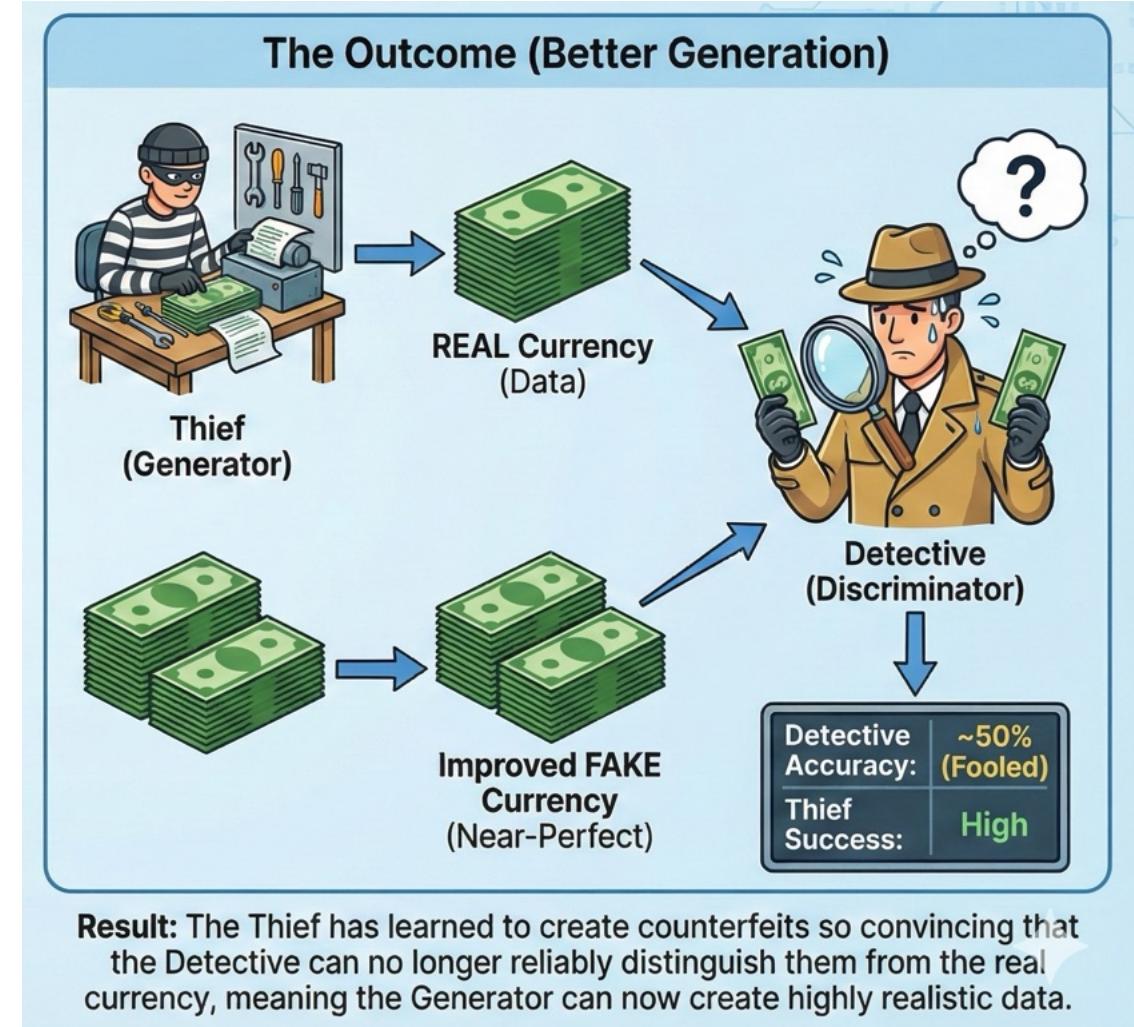
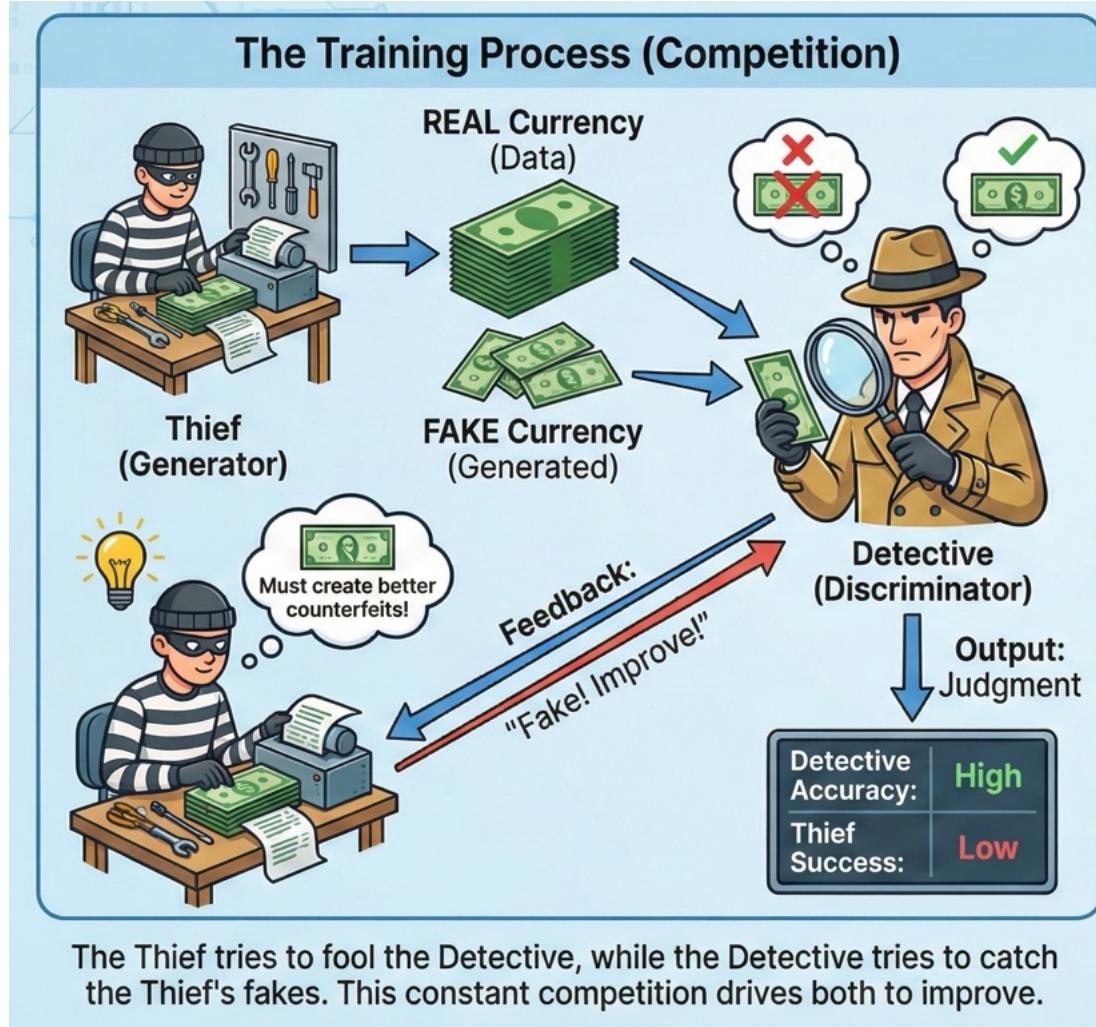
# Implicit Models: What if we just want to sample?

- **Idea:** don't explicitly model density, just sample to generate new instances
- **Problem:** want to sample from complex distribution -- can't do this directly!
- **Solution:** sample from something simple (e.g. noise), learn a transformation to the data distribution



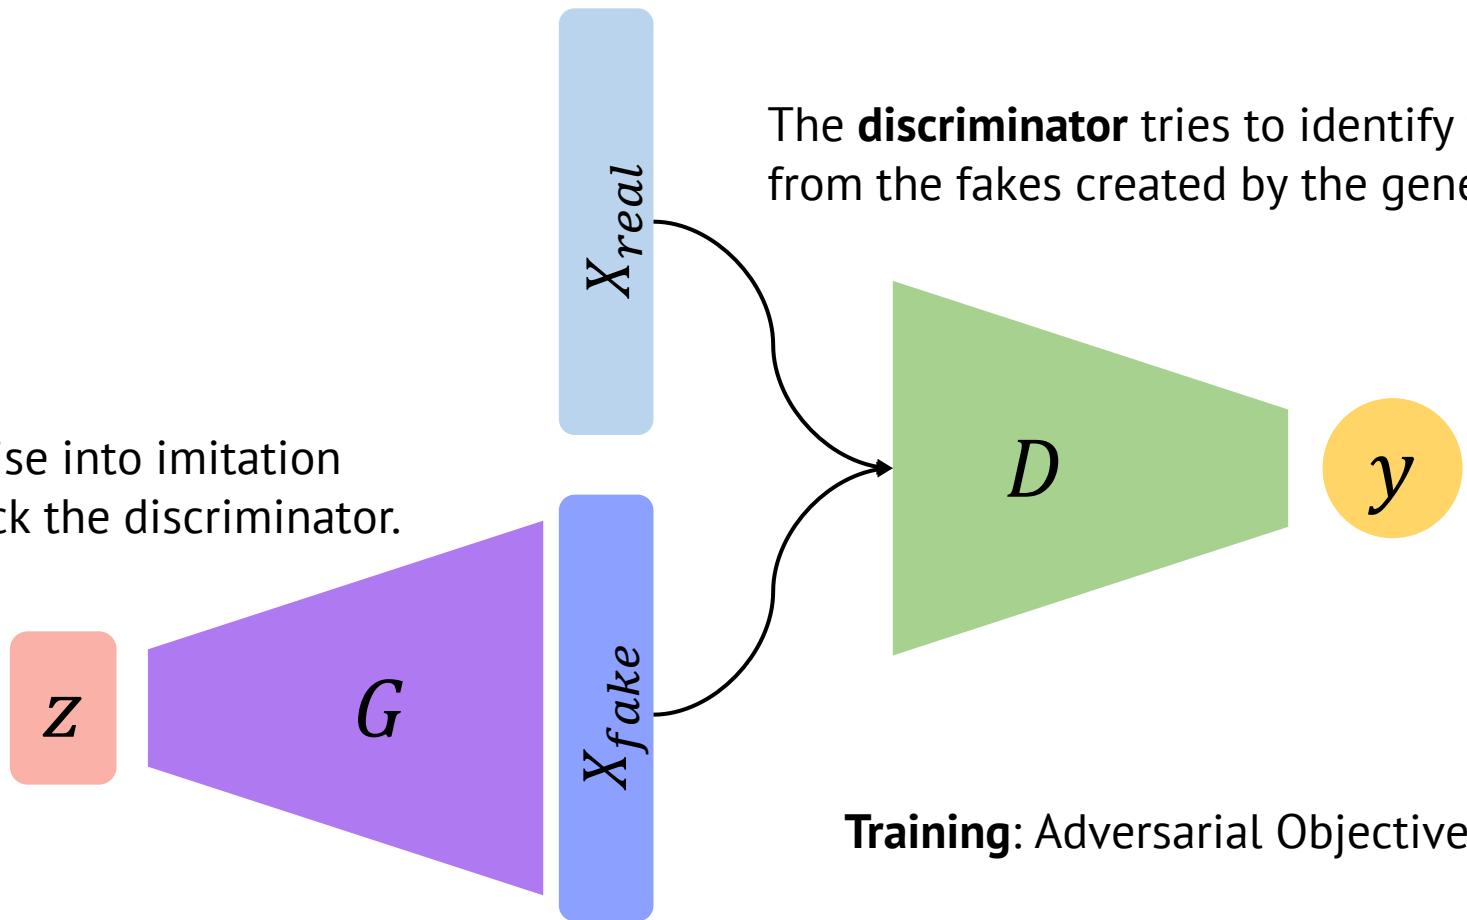
# Generative Adversarial Networks (GANs)

Obtain a generative model by having two neural networks competing with each other



# Generative Adversarial Networks (GANs)

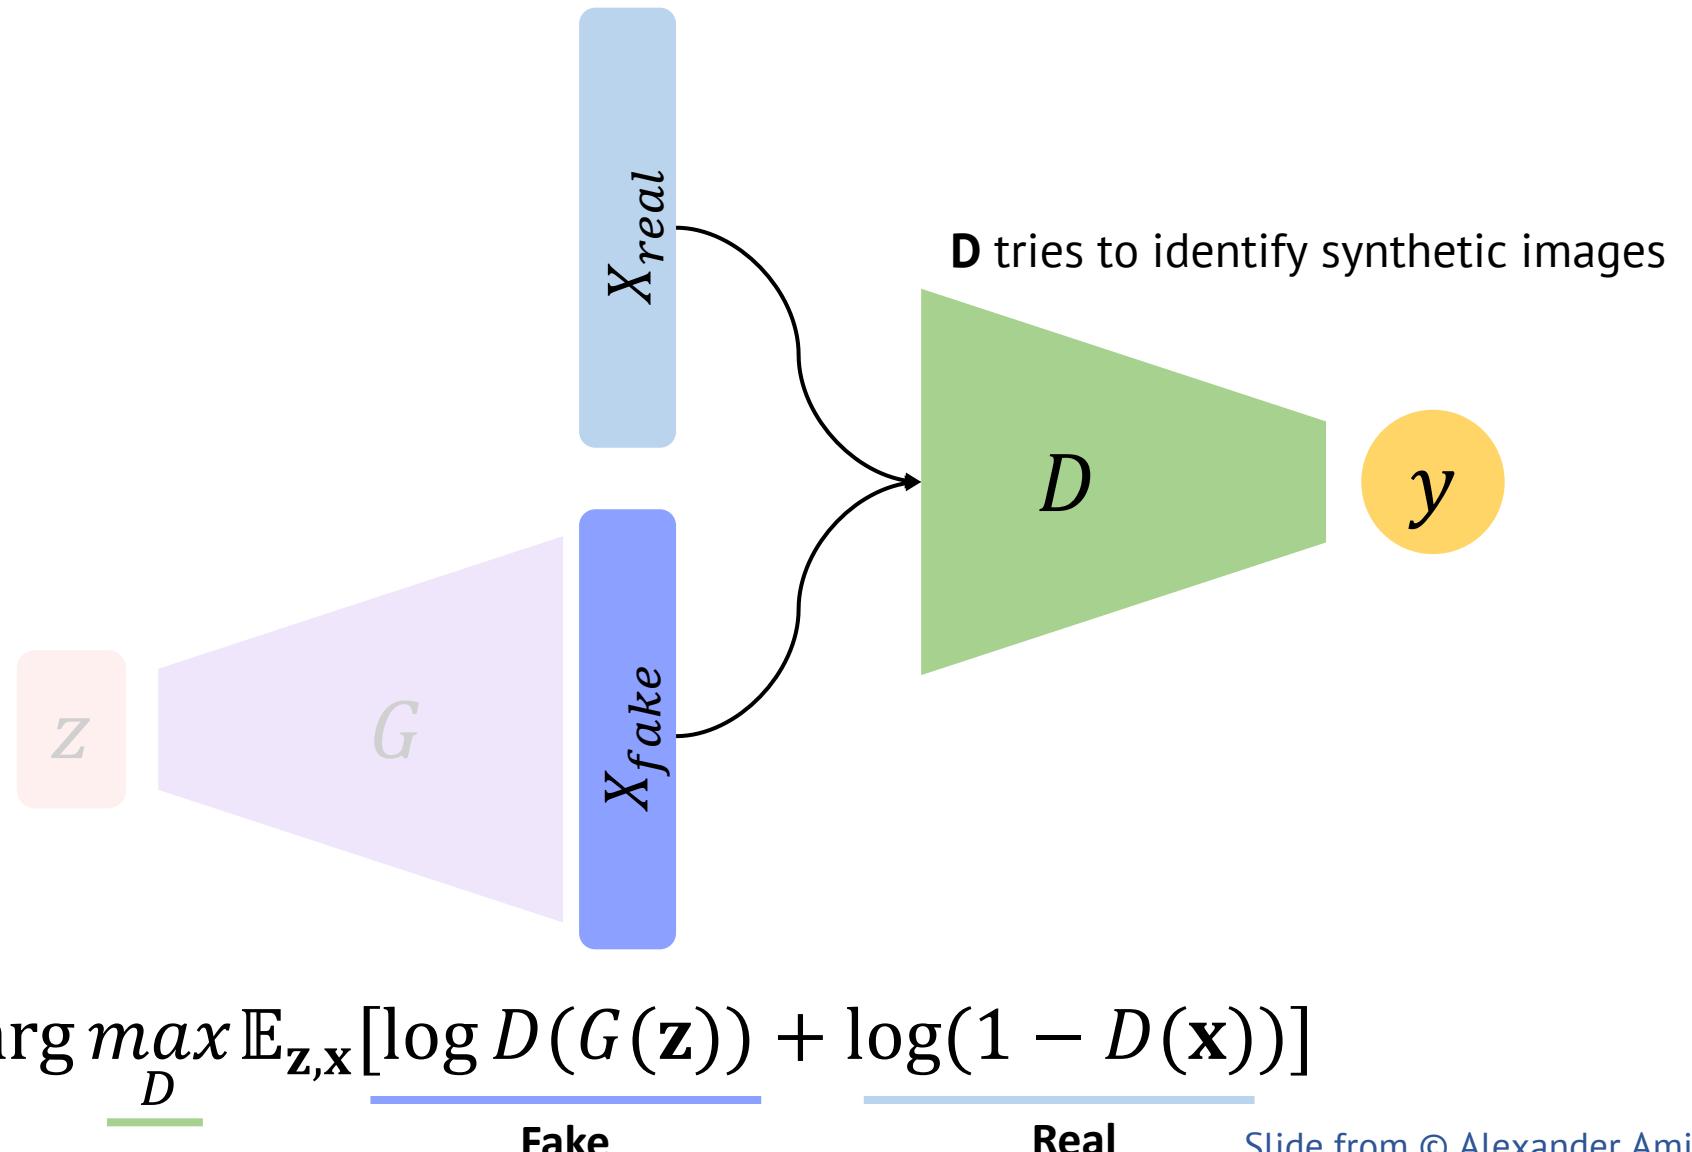
The **generator** turns noise into imitation of the data trying to trick the discriminator.



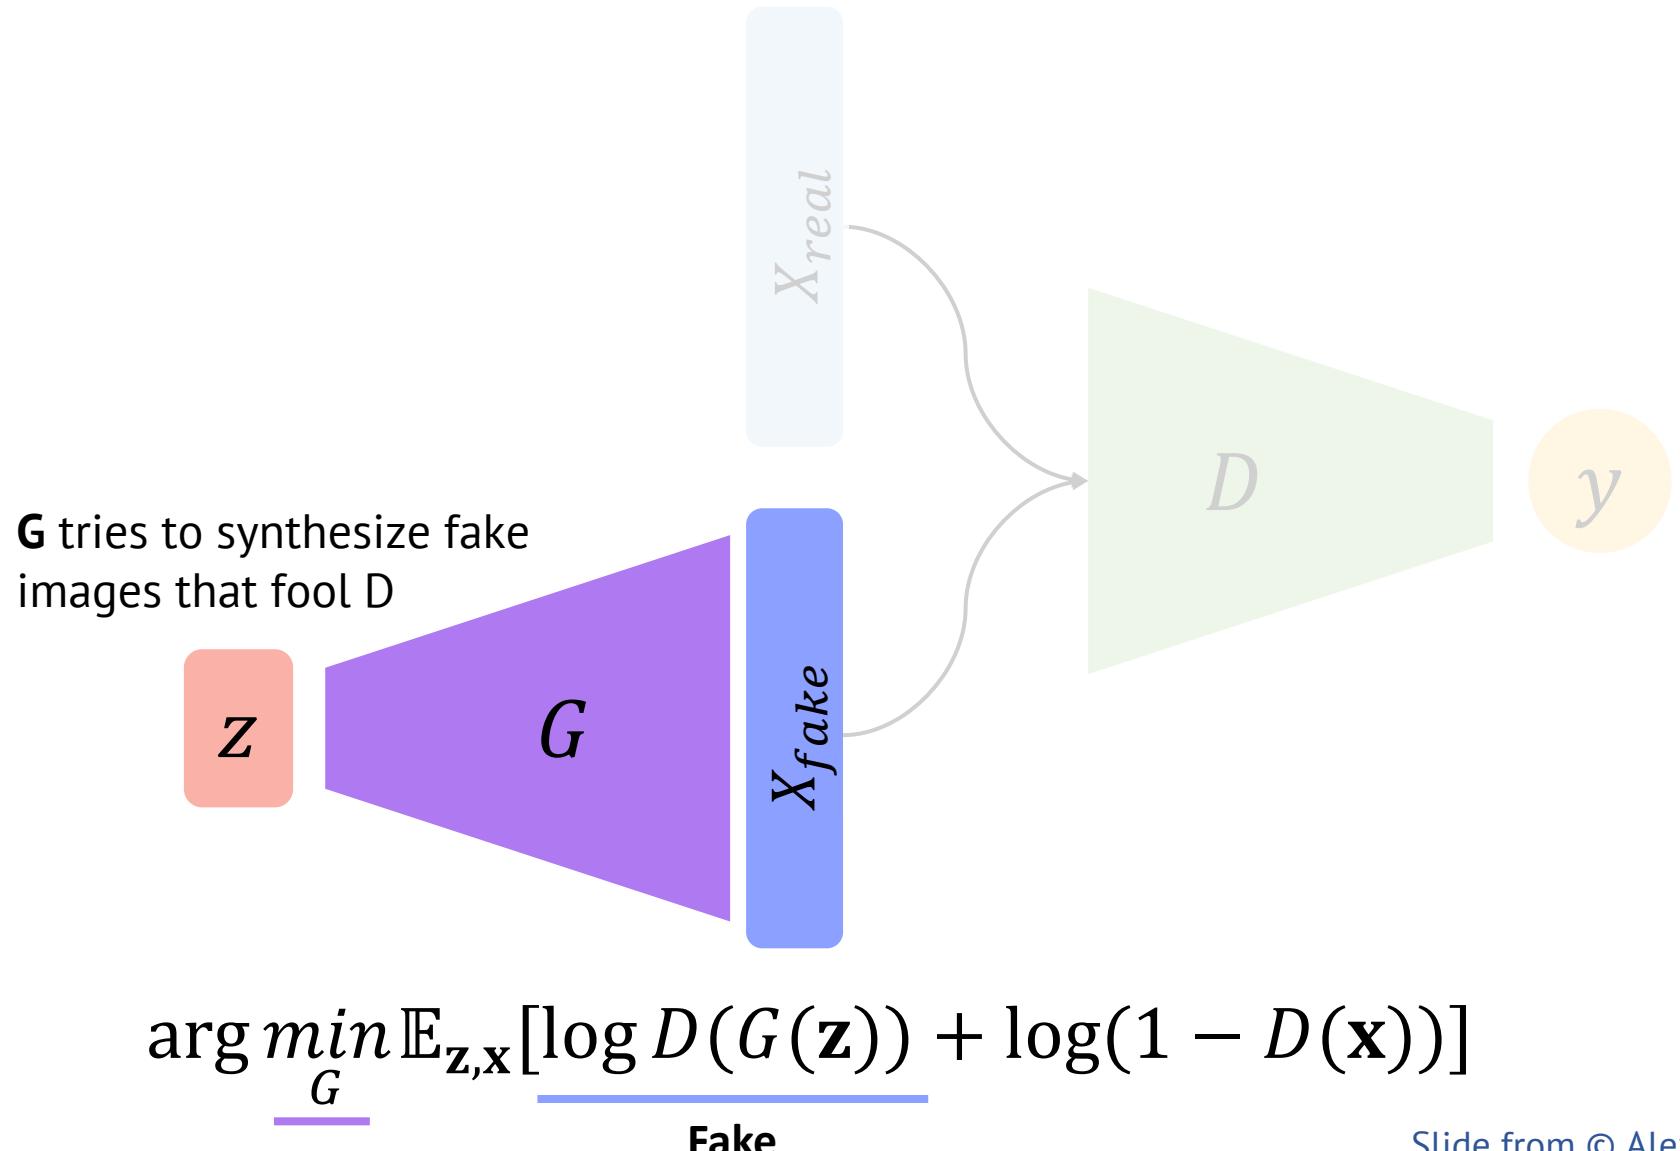
**Training:** Adversarial Objectives for  $G$  and  $D$

**Optimum:**  $G$  reproduces true data distribution

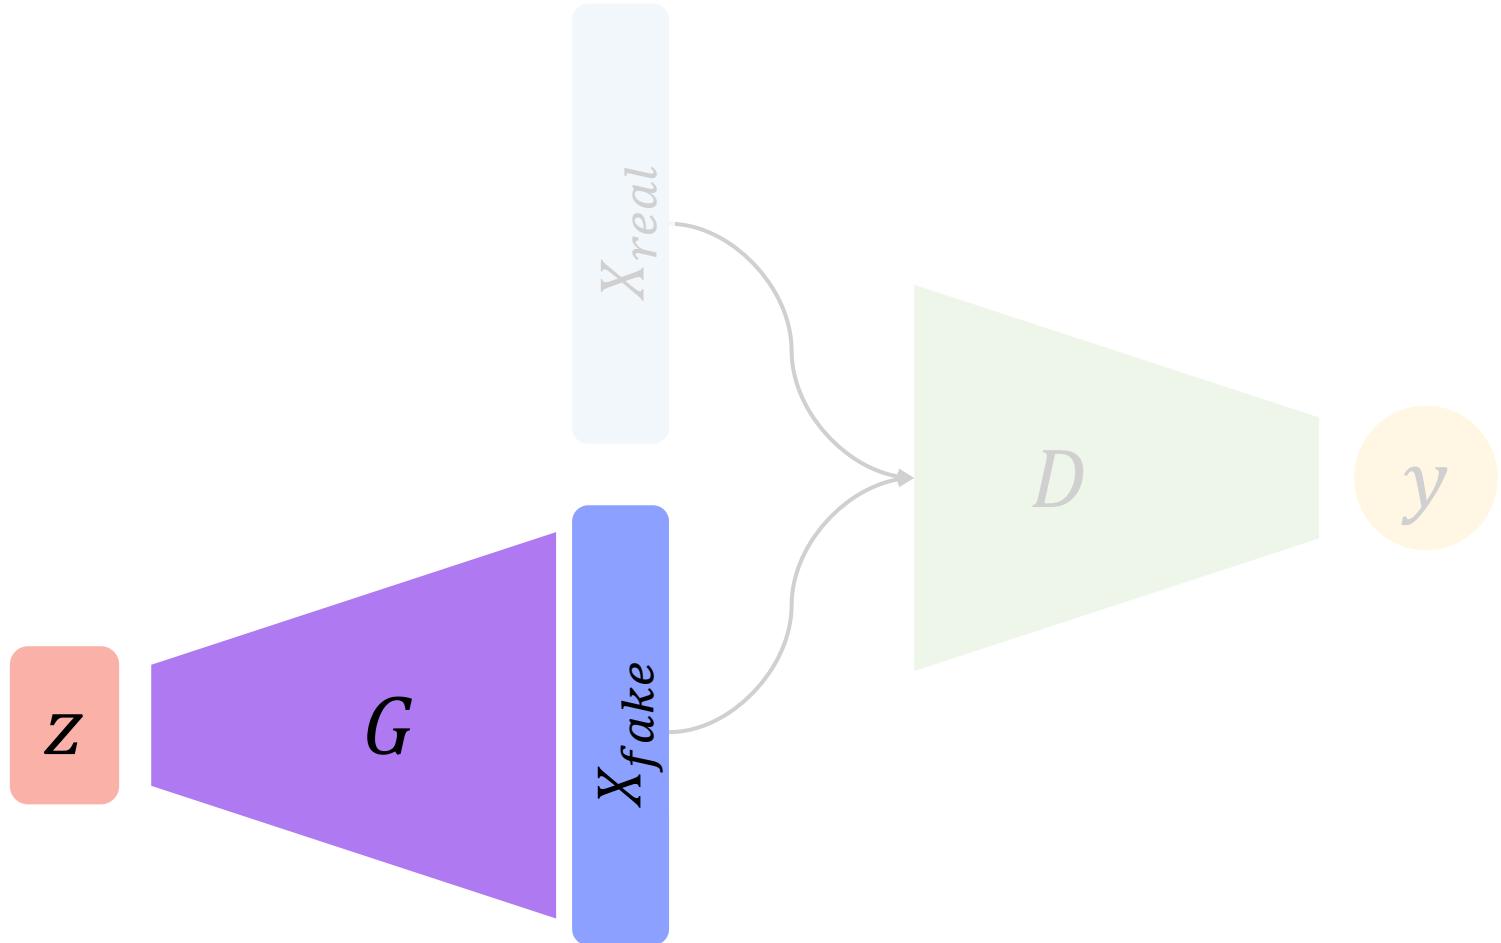
# Generative Adversarial Networks (GANs)



# Generative Adversarial Networks (GANs)



# Generative Adversarial Networks (GANs)



**After training, use the generator to synthesize new data**



# Optimal solution of the minimax game

Optimal **discriminator** cost:

$$V(G, D) = \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p(z)} [\log(1 - D(G(z)))]$$

$$\begin{aligned} \nabla_y [a \log y + b \log(1 - y)] = 0 &\implies y^* = \frac{a}{a + b} \quad \forall \quad [a, b] \in \mathbb{R}^2 \setminus [0, 0] \\ &\implies D^*(x) = \frac{p_{\text{data}}(x)}{(p_{\text{data}}(x) + p_g(x))} \end{aligned}$$



# Generator Objective under Optimal Discriminator

$$V(G, D^*) = \mathbb{E}_{x \sim p_{\text{data}}} [\log D^*(x)] + \mathbb{E}_{x \sim p_g} [\log(1 - D^*(x))]$$



# Generator Objective under Optimal Discriminator

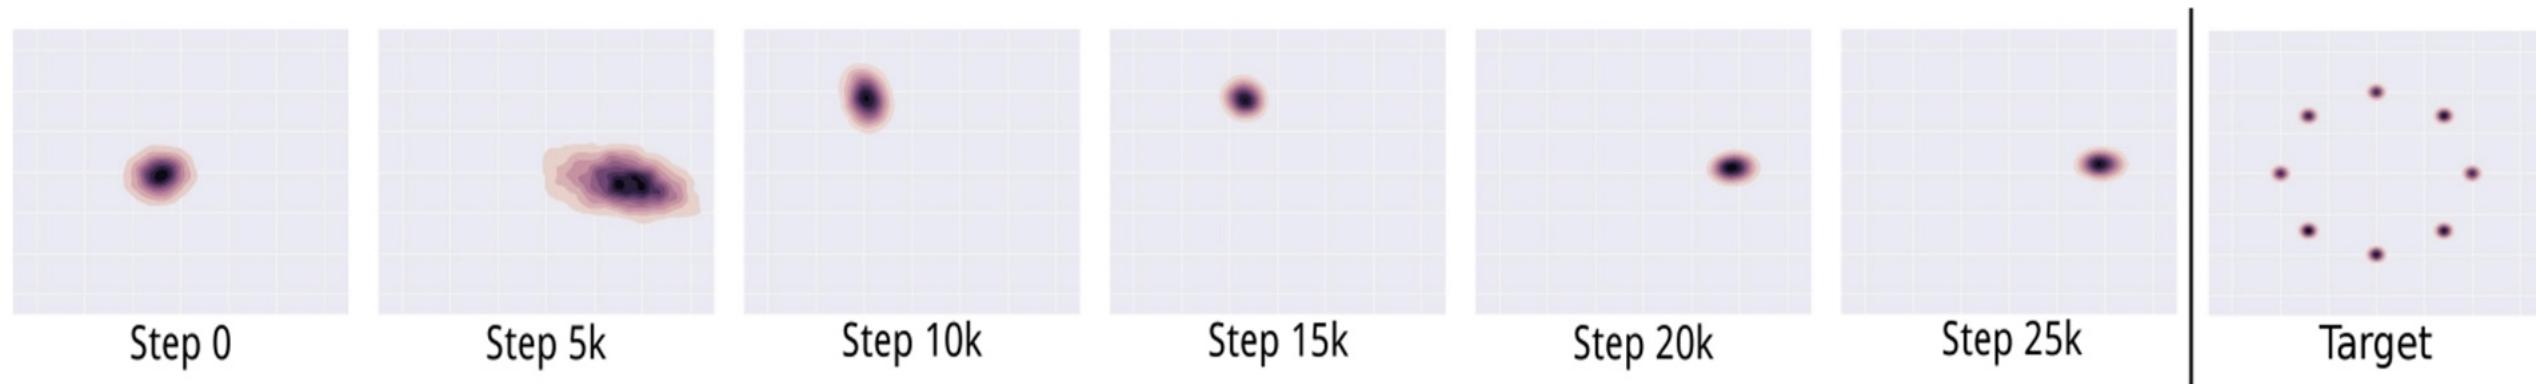
$$\begin{aligned} V(G, D^*) &= \mathbb{E}_{x \sim p_{\text{data}}} [\log D^*(x)] + \mathbb{E}_{x \sim p_g} [\log(1 - D^*(x))] \\ &= \mathbb{E}_{x \sim p_{\text{data}}} \left[ \log \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_g(x)} \right] + \mathbb{E}_{x \sim p_g} \left[ \log \frac{p_g(x)}{p_{\text{data}}(x) + p_g(x)} \right] \\ &= -\log(4) + \underbrace{KL \left( p_{\text{data}} \parallel \left( \frac{p_{\text{data}} + p_g}{2} \right) \right)}_{(\text{Jensen-Shannon Divergence (JSD) of } p_{\text{data}} \text{ and } p_g) \geq 0} + KL \left( p_g \parallel \left( \frac{p_{\text{data}} + p_g}{2} \right) \right) \end{aligned}$$

$$V(G^*, D^*) = -\log(4) \text{ when } p_g = p_{\text{data}}$$

# Mode Collapse

Try:

<https://poloclub.github.io/ganlab/>



Standard GAN training collapses when the true distribution is a mixture of gaussians (Figure from Metz et al 2016)

# Back to our GAN objective

- Is it feasible to run the inner optimization to completion?
- For this specific objective, would it create problems if we were able to do so?

$$V(D, G) = \mathbb{E}_{x \sim p(x)}[\log D(x)] + \boxed{\mathbb{E}_{z \sim q(z)}[\log(1 - D(G(z)))]}$$


---

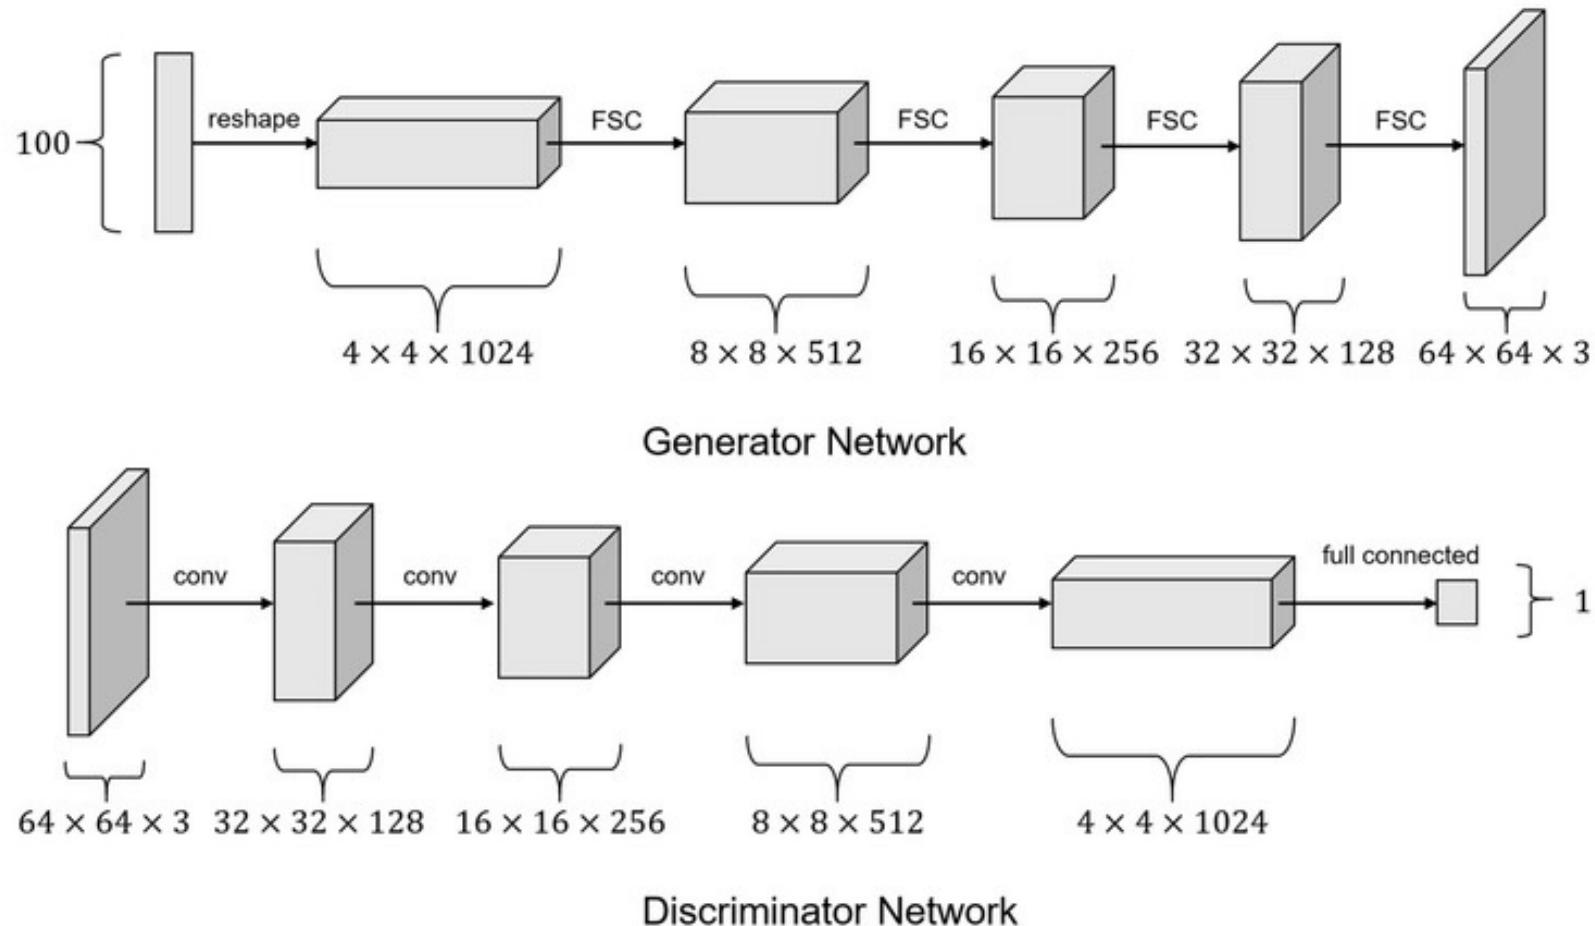
$$\nabla_{\theta_G} V(D, G) = \nabla_{\theta_G} \mathbb{E}_{z \sim q(z)} [\log(1 - D(G(z)))]$$

- $\nabla_a \log(1 - \sigma(a)) = \frac{-\nabla_a \sigma(a)}{1 - \sigma(a)} = \frac{-\sigma(a) (1 - \sigma(a))}{1 - \sigma(a)} = -\sigma(a) = -D(G(z))$

- Gradient goes to 0 if  $D$  is confident, i.e.  $D(G(z)) \rightarrow 0$

- Minimize  $-\mathbb{E}_{z \sim q(z)}[\log D(G(z))]$  for **Generator** instead (keep Discriminator as it is)

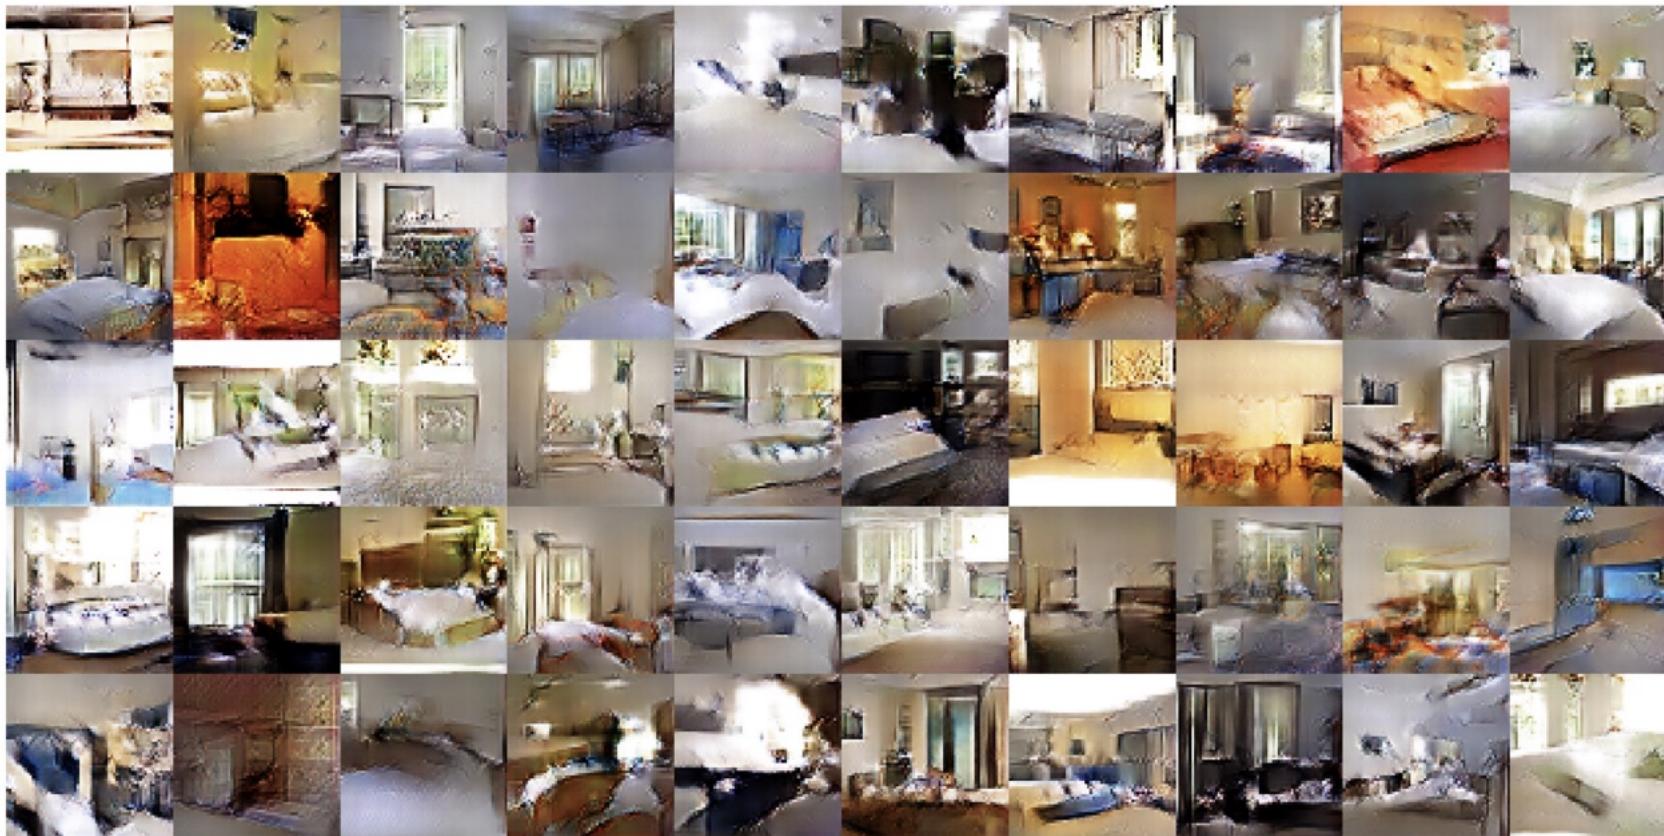
# Deep Convolutional GAN (DCGAN)



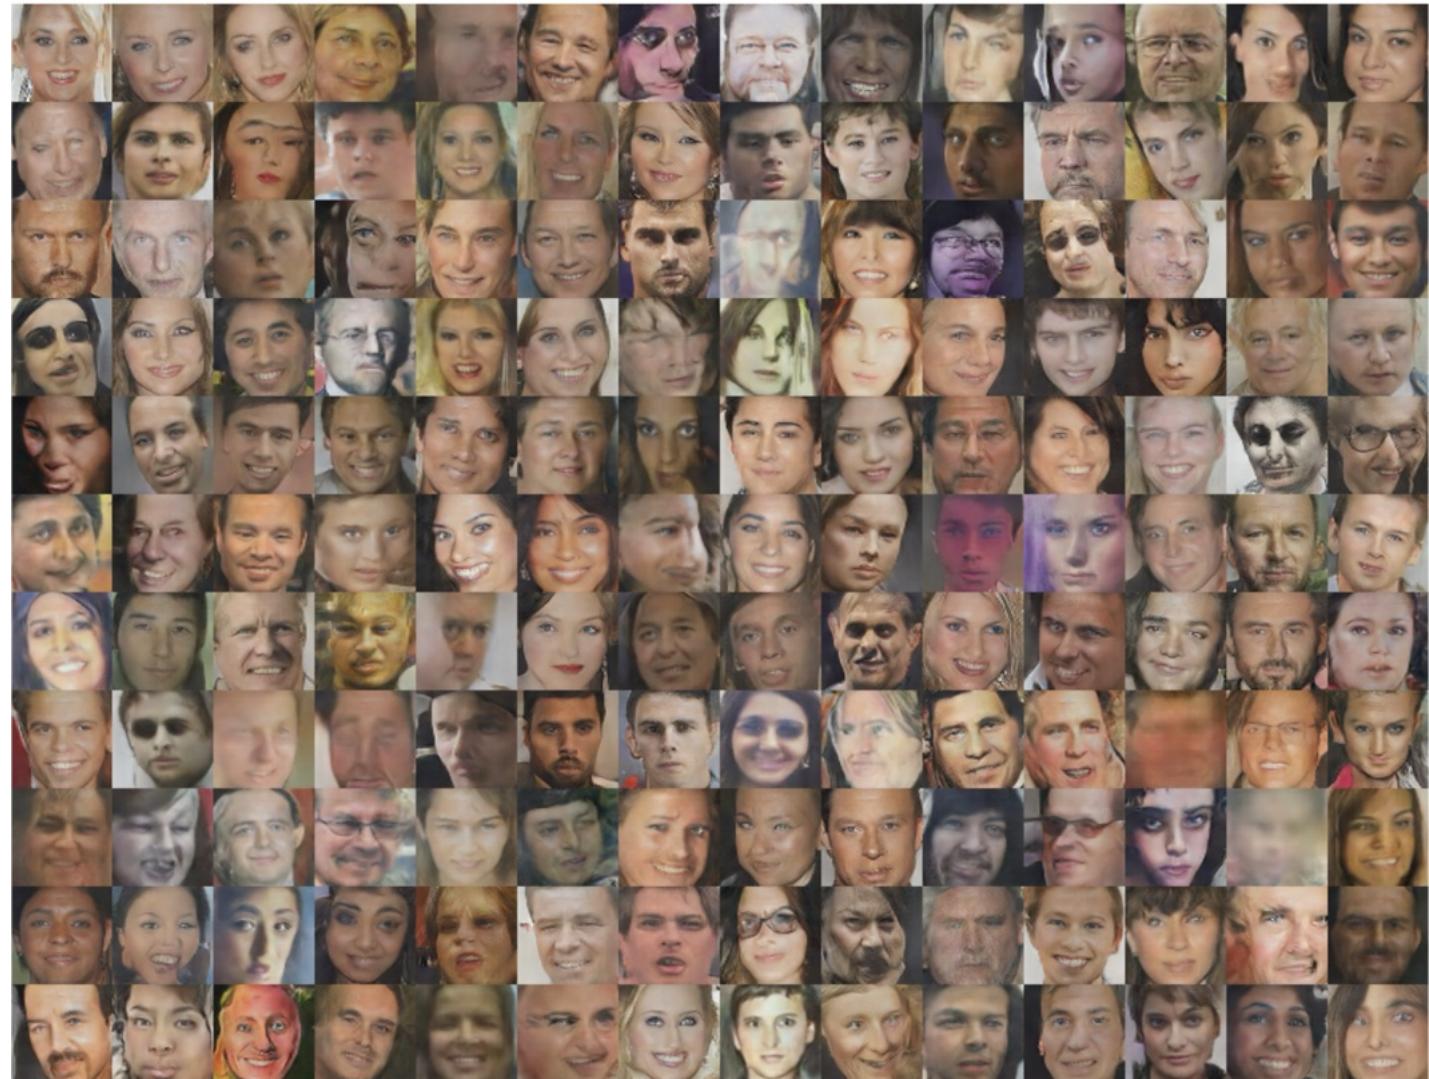
- Remove max-pooling and mean-pooling
- Upsample using transposed convolutions in the generator
- Downsample with strided convolutions and average pooling

# DCGAN - Key Results

- Good samples on datasets with 3M images (Faces, Bedrooms) for the first time

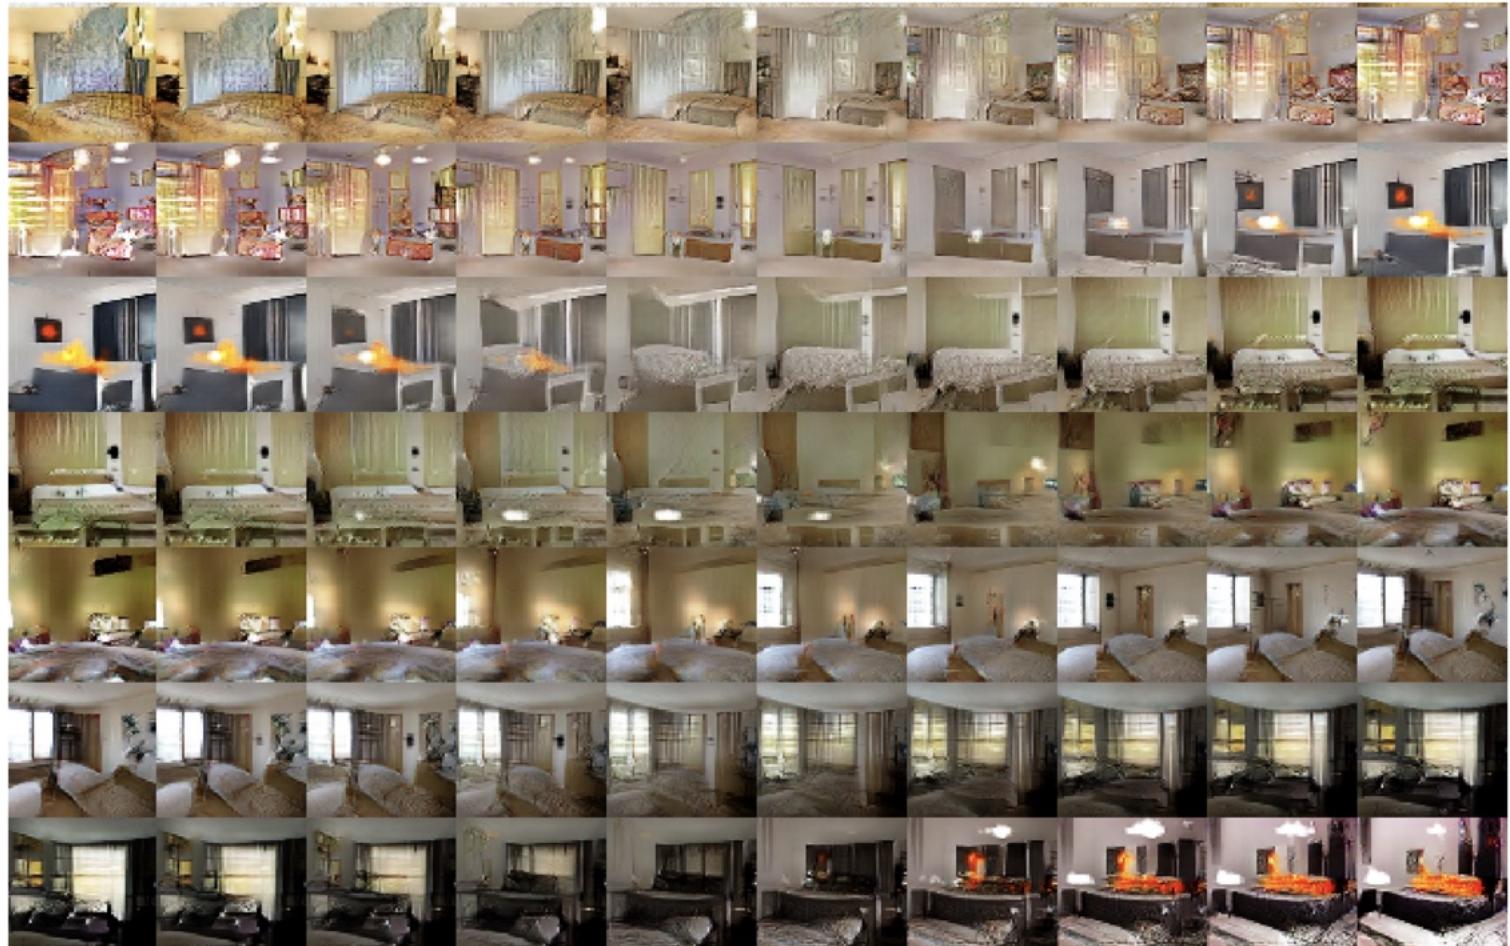


# DCGAN - Key Results



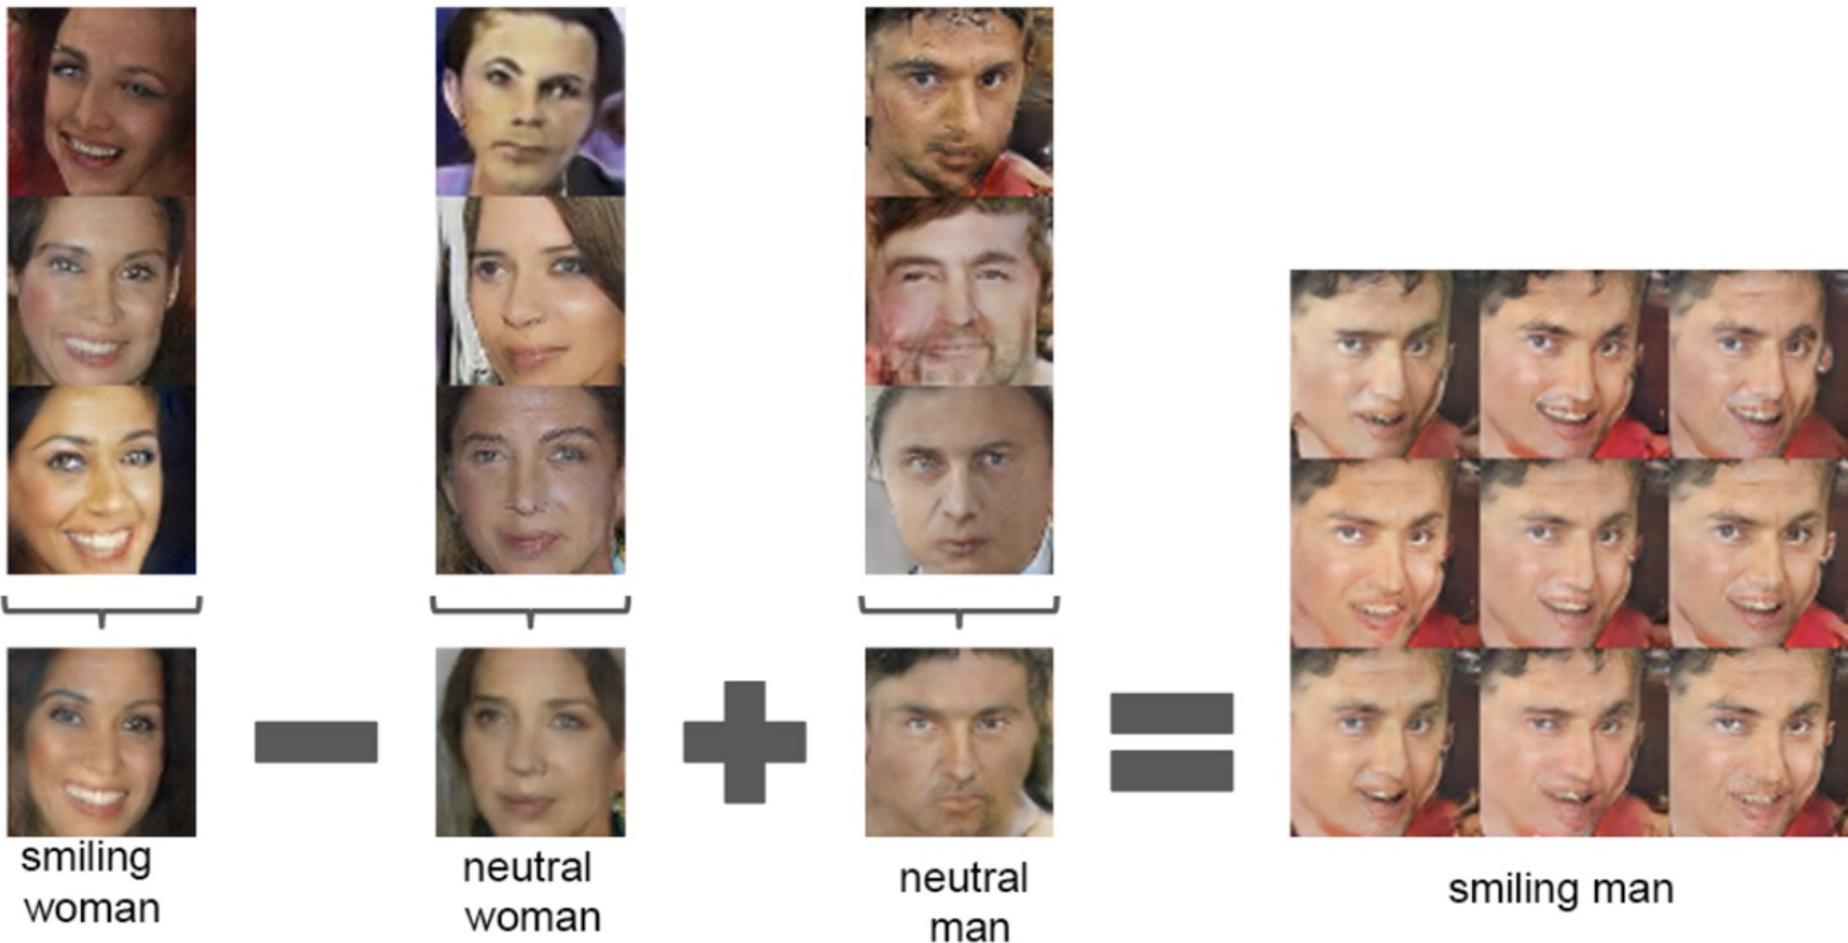
# DCGAN - Key Results

- Smooth interpolations in high dimensions

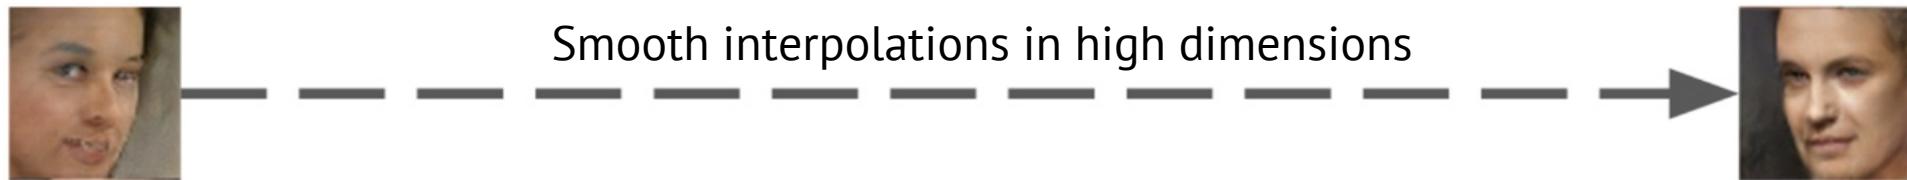


# DCGAN - Key Results

- Vector Arithmetic in Latent Space (z)



# DCGAN - Key Results



# GANs: Progress

## GAN PROGRESS ON FACE GENERATION

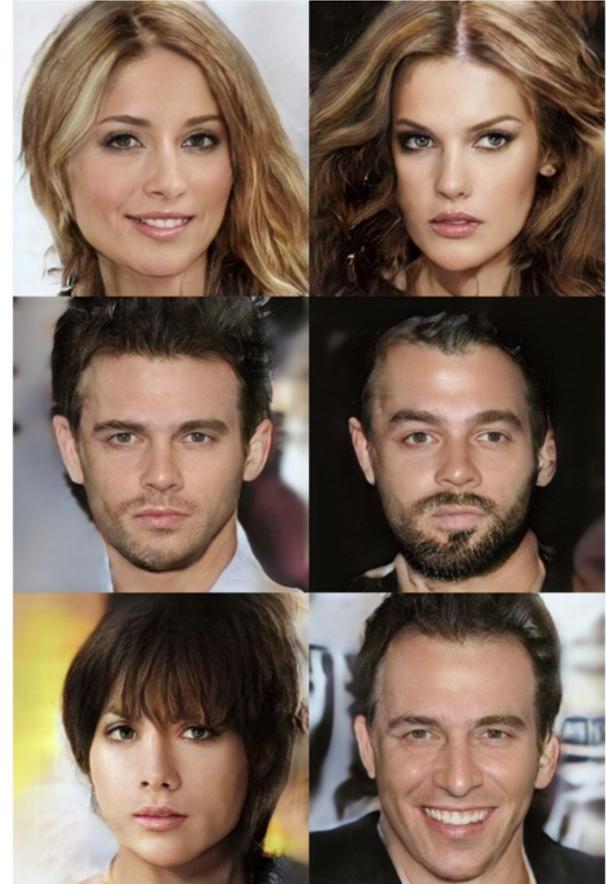
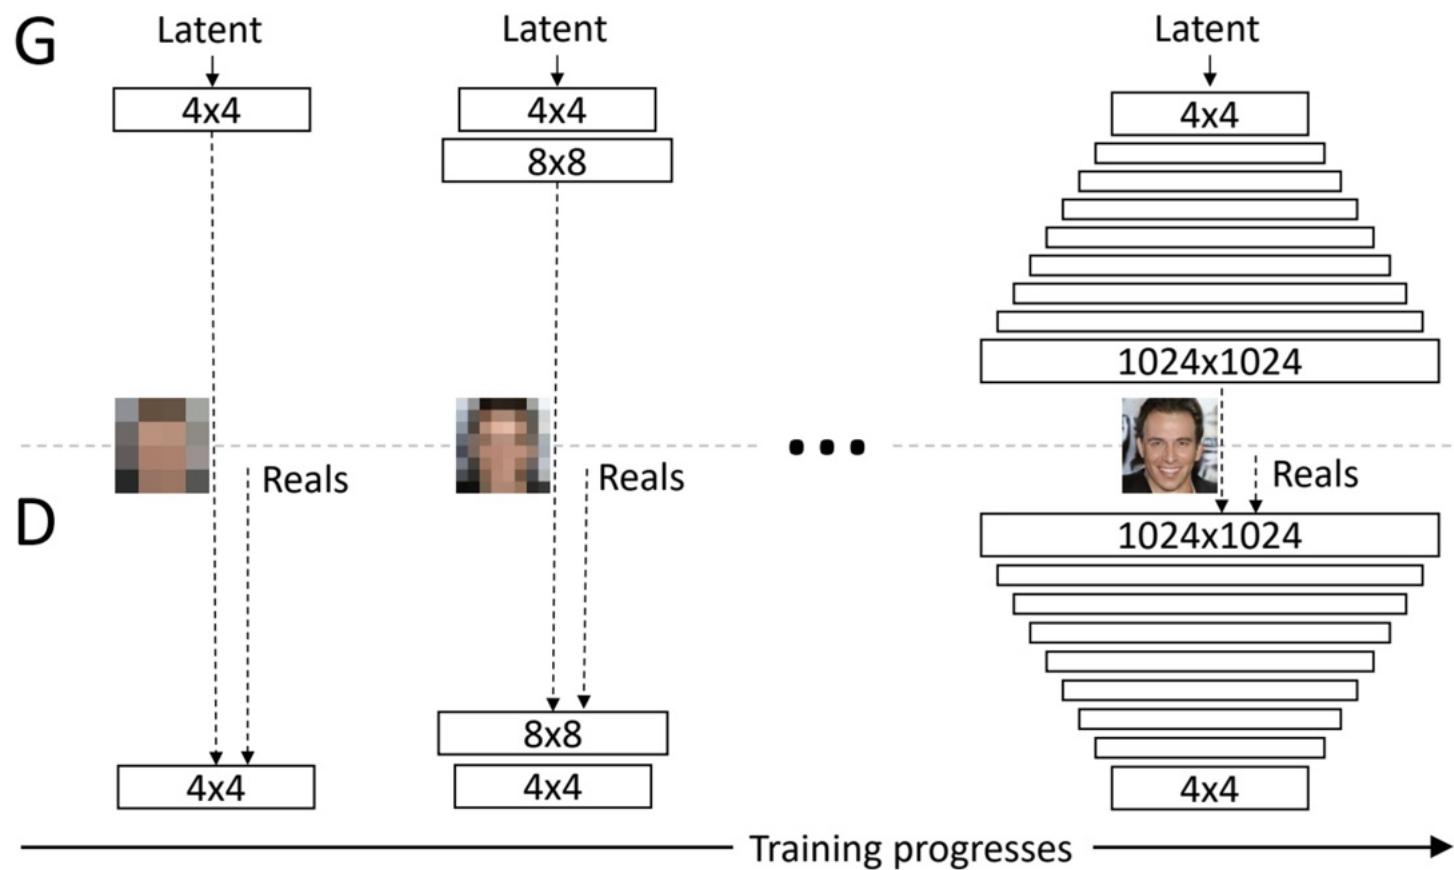
Source: Goodfellow et al., 2014; Radford et al., 2016; Liu & Tuzel, 2016; Karras et al., 2018; Karras et al., 2019; Goodfellow, 2019; Karras et al., 2020; AI Index, 2021



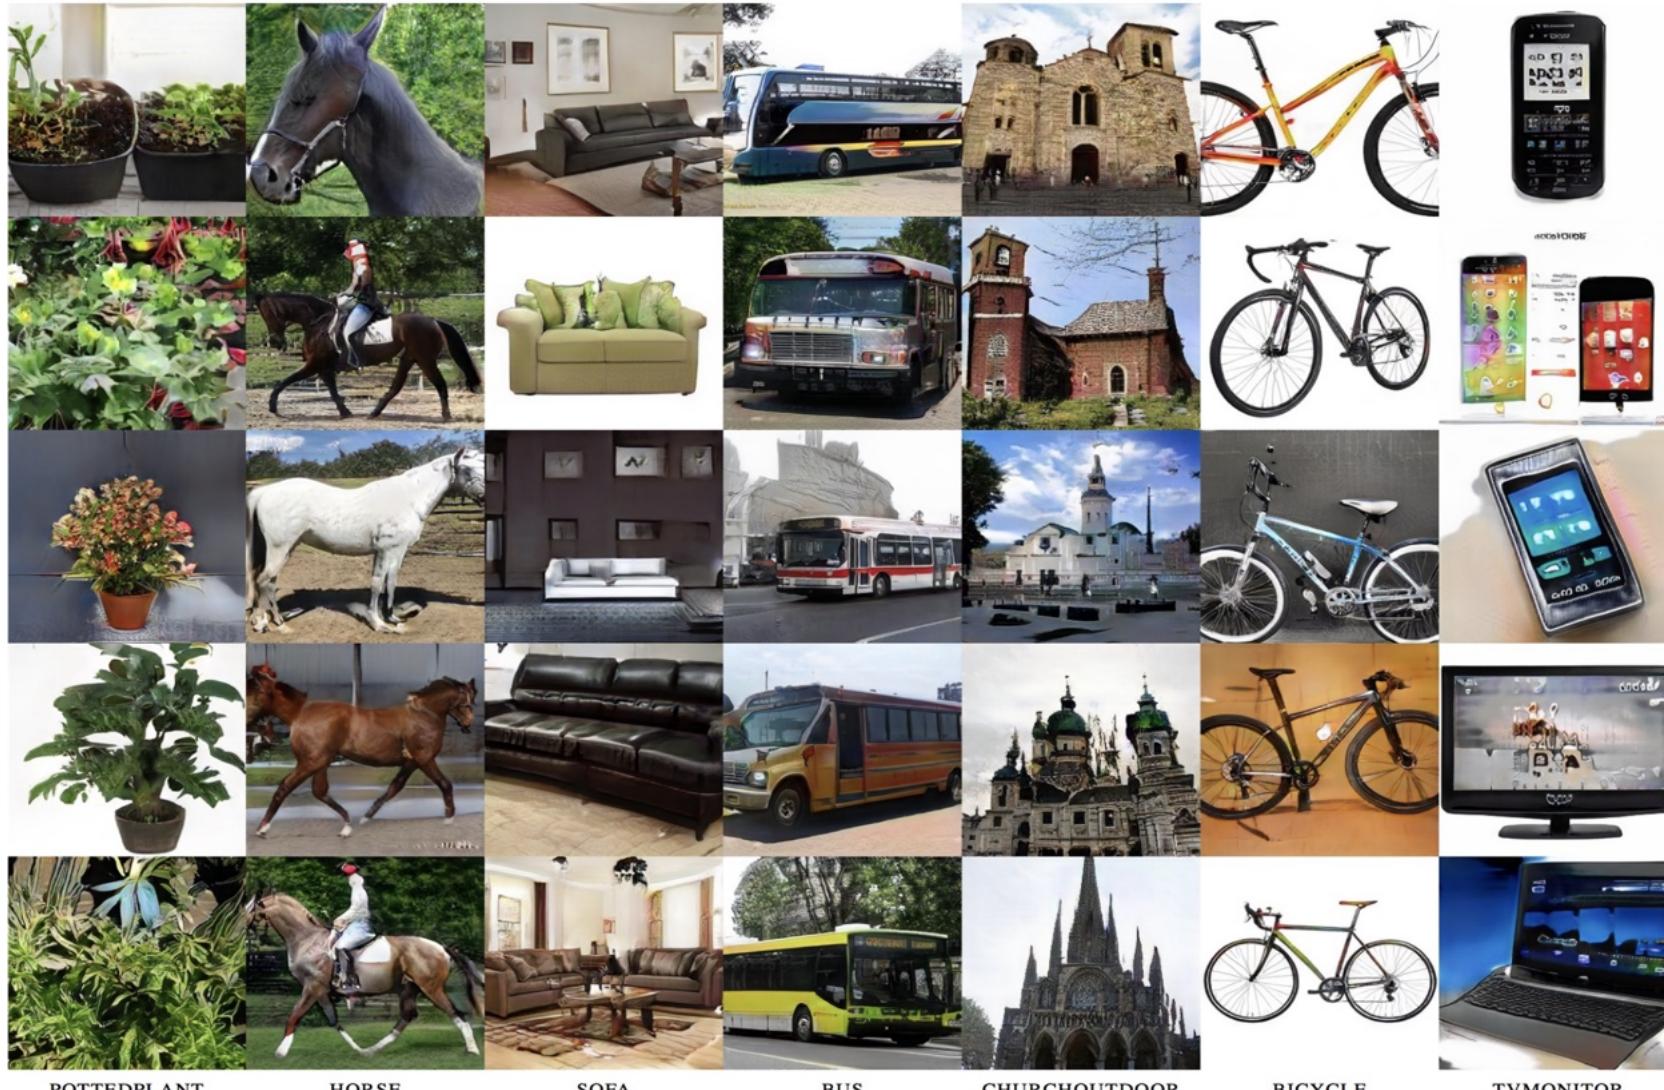
# Progressive growing of GANs



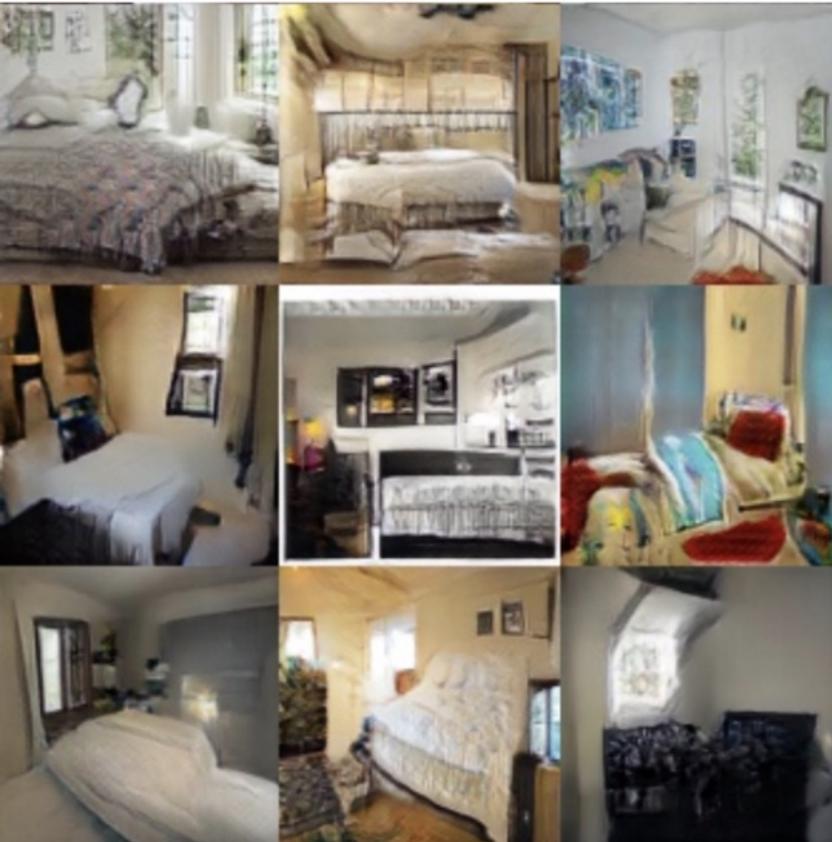
First high quality samples from any generative models



# Progressive growing of GANs



# Progressive growing of GANs



Mao et al. (2016b) (128 × 128)



Gulrajani et al. (2017) (128 × 128)



Our (256 × 256)

Karras et al  
2017

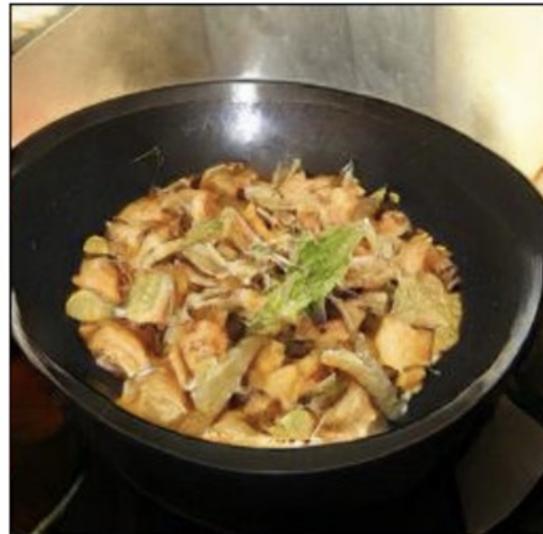
# Progressive growing of GANs



# Progressive growing of GANs



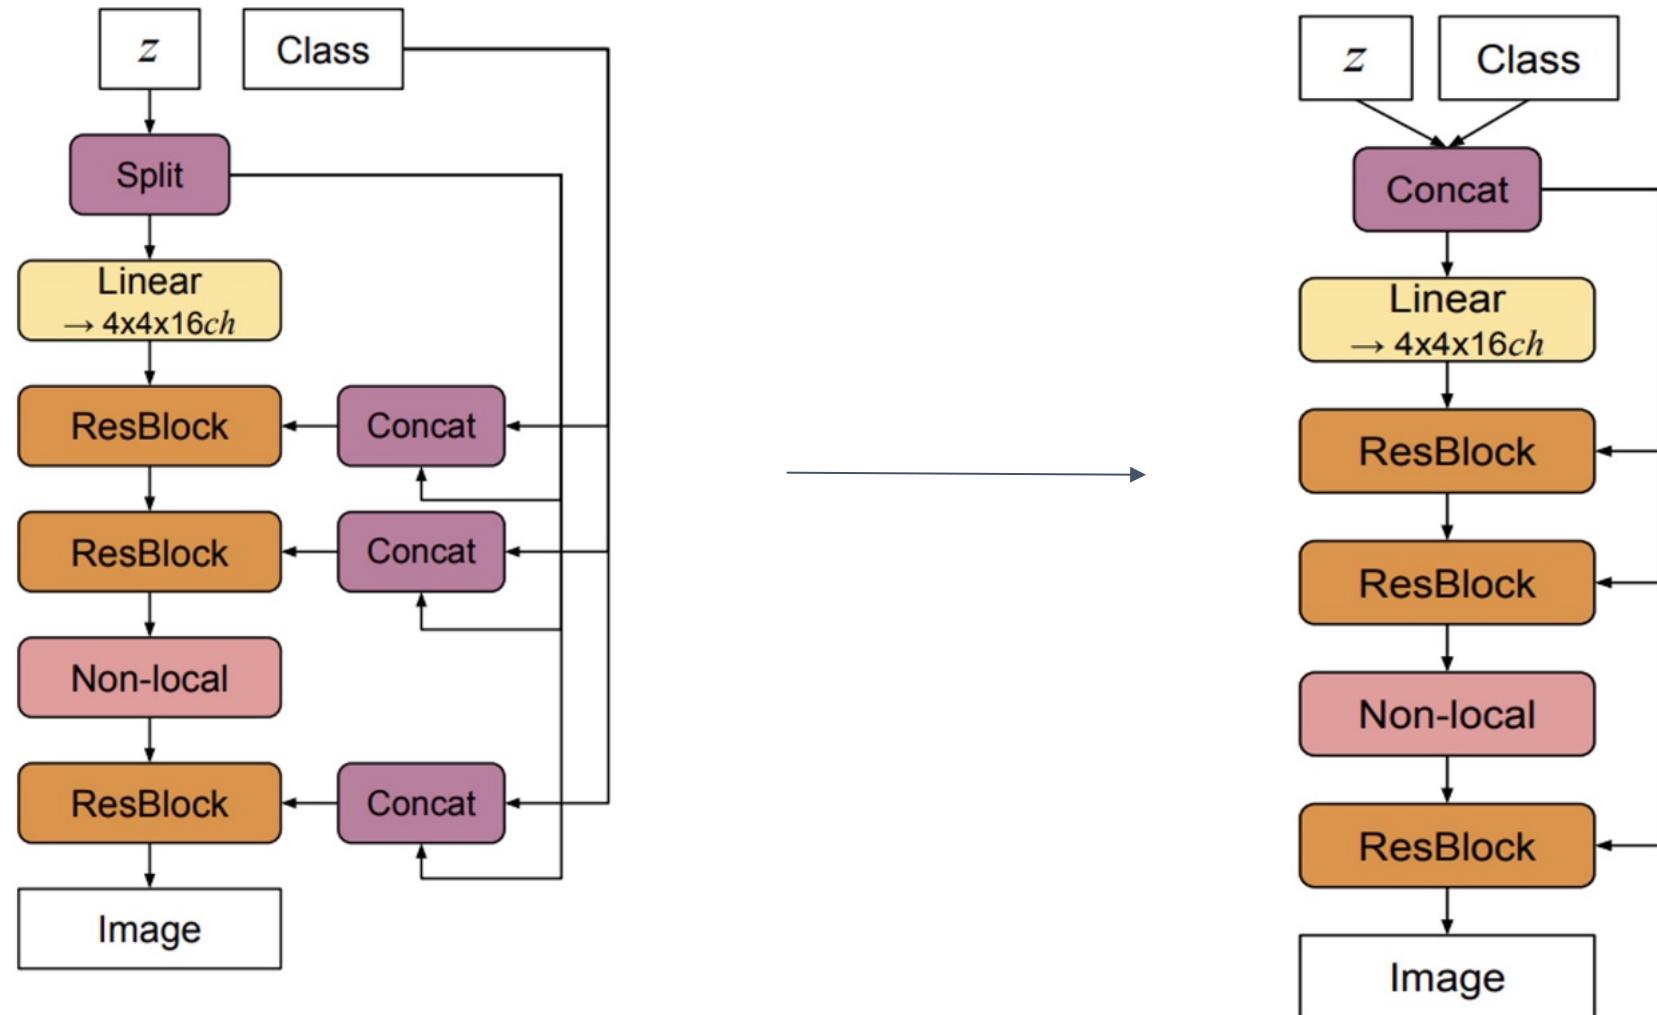
# BigGAN



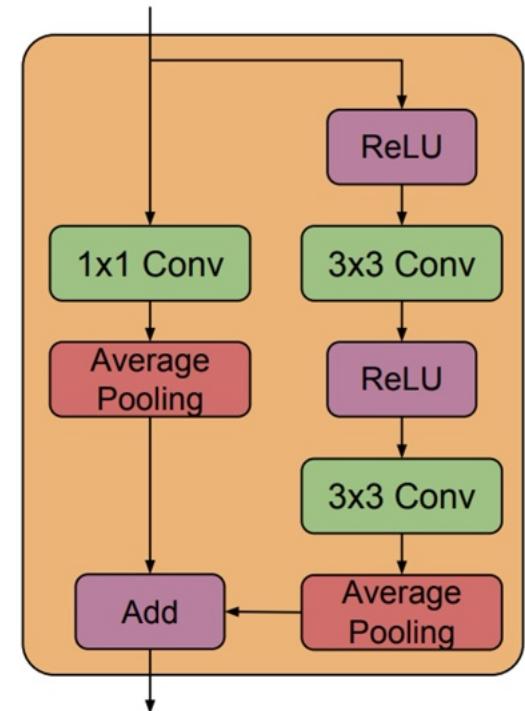
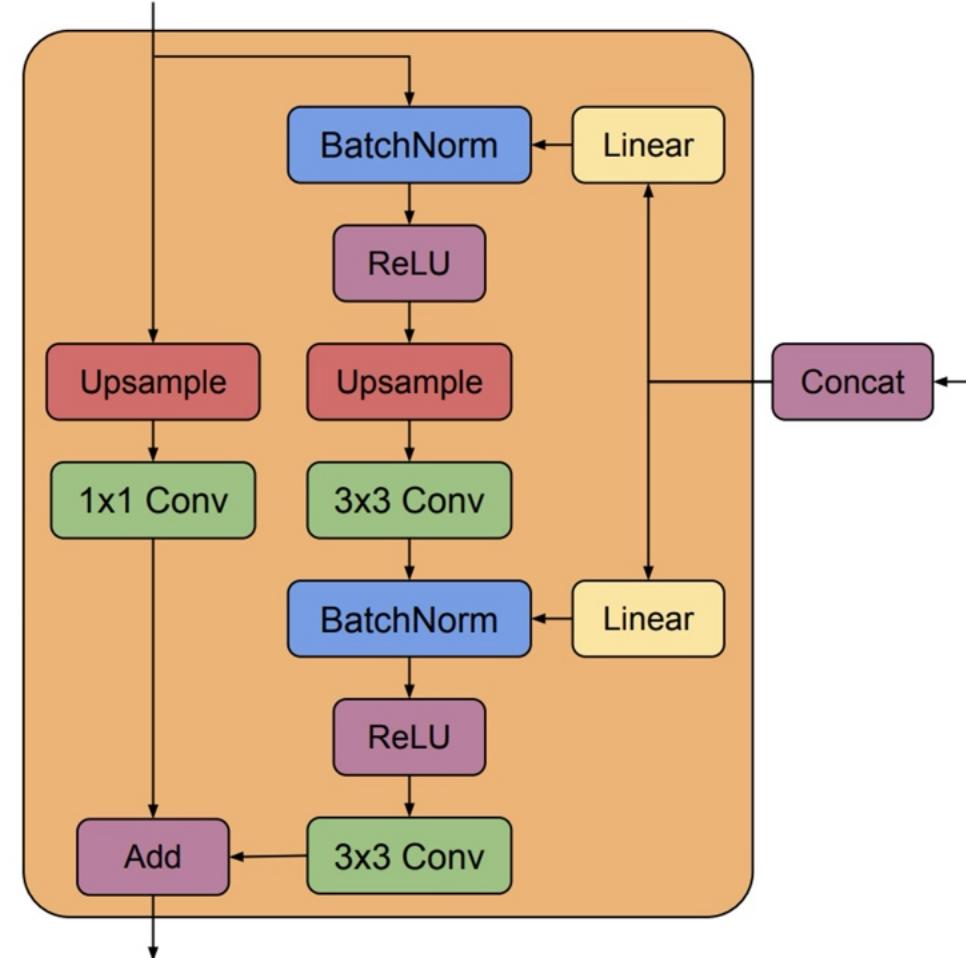
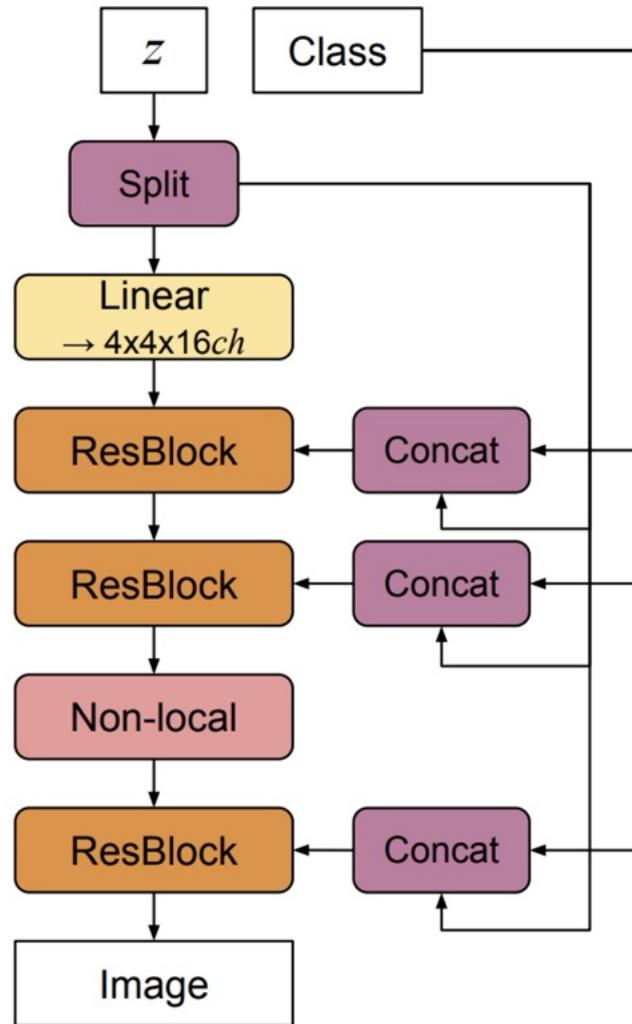
# BigGAN



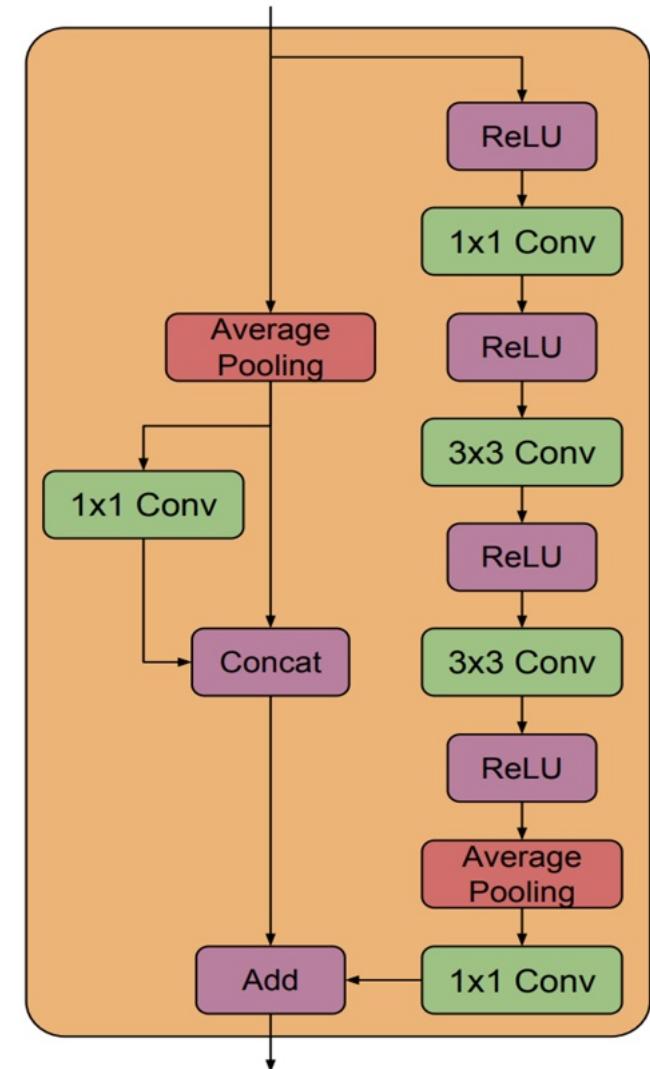
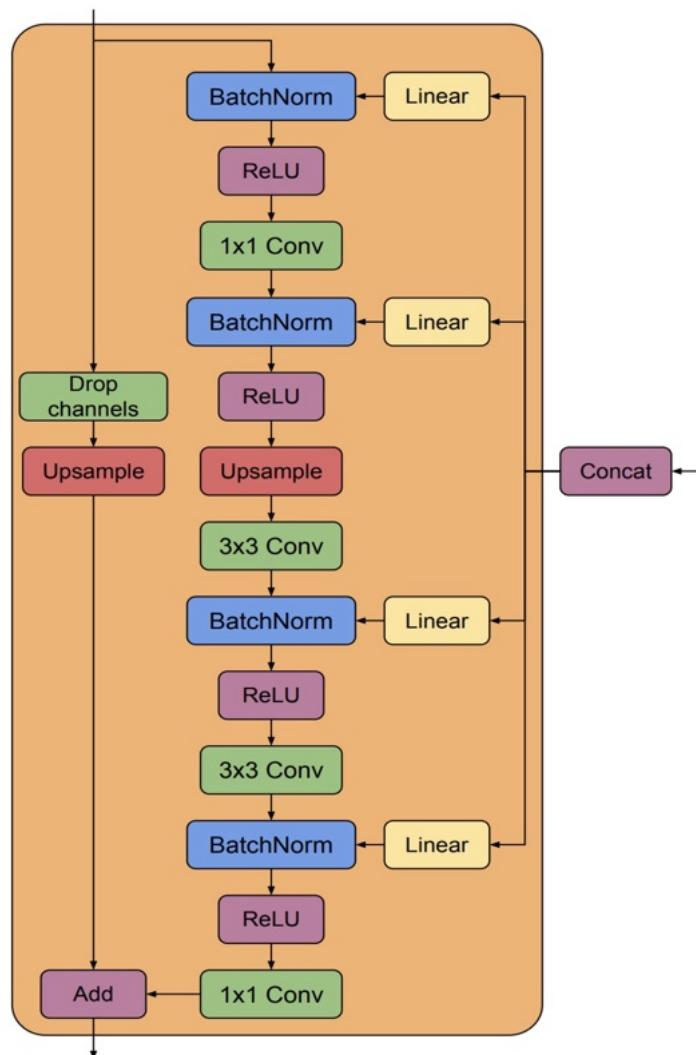
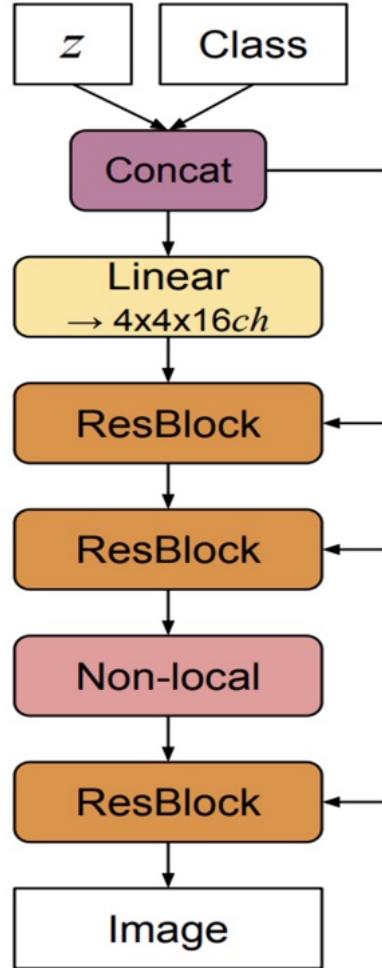
# BigGAN and BigGAN-deep



# BigGAN



# BigGAN-deep





# BigGAN

- Salient bits
  - Increase your batch size (as much as you can)
  - Use Cross-Replica (Sync) Batch Norm
  - Increase your model size
  - Wider helps as much as deeper
  - Fuse class information at all levels
  - Hinge Loss
  - Orthonormal regularization & Truncation Trick (on  $z$ )

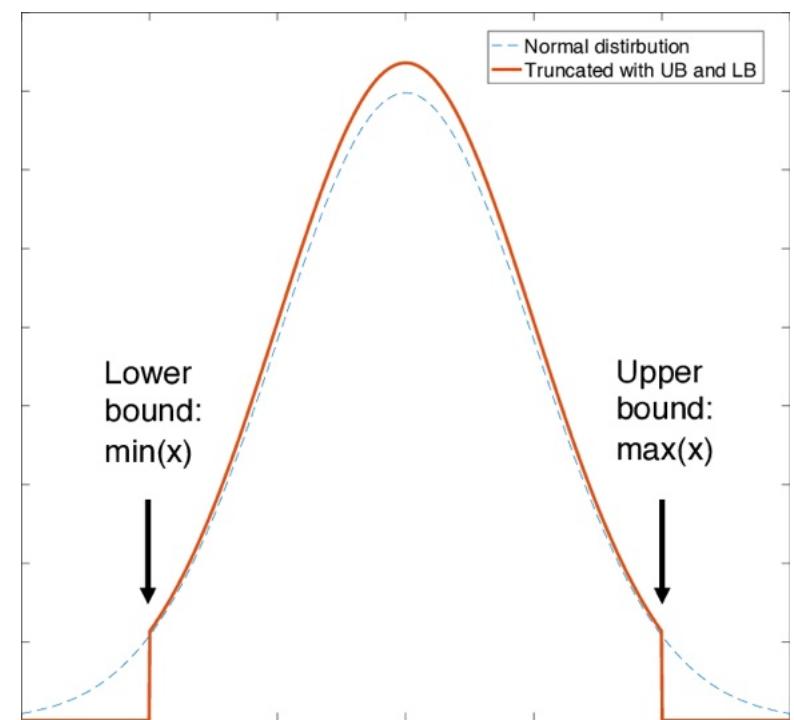
# BigGAN - Truncation Trick



(a)



(b)



# BigGAN



(a)  $128 \times 128$



(b)  $256 \times 256$

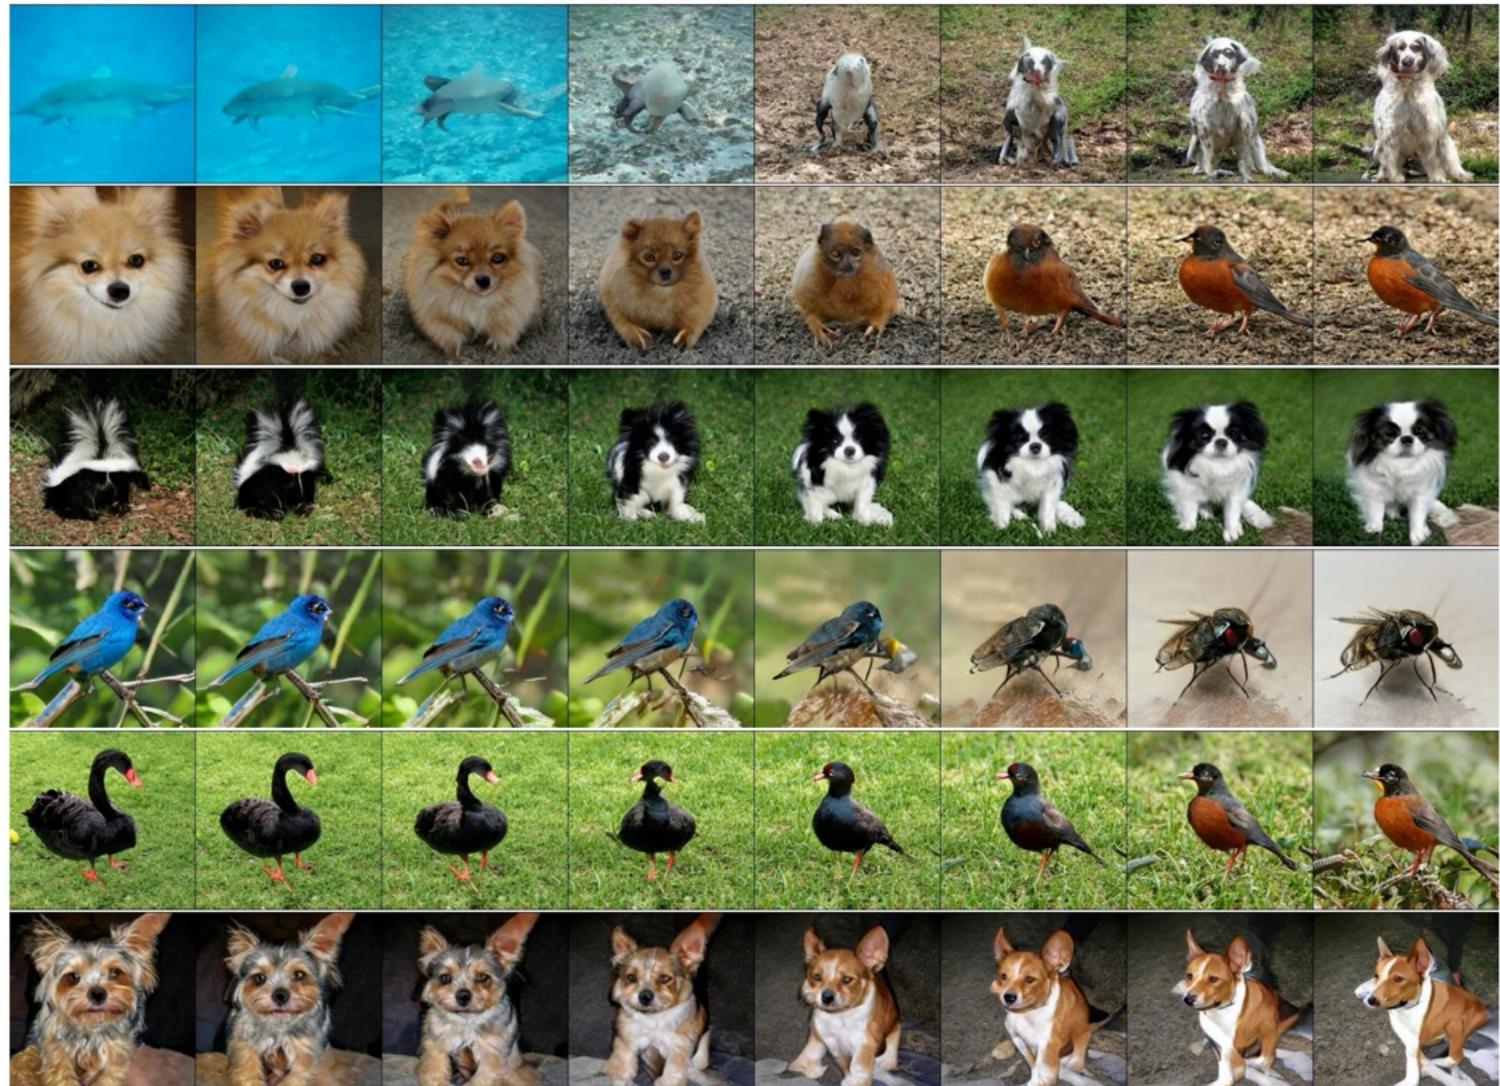


(c)  $512 \times 512$

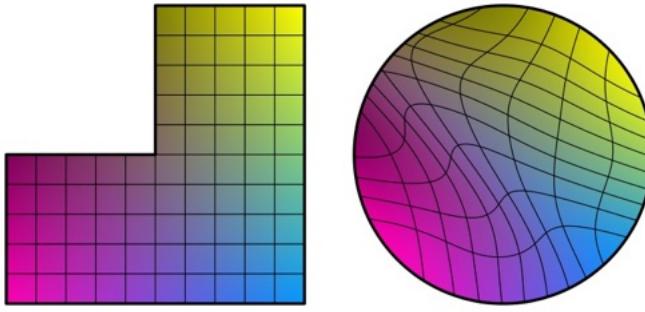
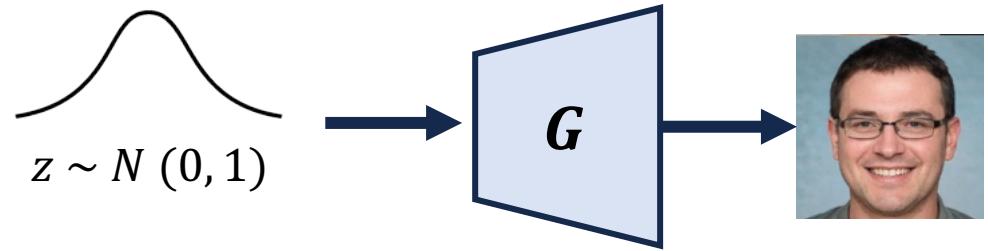


(d)

# BigGAN



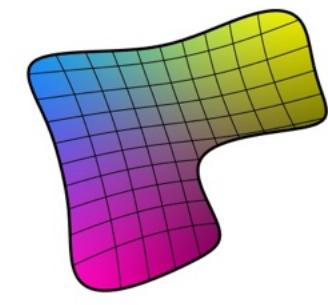
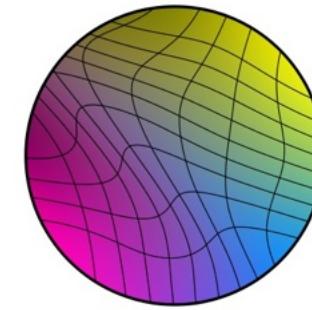
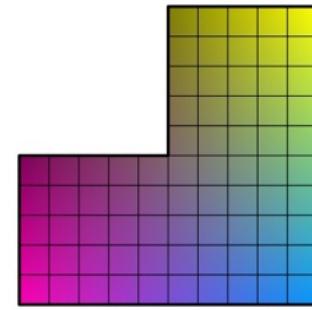
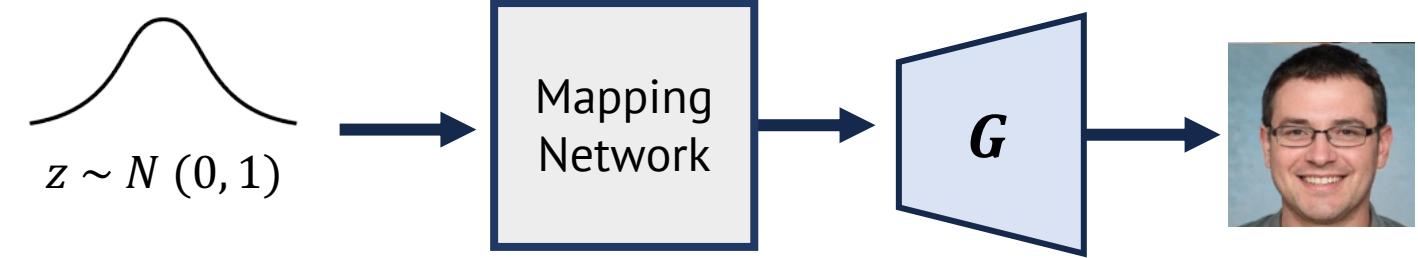
# StyleGAN



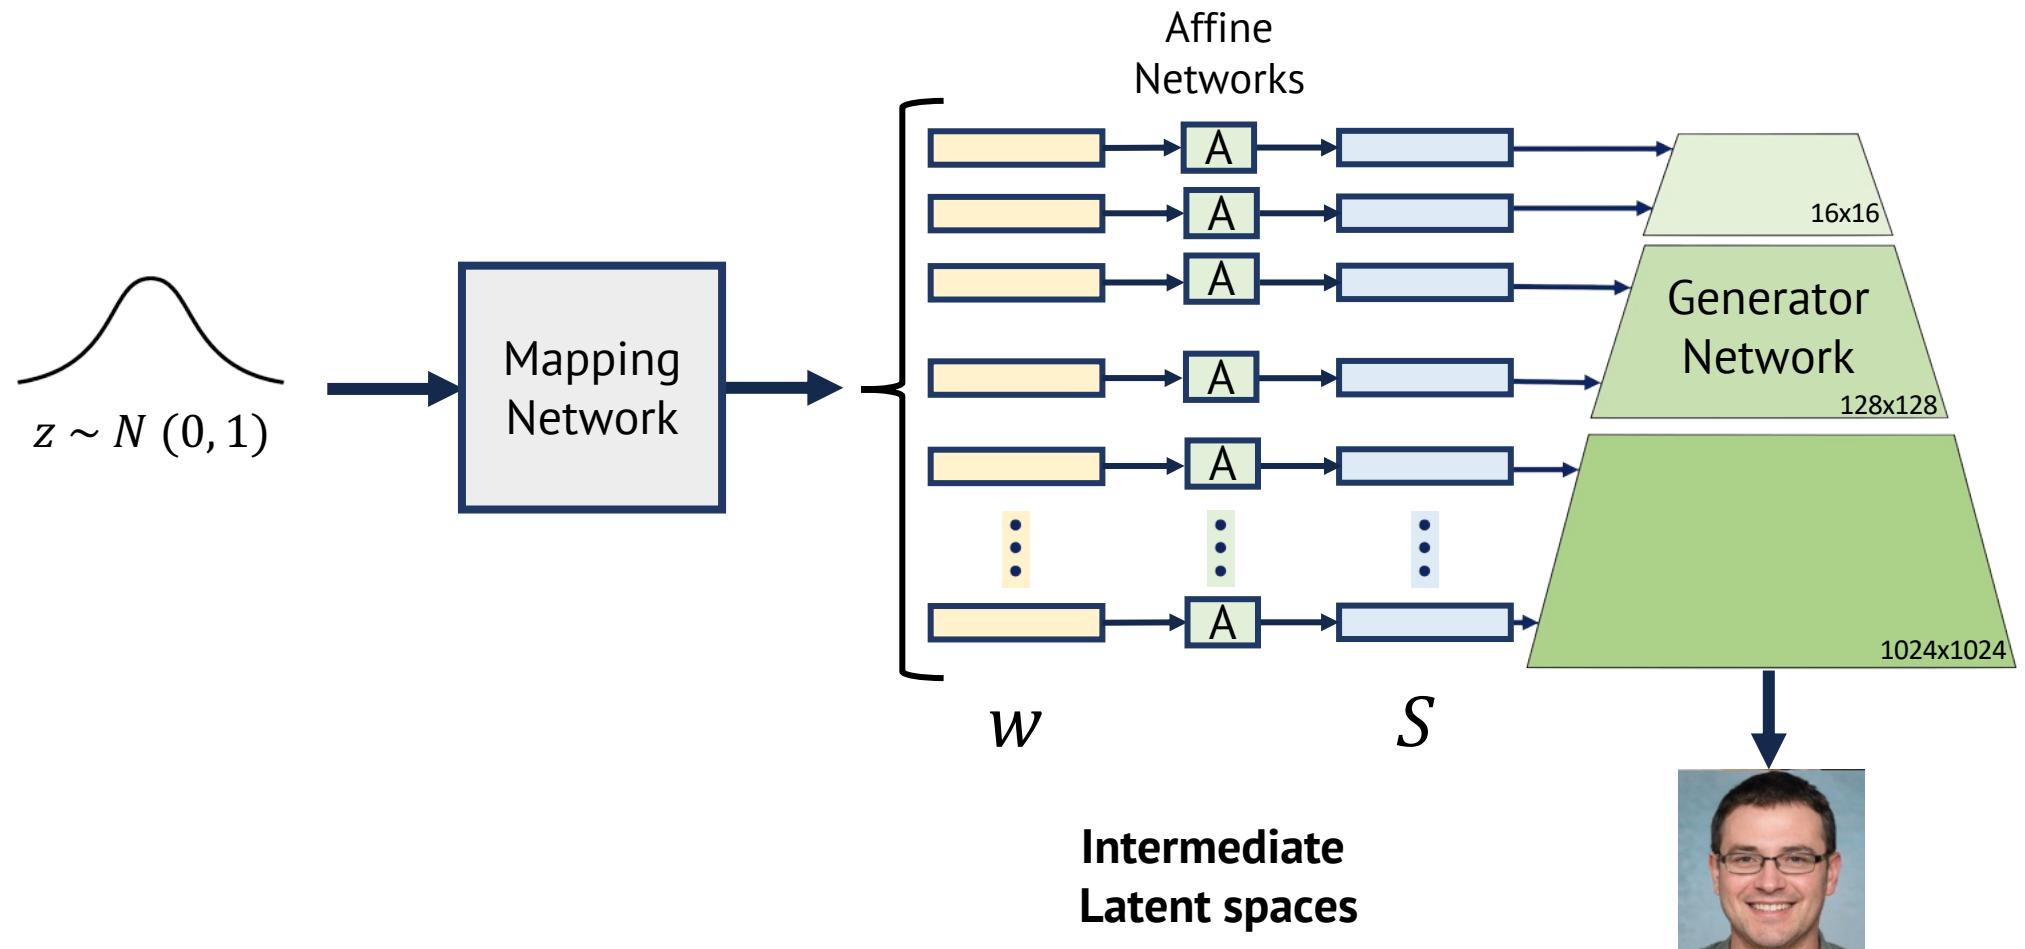
(a) Distribution of features in training set

(b) Mapping from  $\mathcal{Z}$  to features

# StyleGAN

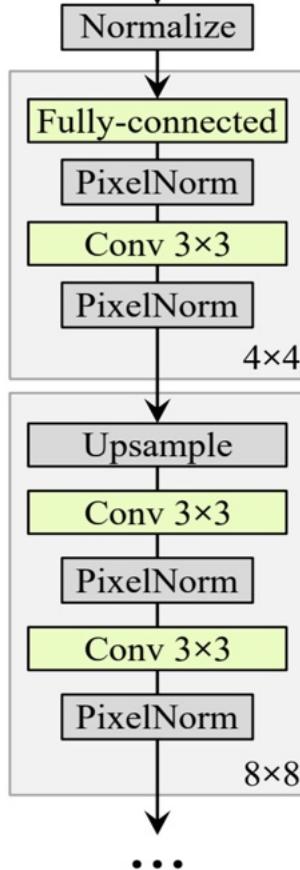


# StyleGAN



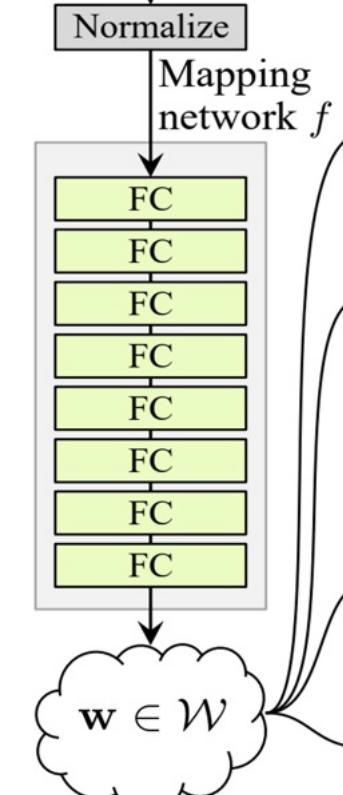
# StyleGAN

Latent  $z \in \mathcal{Z}$



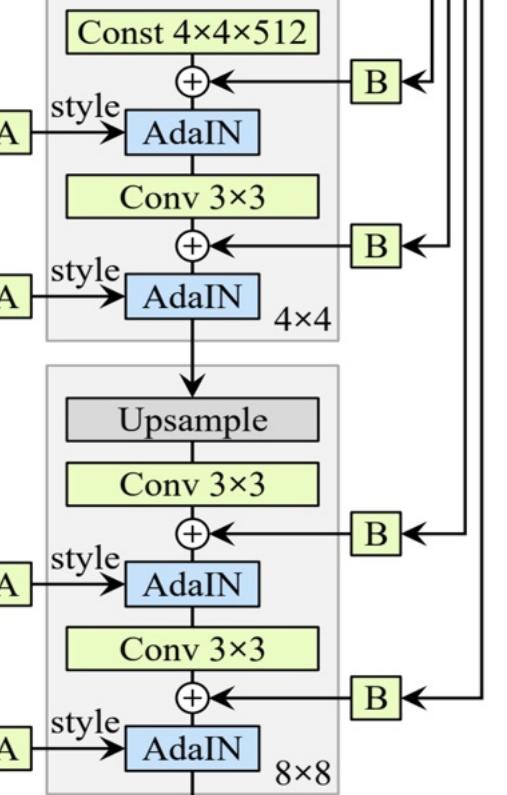
(a) Traditional

Latent  $z \in \mathcal{Z}$



(b) Style-based generator

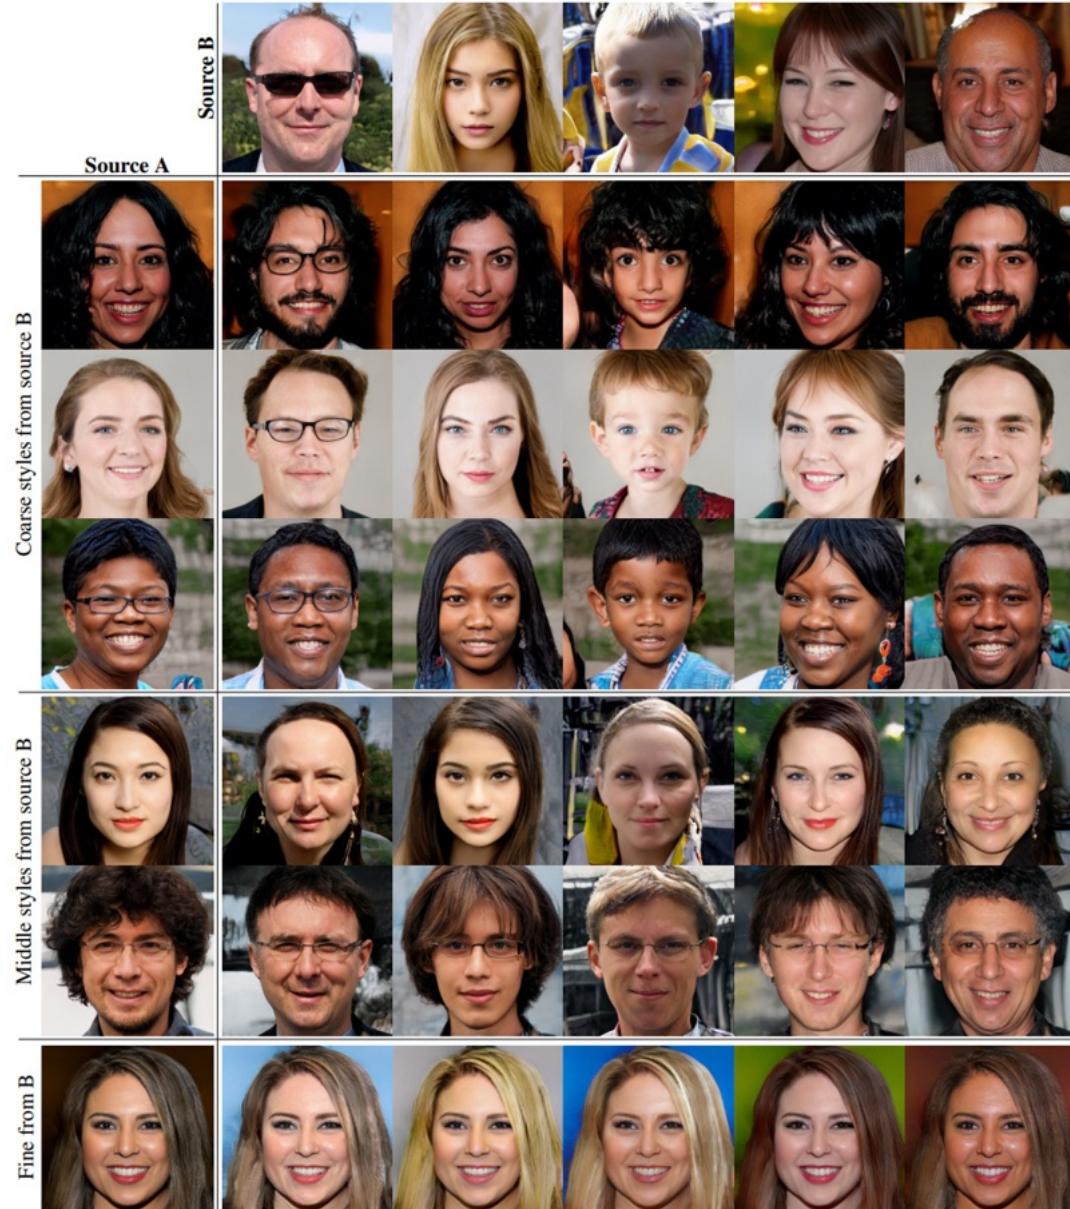
Latent  $z \in \mathcal{Z}$



$$\text{AdaIN}(\mathbf{x}_i, \mathbf{y}) = \mathbf{y}_{s,i} \frac{\mathbf{x}_i - \mu(\mathbf{x}_i)}{\sigma(\mathbf{x}_i)} + \mathbf{y}_{b,i}$$

Instance normalization =  
normalize per channel per sample

# StyleGAN – “Style” Transfer

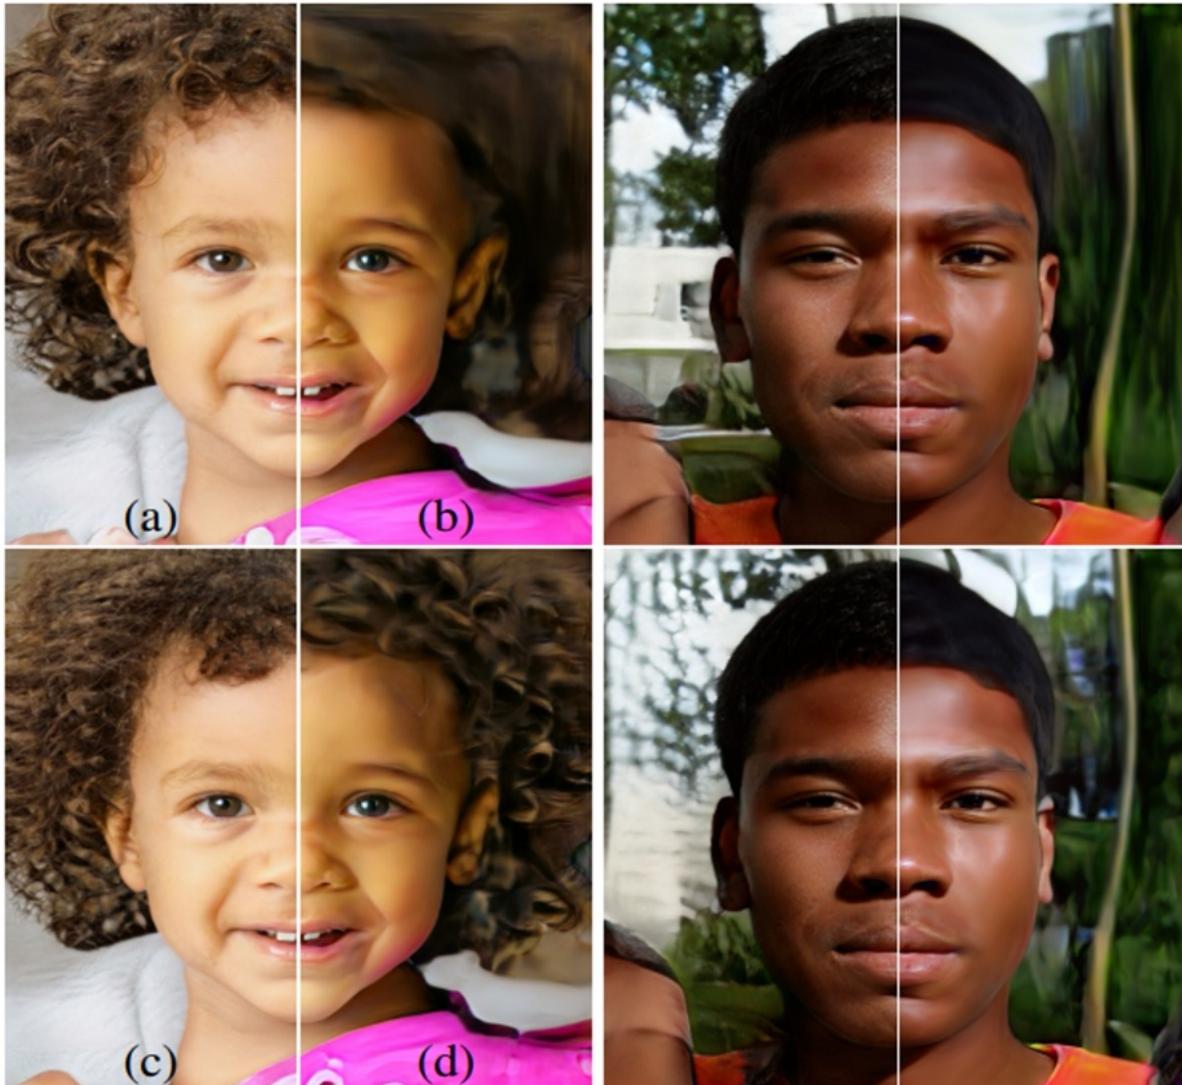


Ref: StyleGAN, Karras et al. CVPR '19;  
StyleGAN2, Karras et al. CVPR '20

# StyleGAN



# StyleGAN - Effect of adding noise

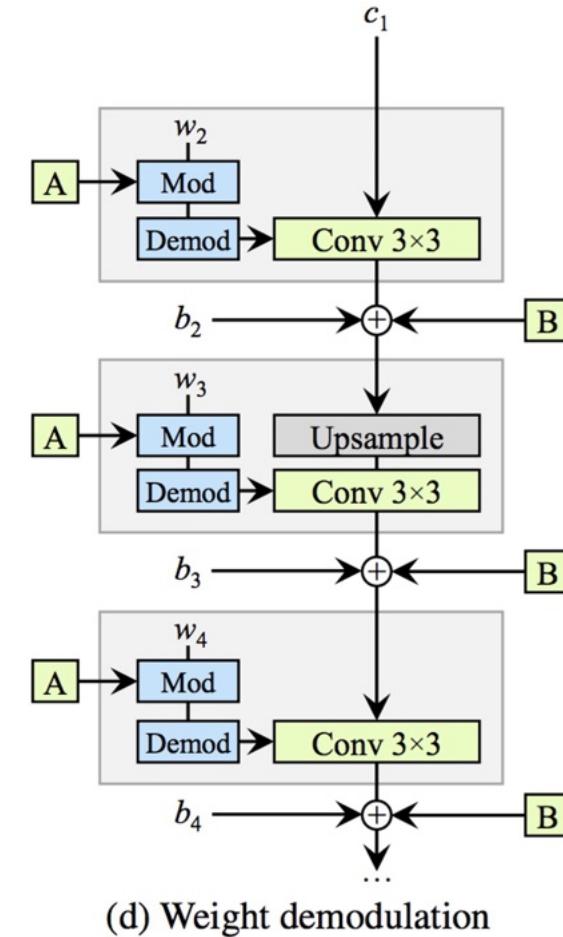
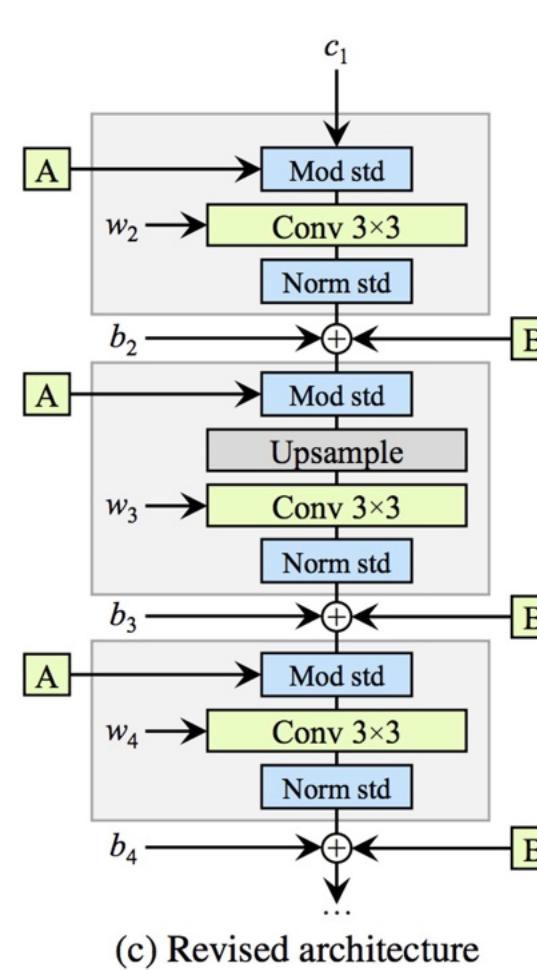
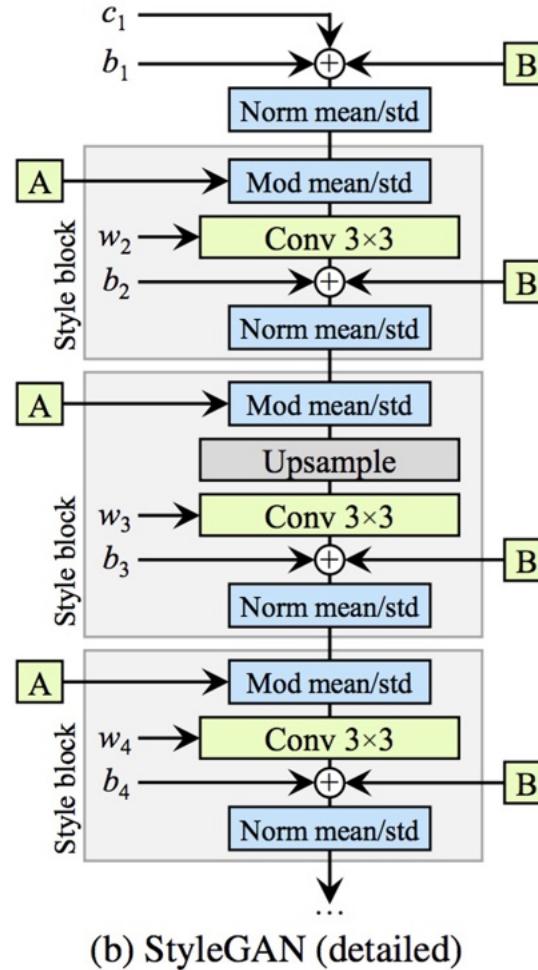
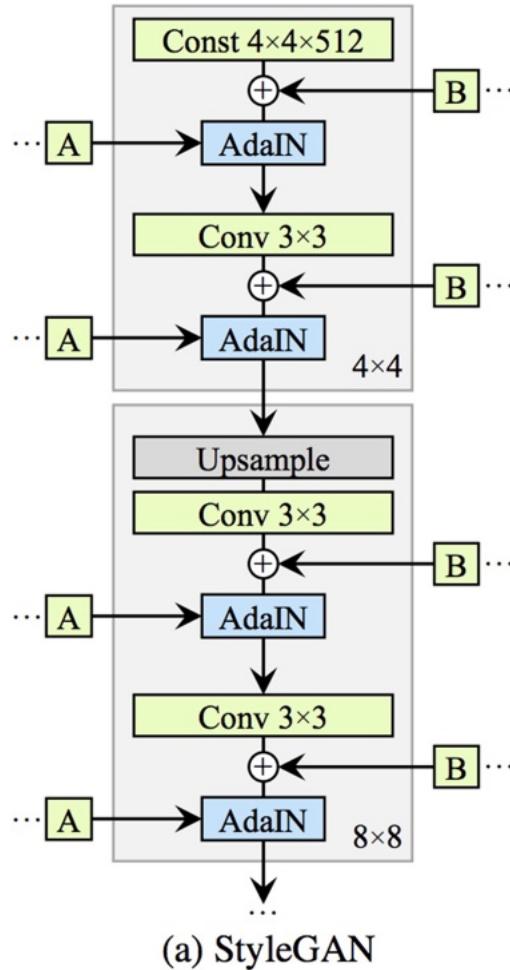


# StyleGAN-v2



Figure 1. Instance normalization causes water droplet -like artifacts in StyleGAN images. These are not always obvious in the generated images, but if we look at the activations inside the generator network, the problem is always there, in all feature maps starting from the 64x64 resolution. It is a systemic problem that plagues all StyleGAN images.

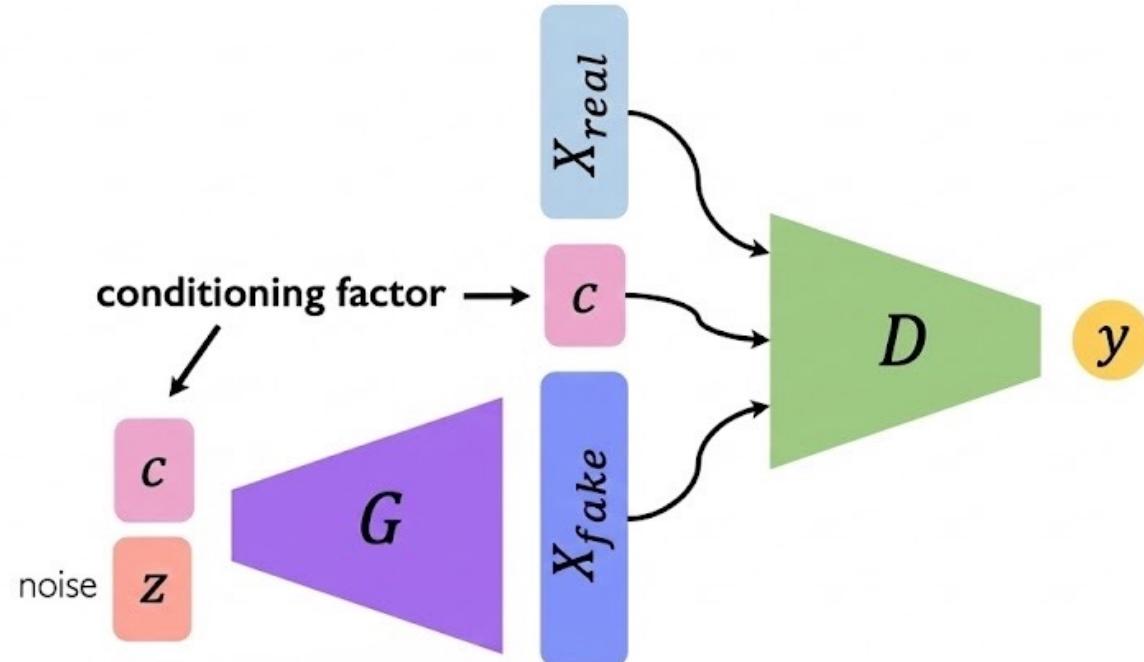
# StyleGAN-v2



# Conditional Generation using GANs

## Conditional GANs

What if we want to control the nature of the output, by **conditioning** on a label?



# Pix2Pix: Conditional Image-to-Image Generation

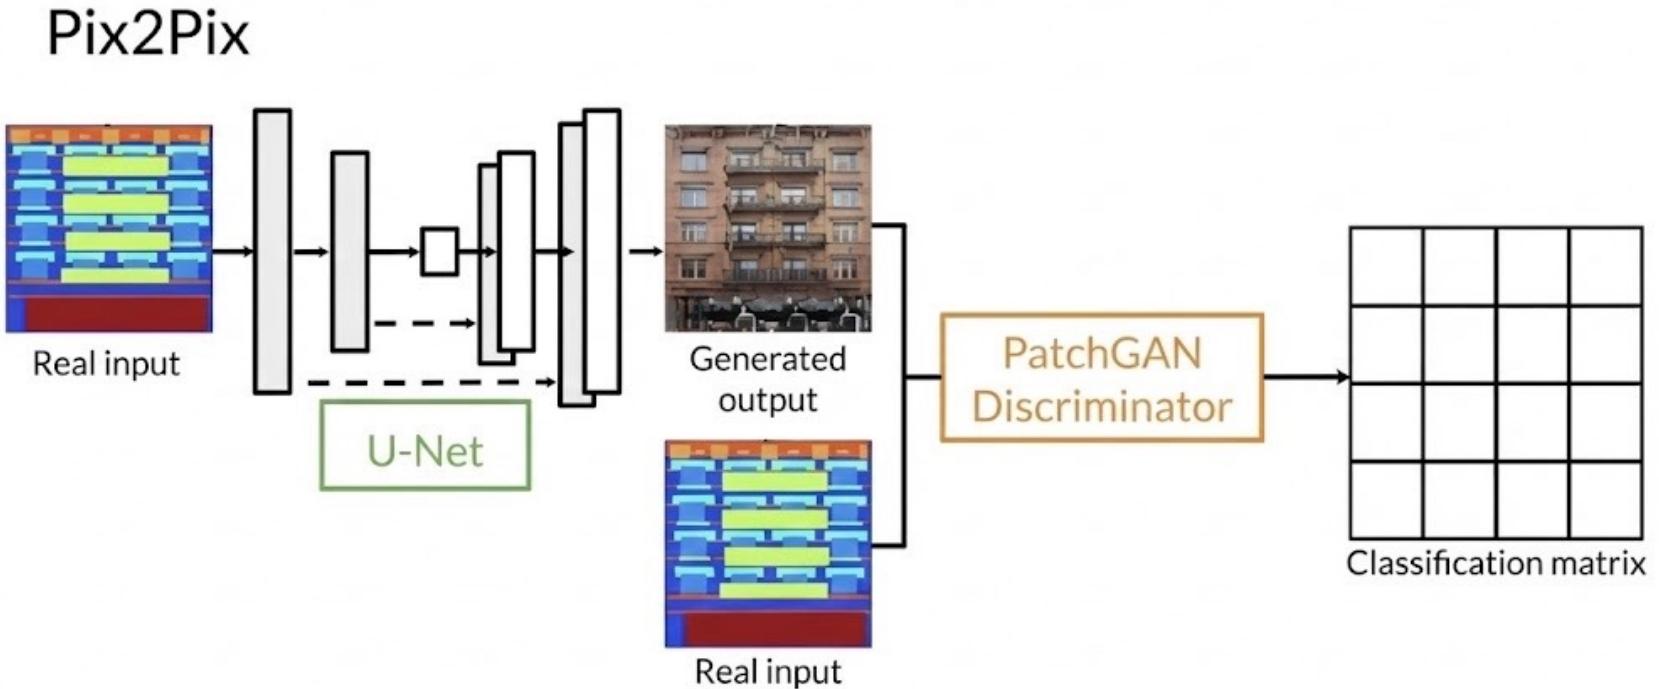
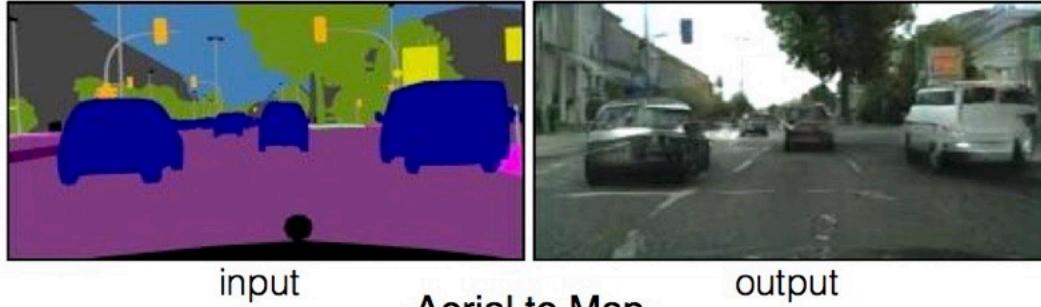


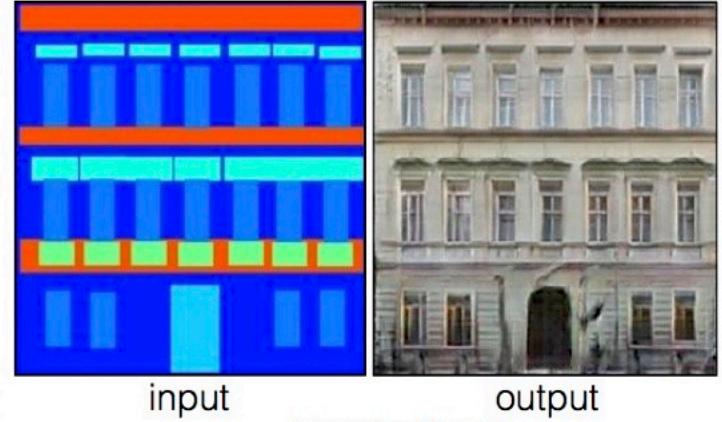
Image available from: <https://arxiv.org/abs/1611.07004>

# Pix2Pix: Results

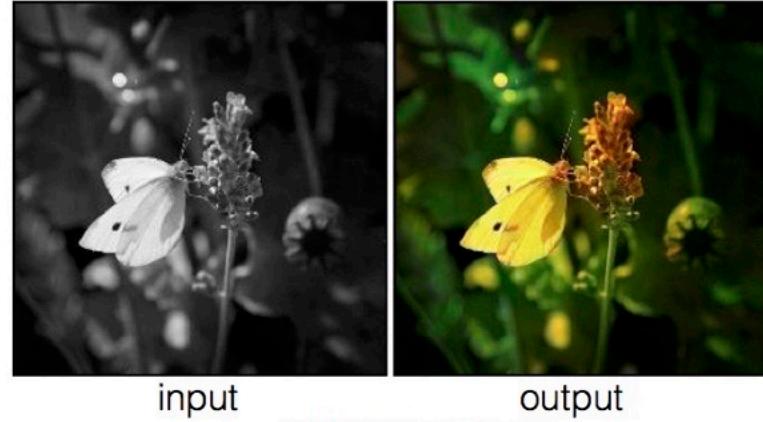
Labels to Street Scene



Labels to Facade



BW to Color



Aerial to Map



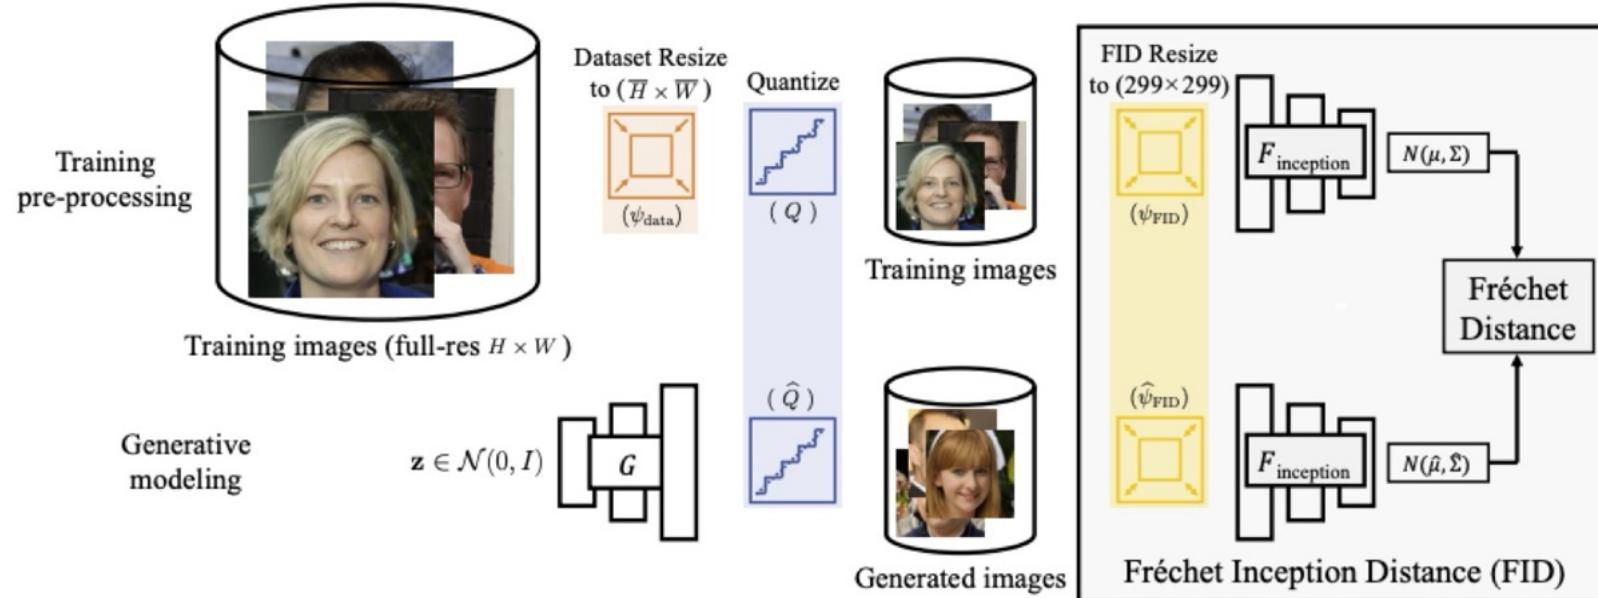
Day to Night



Edges to Photo



# How to evaluate performance of GAN?



Fréchet Inception Distance (FID)

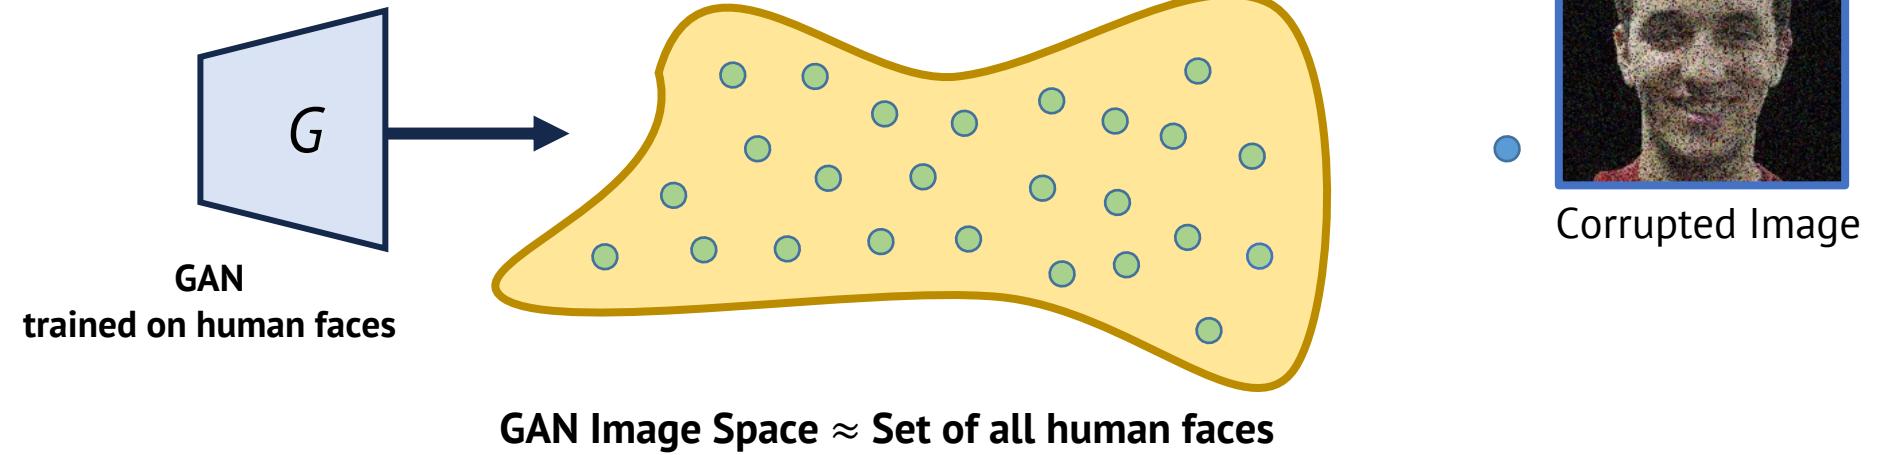
$$\text{FID} = \|\mu - \hat{\mu}\|_2^2 + \text{Tr}(\Sigma + \hat{\Sigma} - 2(\Sigma \hat{\Sigma})^{1/2})$$



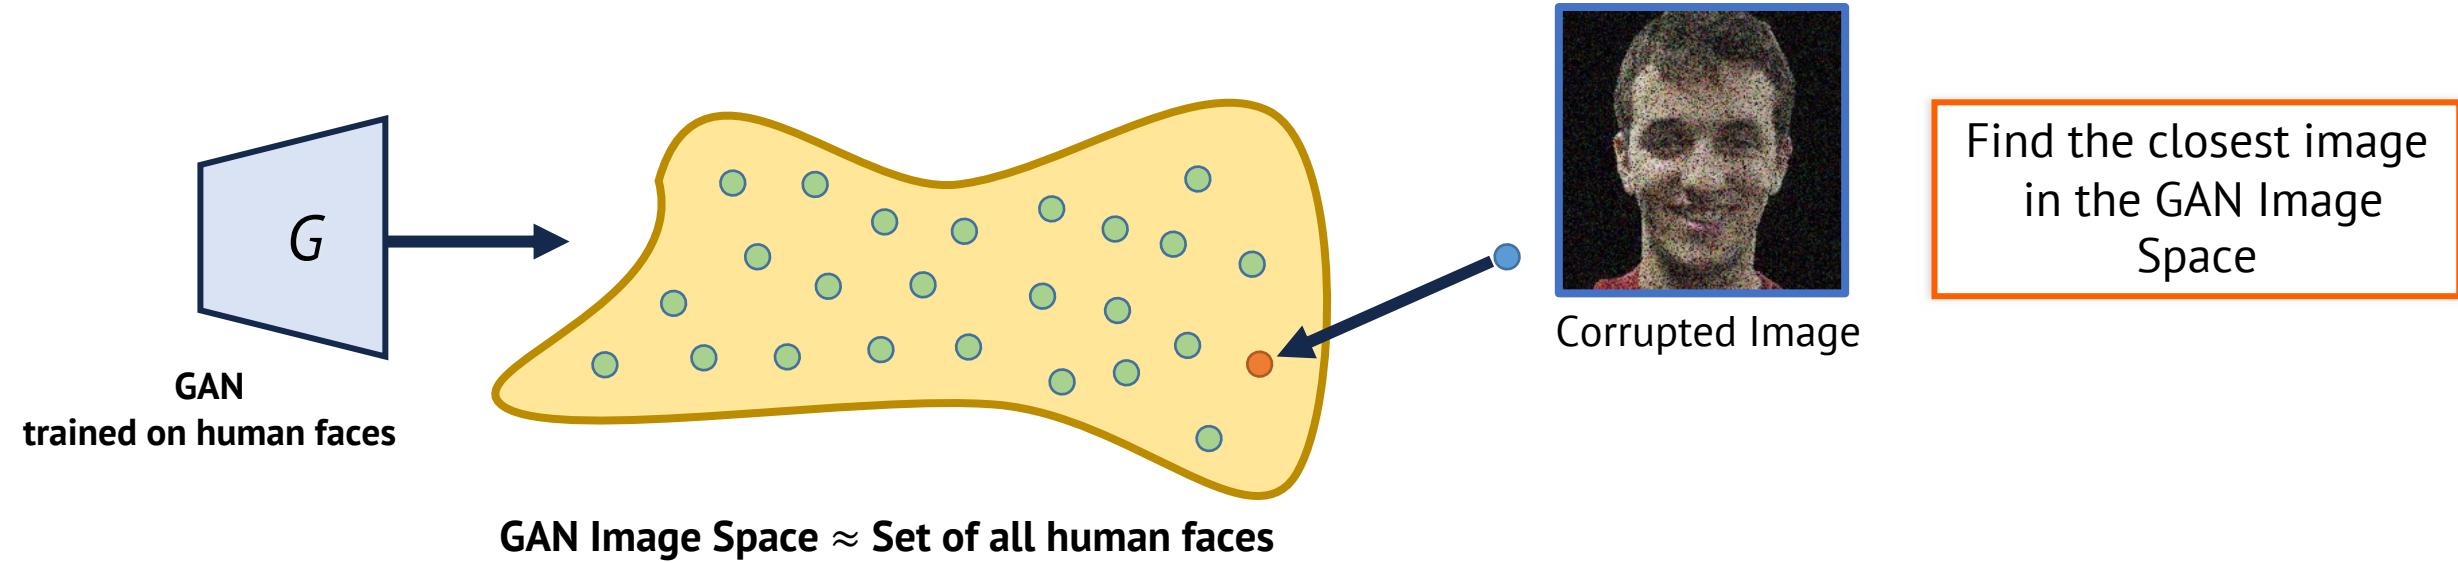
# GANs: Applications

- Rich representation of the image data
- Smooth latent space with interpolation capabilities
- Fast, single step inference

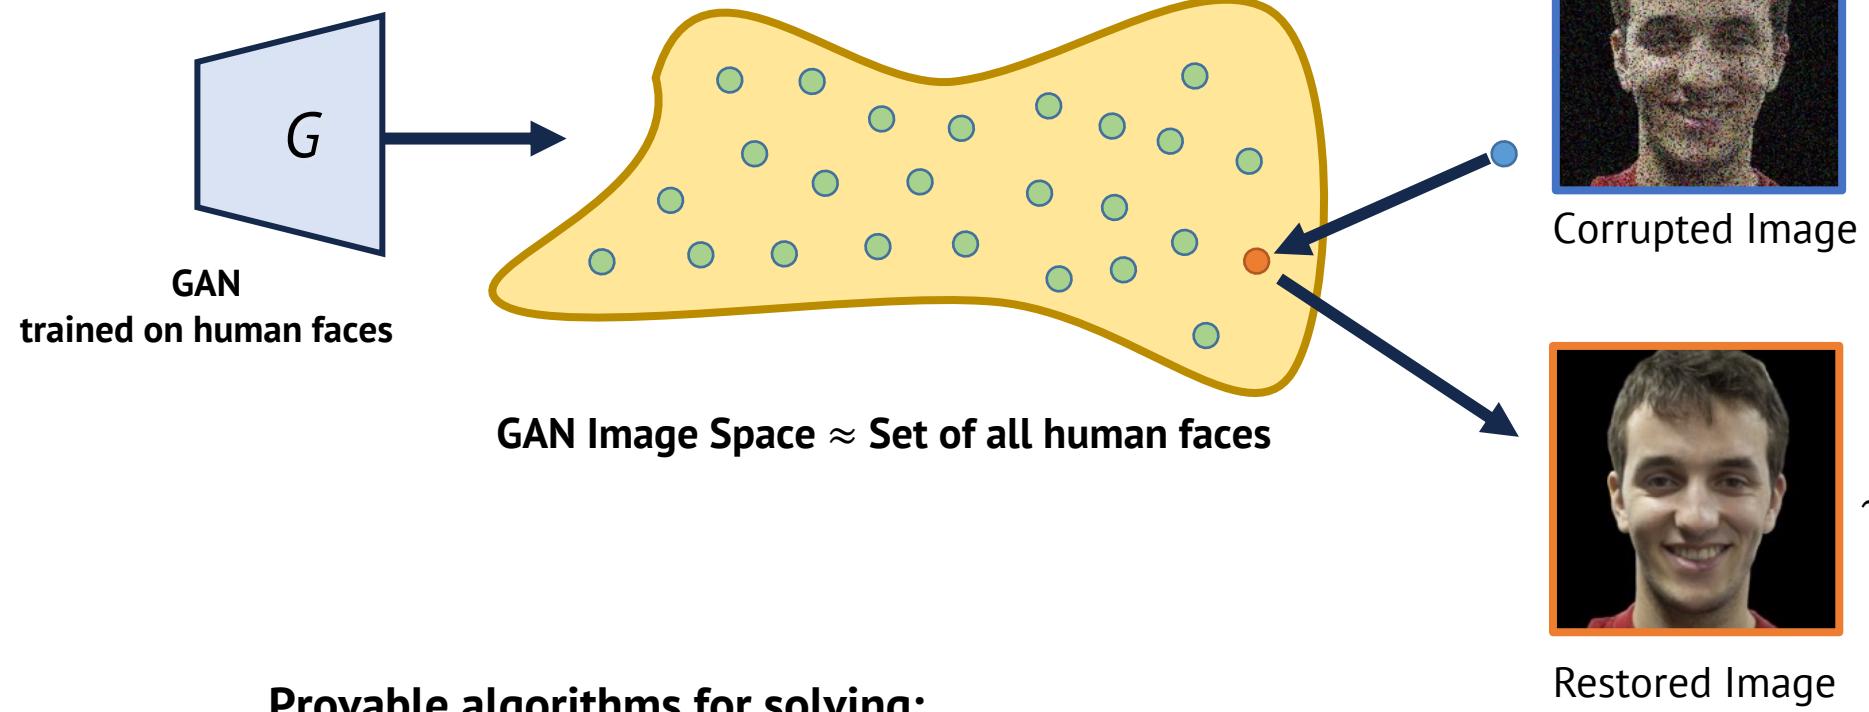
# Image Restoration using GANs



# Image Restoration using GANs



# Image Restoration using Generative Models



Find the closest image in the GAN Image Space

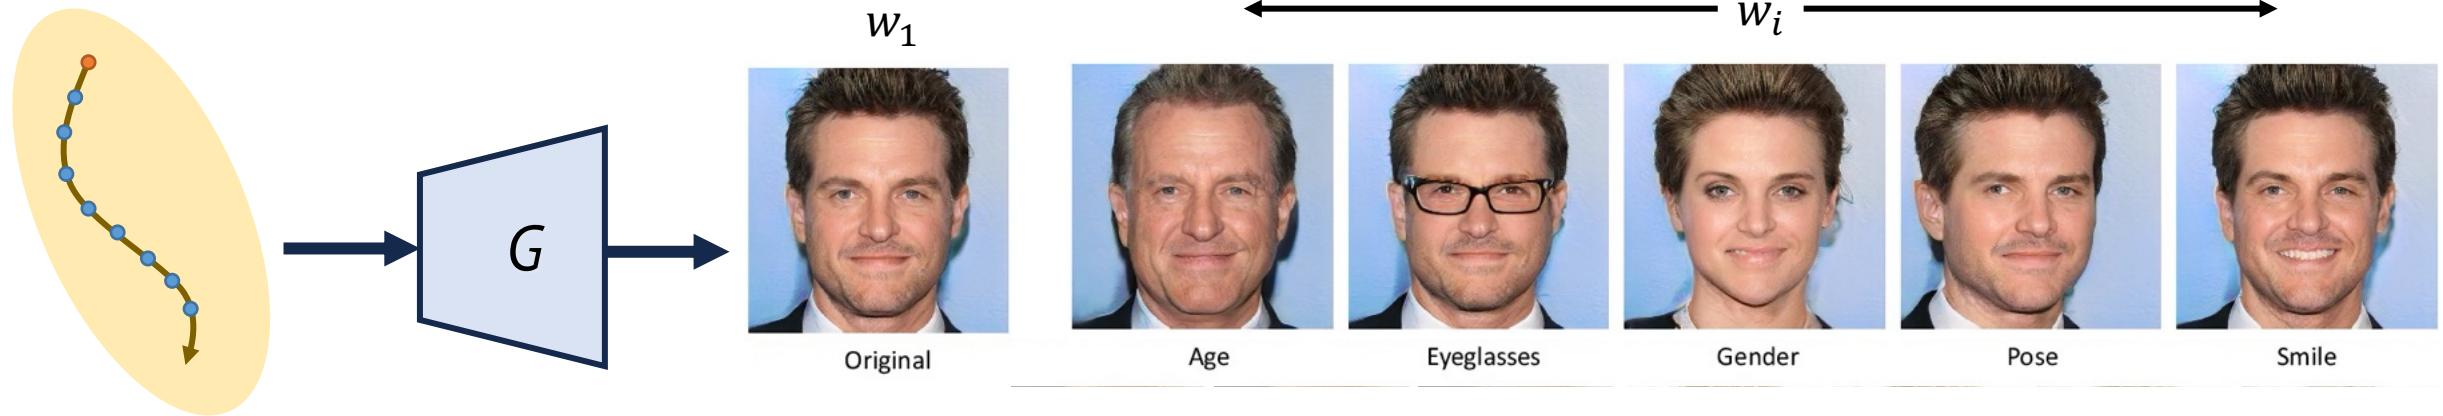
$$w^* = \underset{z}{\operatorname{argmin}} \|y - G(w)\|_2$$

$$\sim G(w^*)$$

## Provable algorithms for solving:

- Image Restoration
- Image Denoising, Inpainting etc.

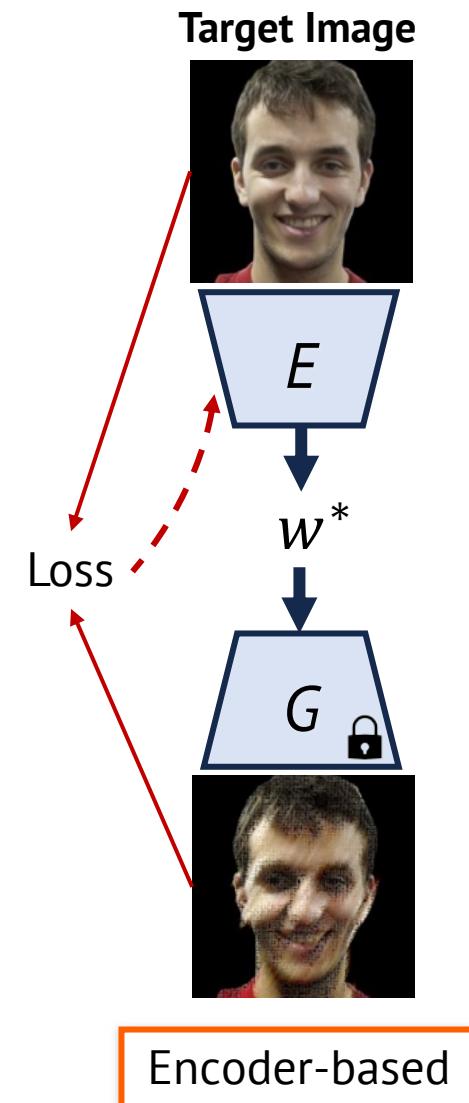
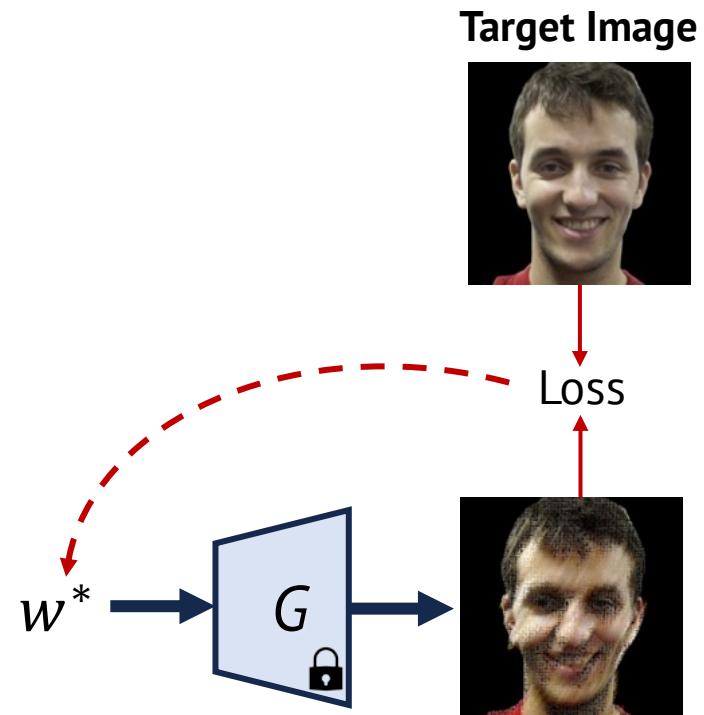
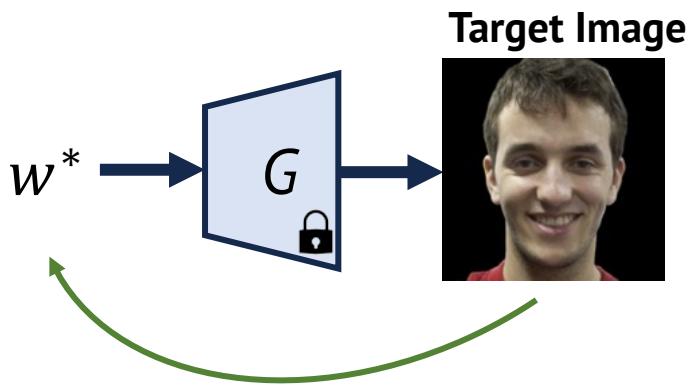
# GAN Inversion is necessary for Image Editing



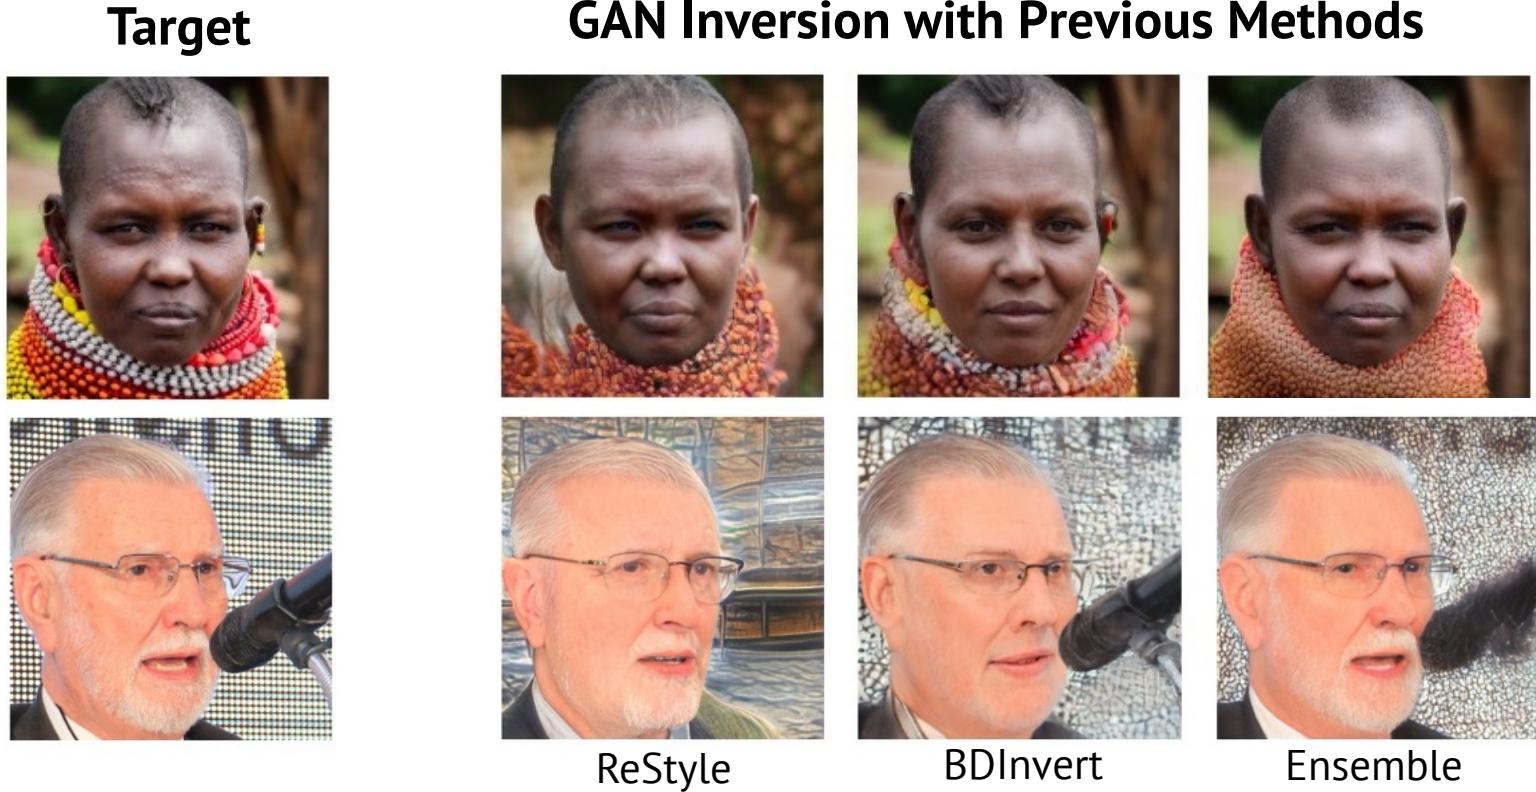
Edit a real image?

**Corresponding Latent code must be known!**

# GAN Inversion

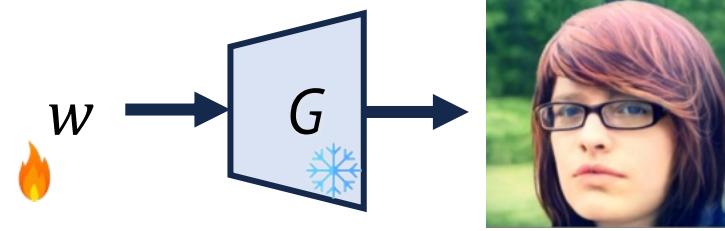
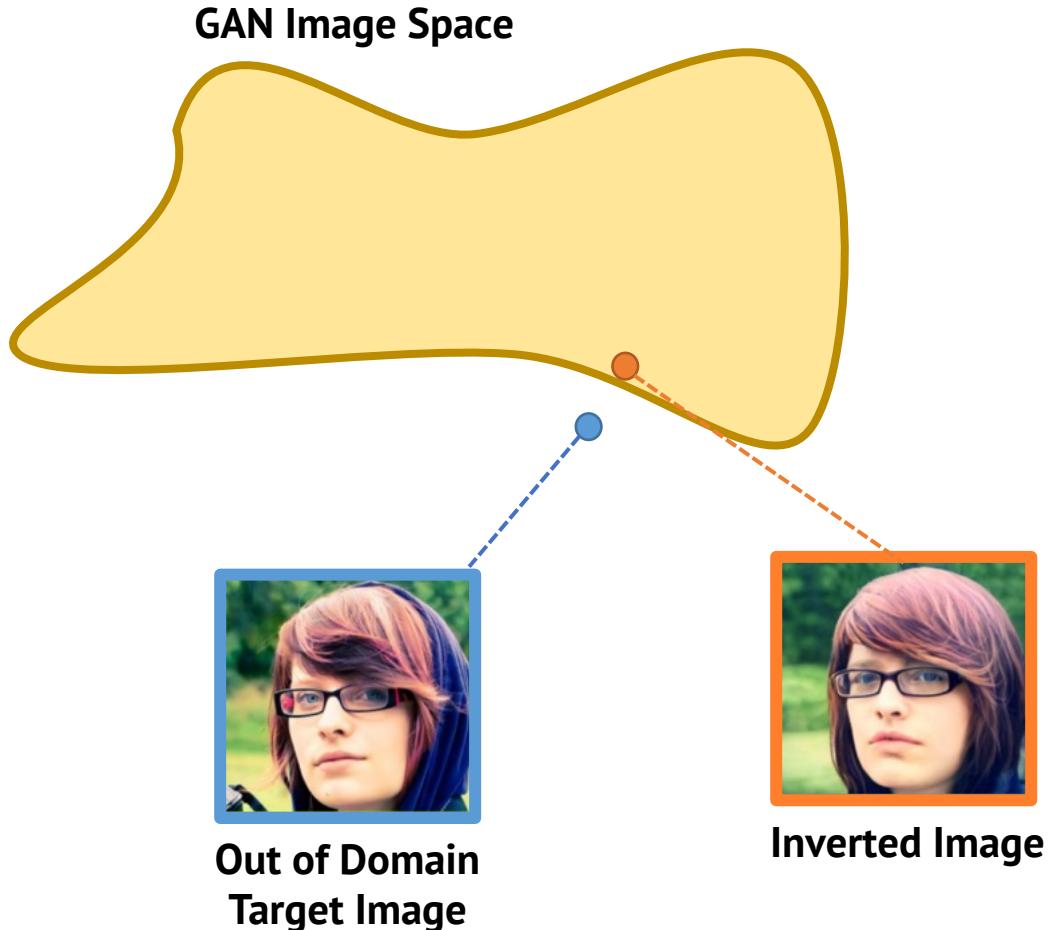


# GAN Inversion: Most Techniques Fail on out-of-domain Images



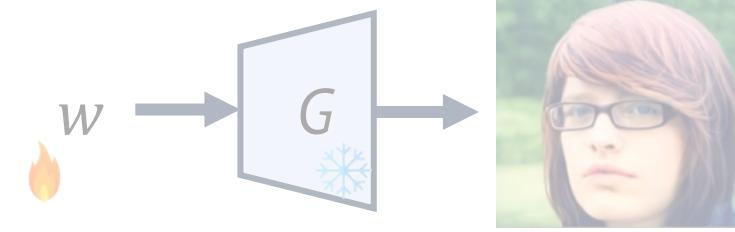
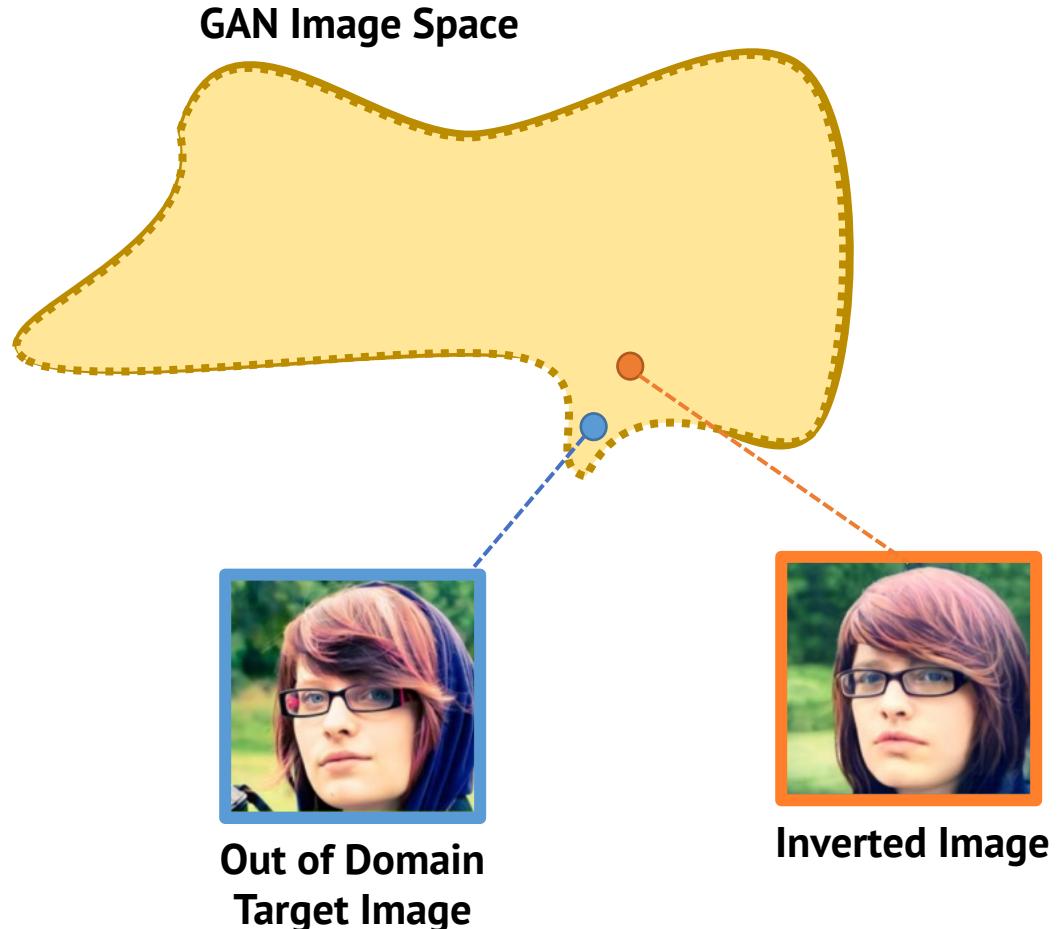
Ref: ReStyle, Alauf et al., ICCV '21  
BDInvert, Kang et al., ICCV '21  
Ensemble, Cai et al., CVPR '21

# Key Challenge in Inverting a Frozen GAN



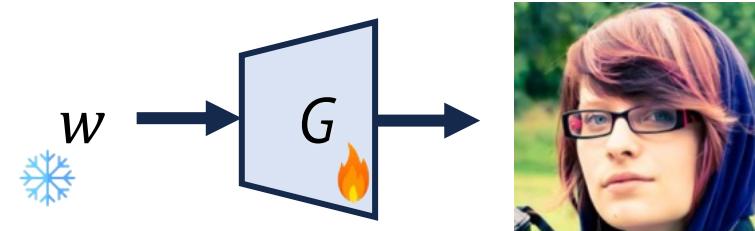
Generator is *Fixed*, Latent is *Flexible*

# Key Challenge in Inverting a Frozen GAN



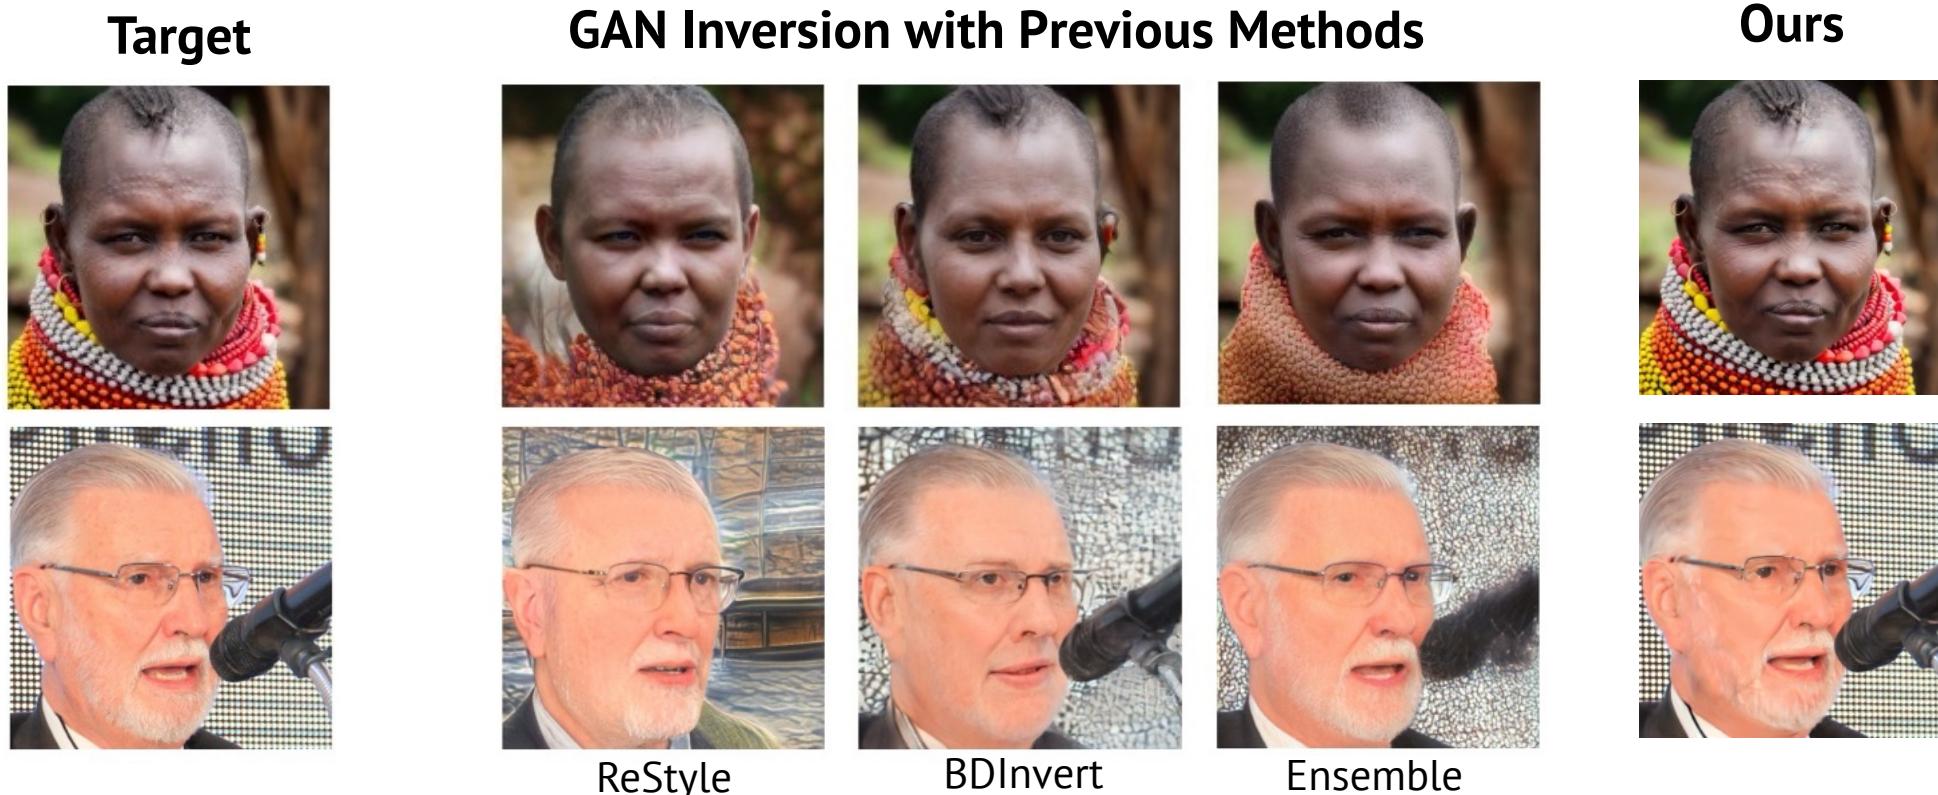
Generator is *Frozen*, Latent is *Trainable*

**Proposed Idea:**



Generator is *Flexible*, Latent is *Fixed*

# Our Method Achieves Near-perfect GAN Inversion

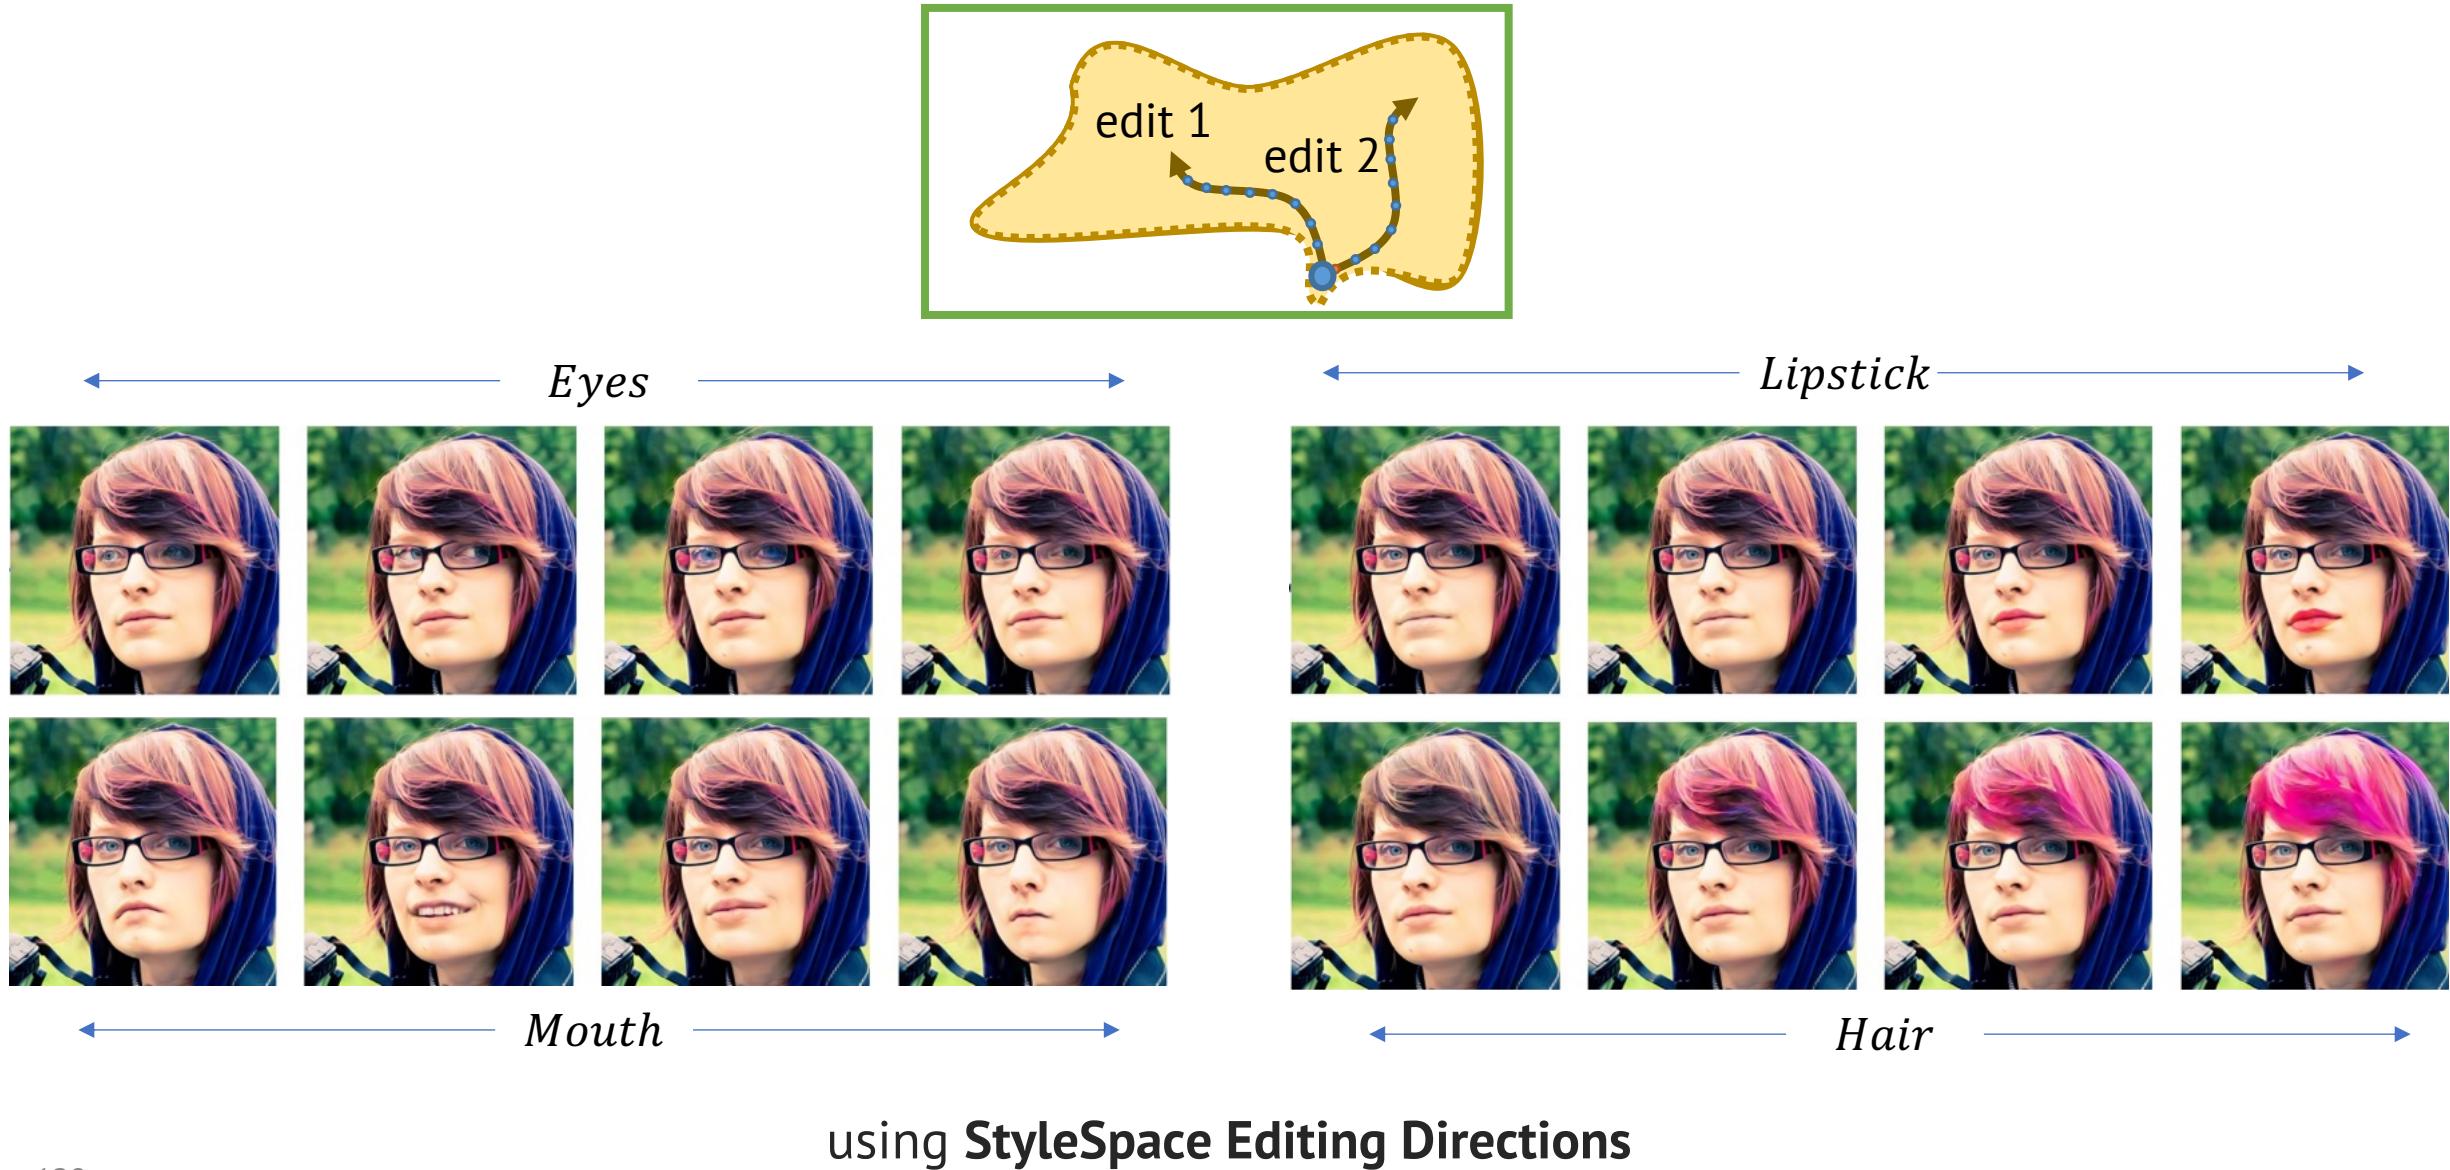


ReStyle

BDInvert

Ensemble

# Most off-the-shelf Editing methods work!



# One-shot Image Stylization



Style Reference



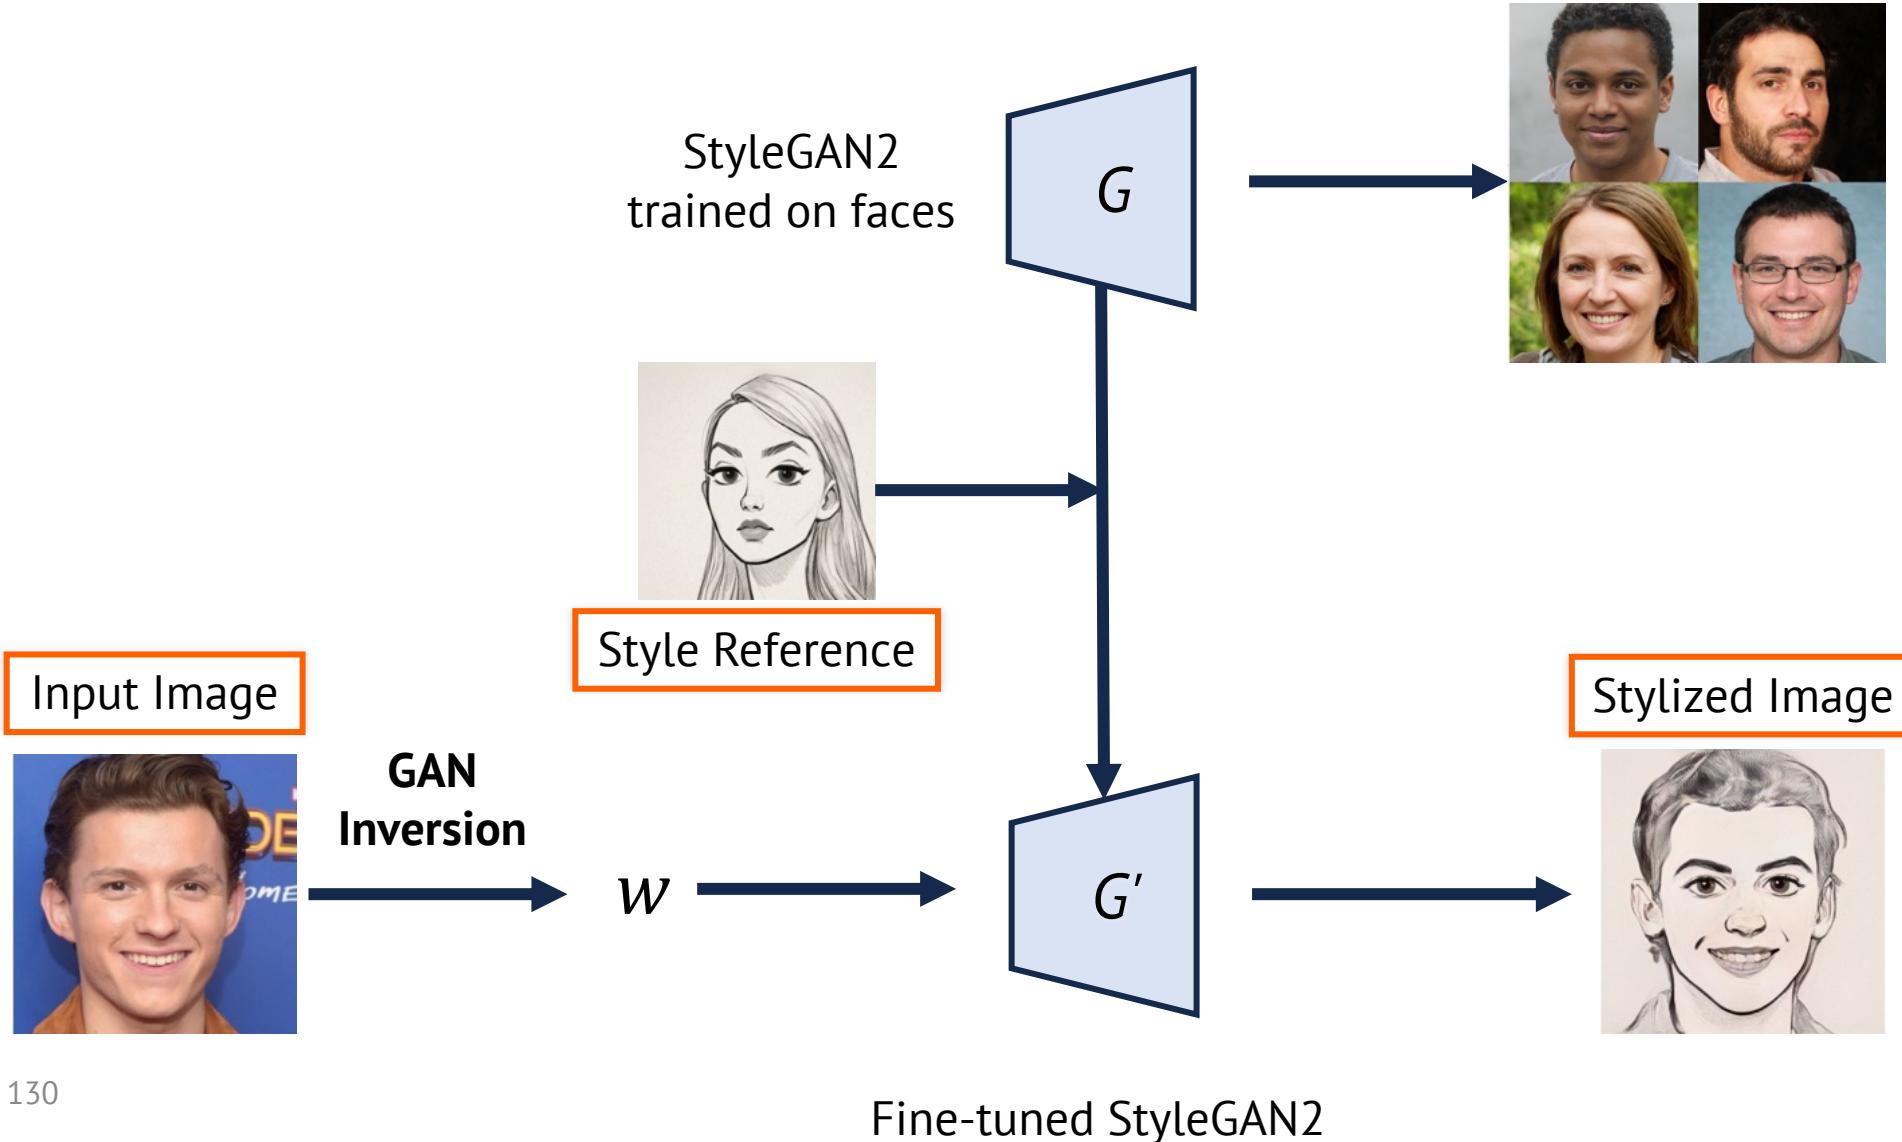
Input Image



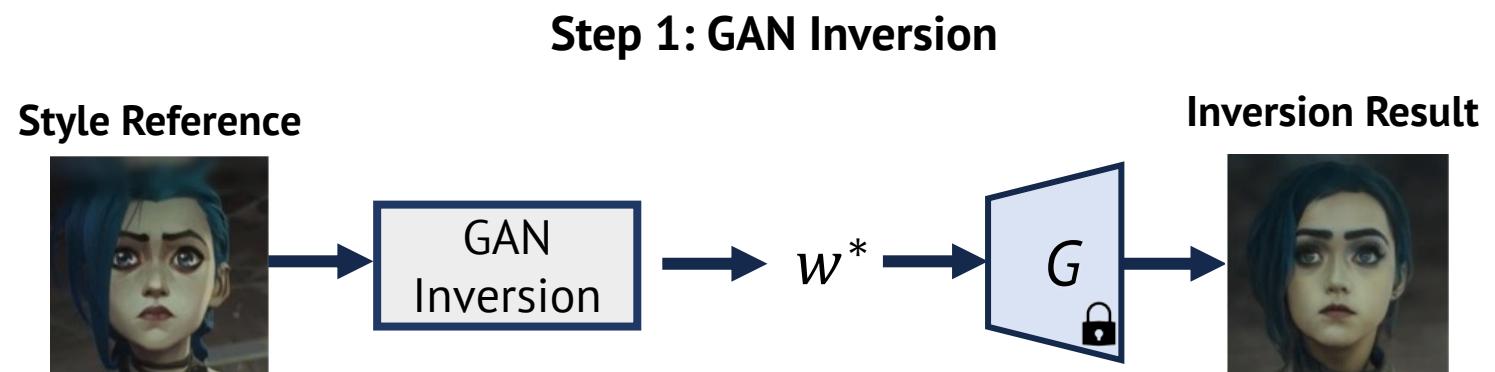
Stylized Image

**Stylizing an Input Image in the style of a Reference Style using only *one* example**

# Pre-trained GAN for one-shot Stylization



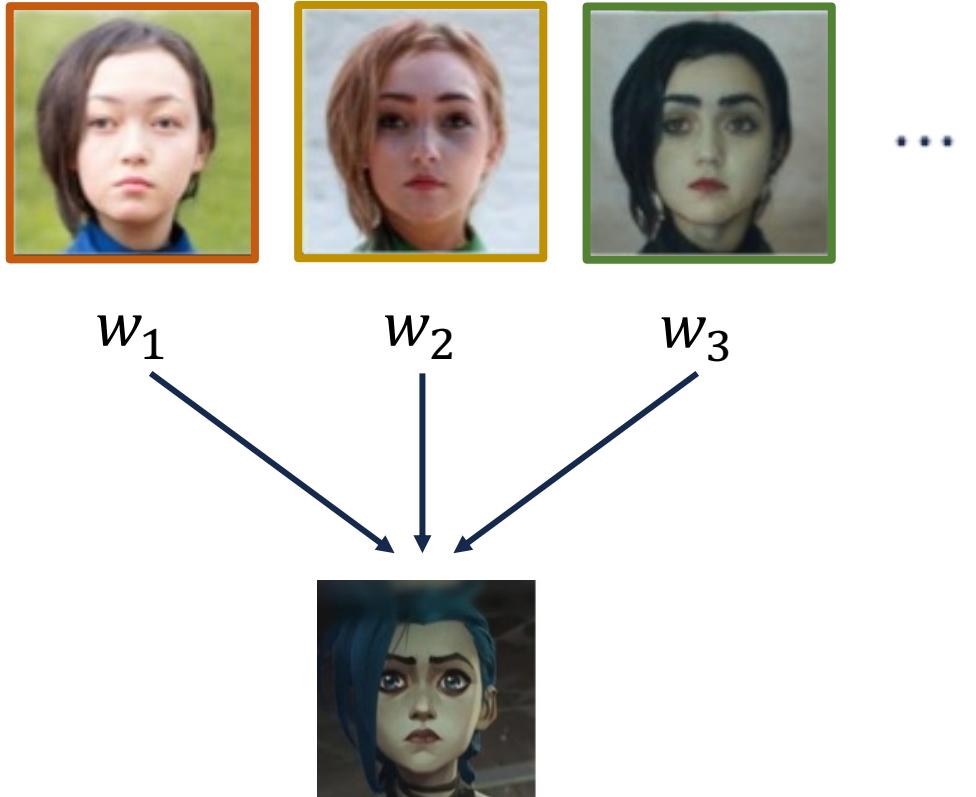
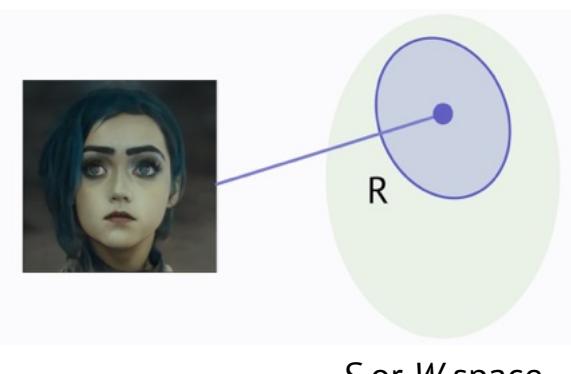
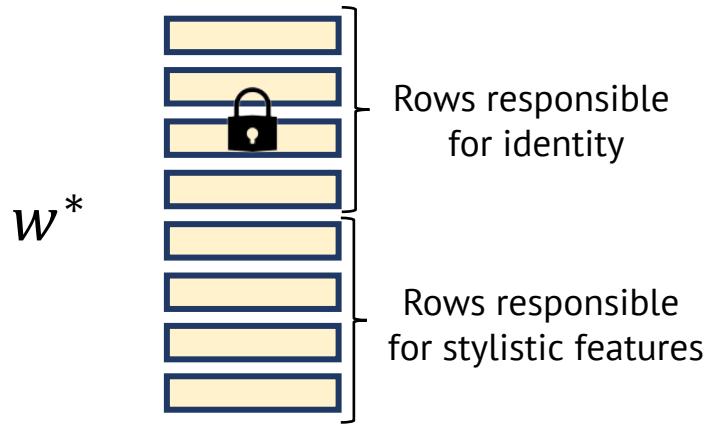
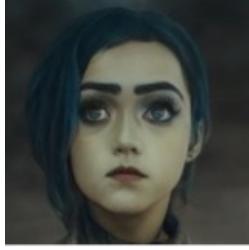
# How to perform global customization of StyleGAN2?



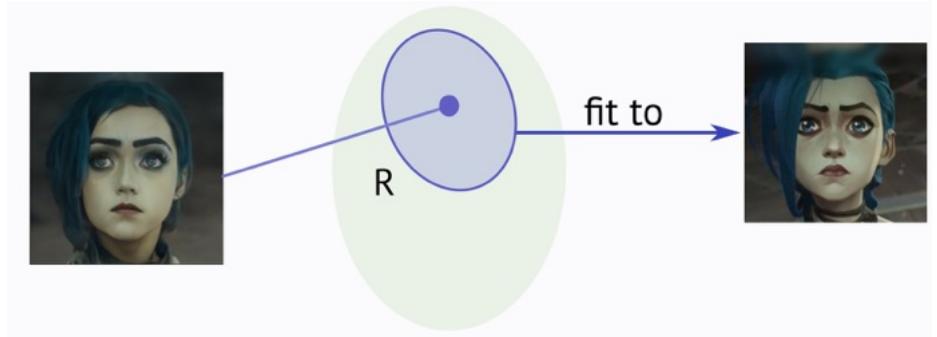
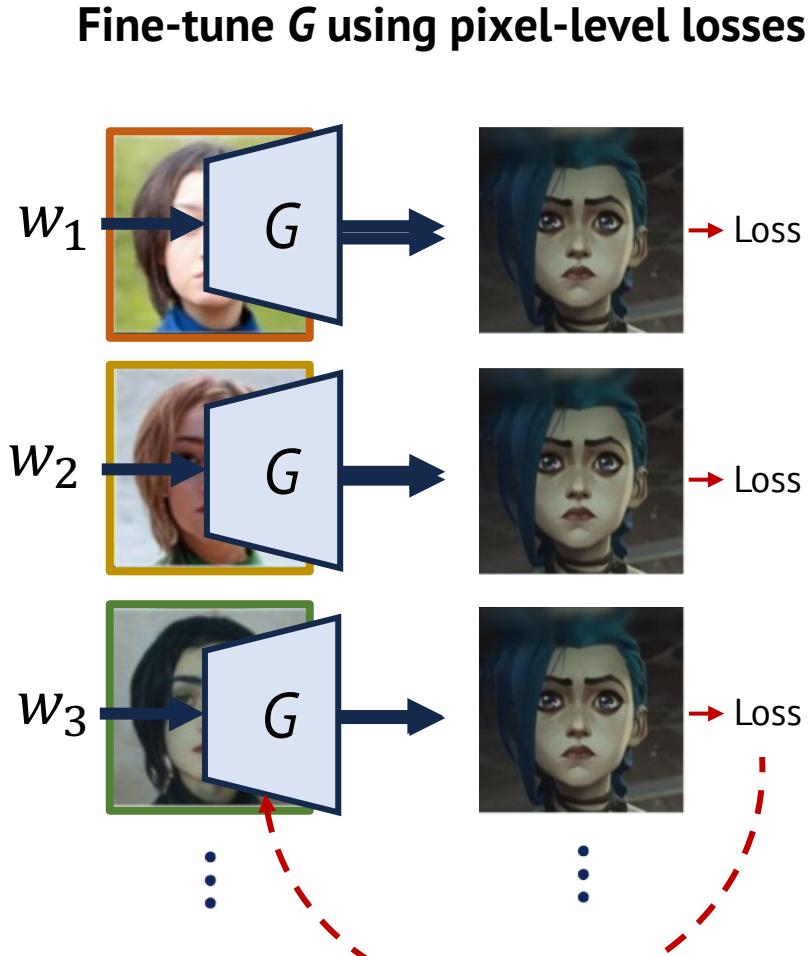
# JoJoGAN: Style-mixing

## Step 2. Use Style-mixing property to Create a Training Set

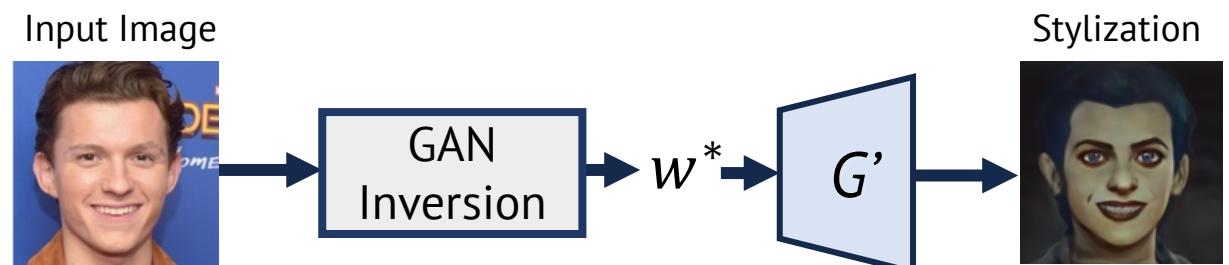
**Inversion result**



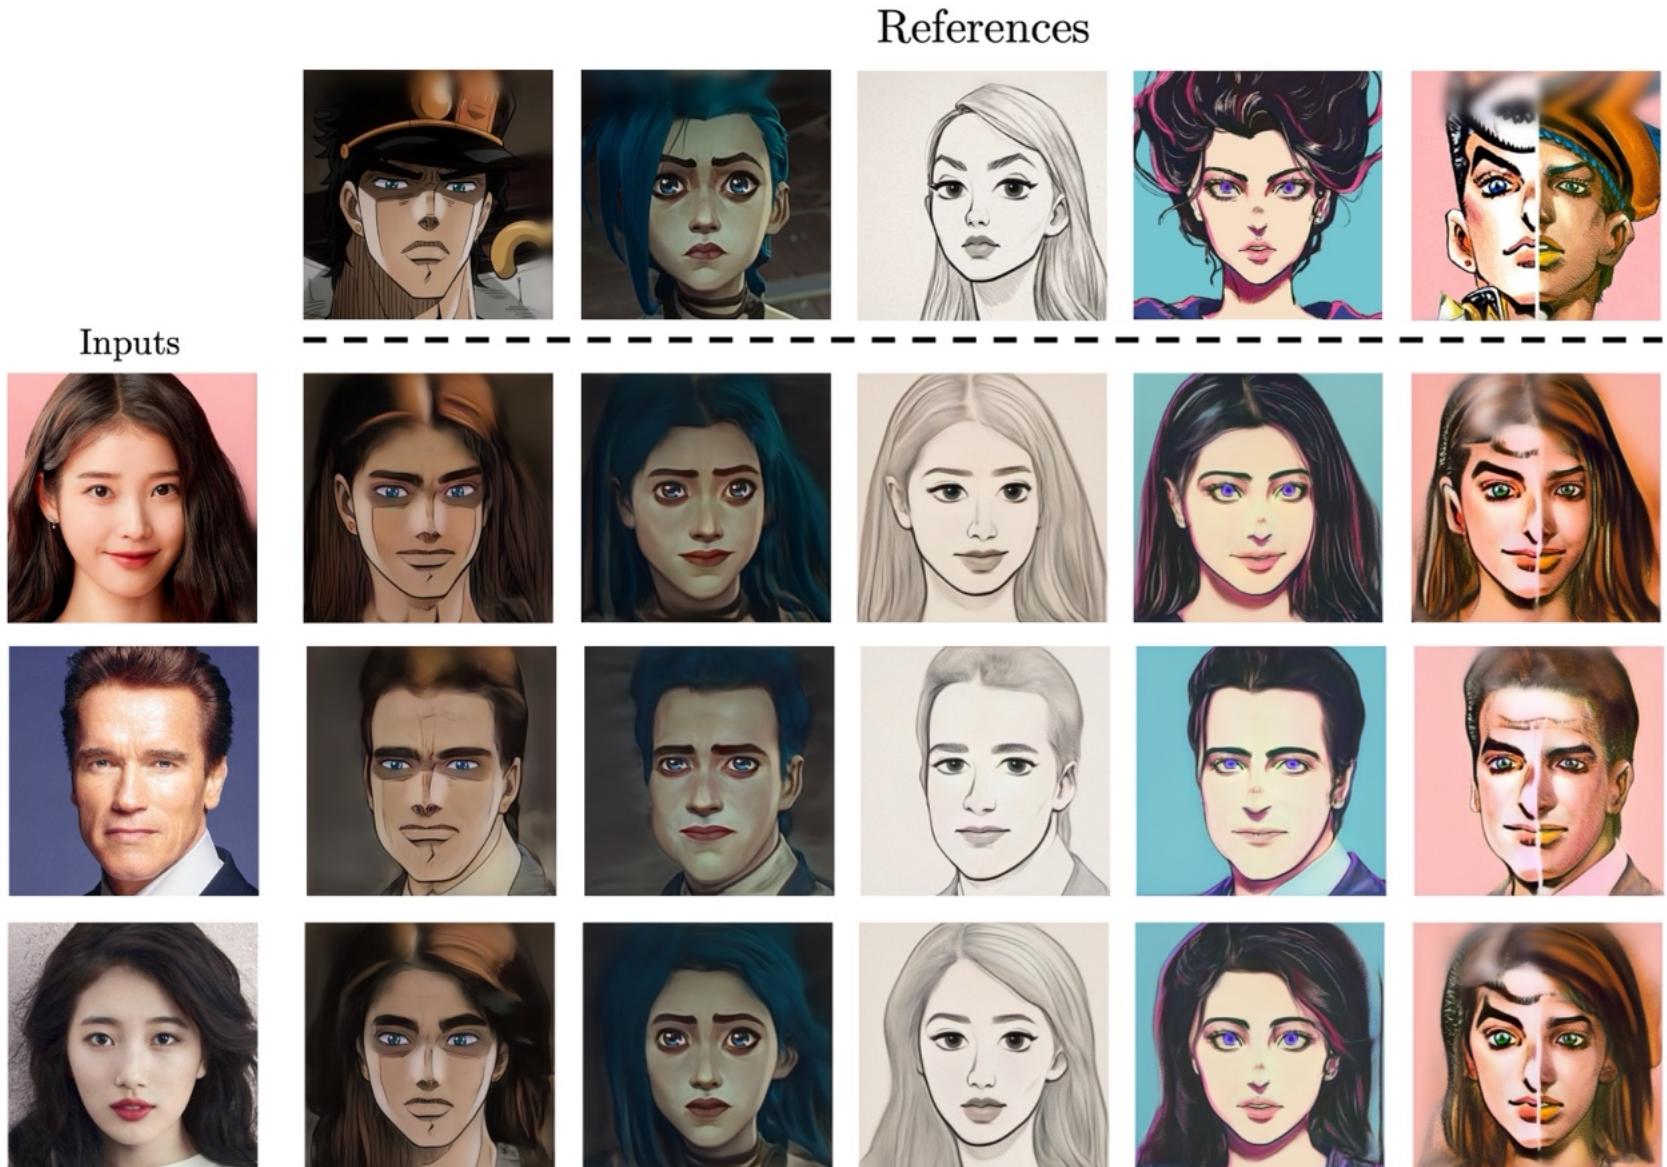
# How to Fine-tune StyleGAN2 ?: JoJoGAN



**Inference on fine-tuned  $G$**



# JoJoGAN: Results





# GANs: Summary

- Tremendous progress in training algorithm, architecture, and conditioning mechanism
- Rich understanding of image features with disentangled representation
- Smooth interpolations in latent space
- Vast variety of applications of pre-trained GANs in all engineering domains through,
  - GAN conditioning
  - GAN Inversion + latent space traversal
  - GAN customization (fine-tuning)
- Are GANs still relevant?! : [current and future trends](#)