

# Practical GAN Training

Neural Networks Design And Application

# Tricks


## Architecture guidelines for stable Deep Convolutional GANs

- Replace any pooling layers with strided convolutions (discriminator) and fractional-strided convolutions (generator).
- Use batchnorm in both the generator and the discriminator.
- Remove fully connected hidden layers for deeper architectures.
- Use ReLU activation in generator for all layers except for the output, which uses Tanh.
- Use LeakyReLU activation in the discriminator for all layers.



Radford, Alec, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks." *arXiv preprint arXiv:1511.06434* (2015).

<https://arxiv.org/pdf/1511.06434.pdf>

# Replace pooling with strided conv layer



# Replace pooling with strided conv layer



Stride 2 Padding 0  
Strided Convolution

# Batch normalization [BN]

**Input:** Values of  $x$  over a mini-batch:  $\mathcal{B} = \{x_{1\dots m}\}$ ;

Parameters to be learned:  $\gamma, \beta$

**Output:**  $\{y_i = \text{BN}_{\gamma, \beta}(x_i)\}$

$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{ mini-batch mean}$$

$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{ mini-batch variance}$$

$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{ normalize}$$


$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma, \beta}(x_i) \quad // \text{ scale and shift}$$

Rescaling for a batch

A linear model as output:  
There are two learnable parameters

**Algorithm 1:** Batch Normalizing Transform, applied to activation  $x$  over a mini-batch.

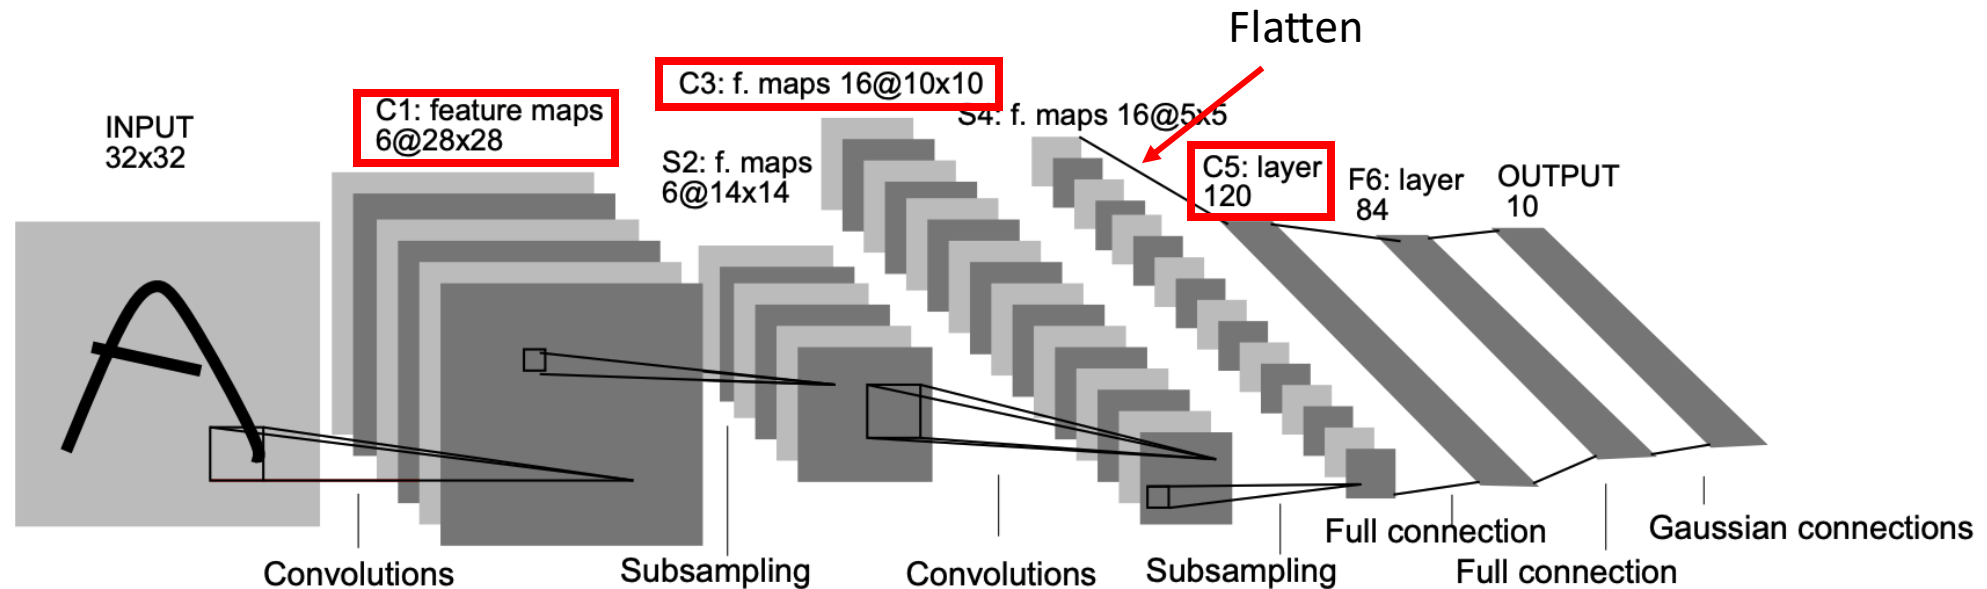
# LeNet-5 in 1999



**Fig. 1.** Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

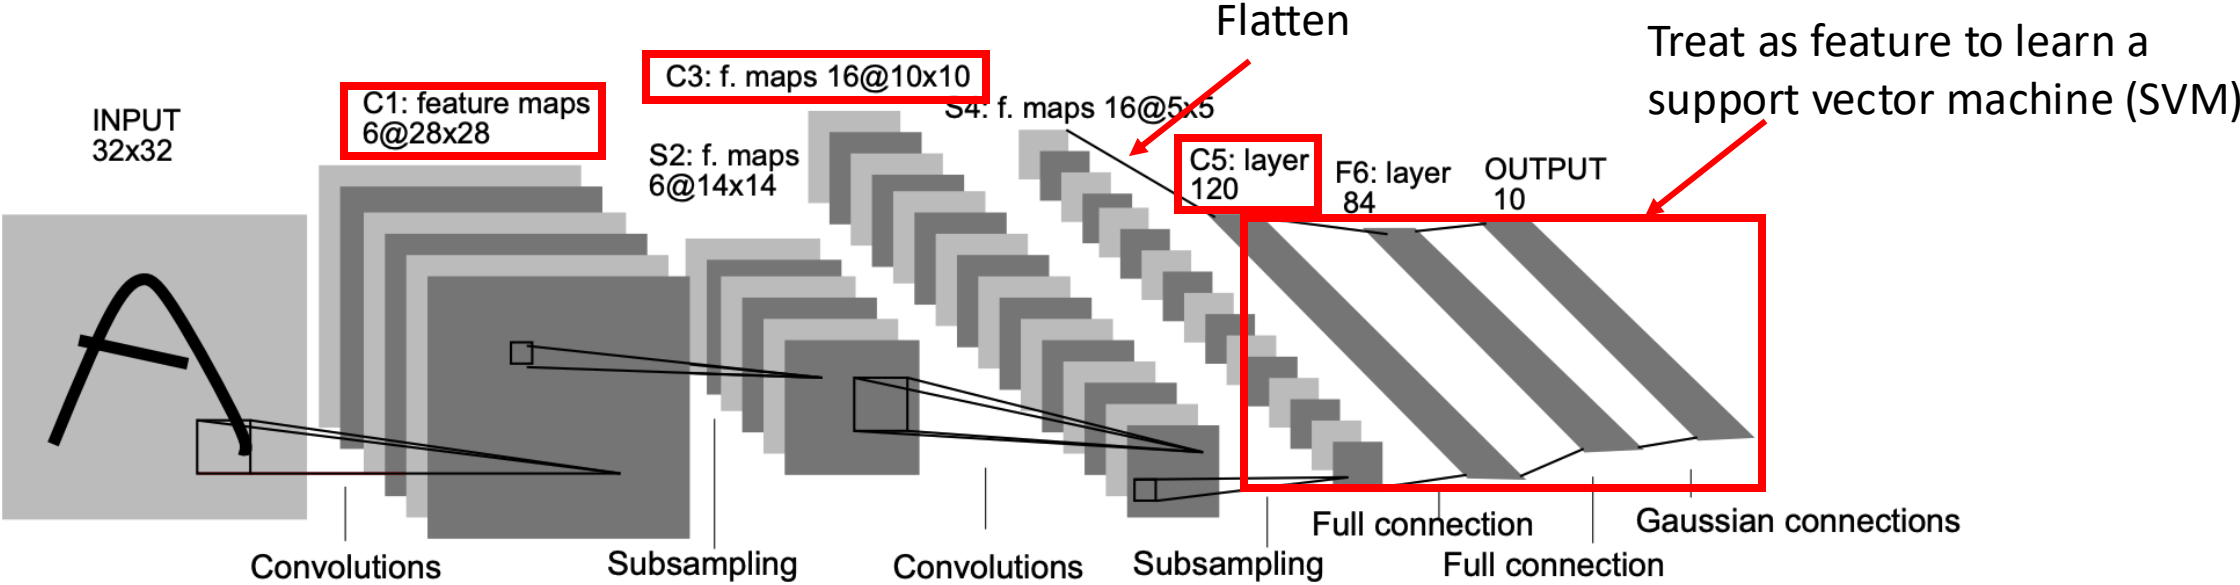
LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

# Remove full connected layer



**Fig. 1.** Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

# Remove full connected layer



**Fig. 1.** Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

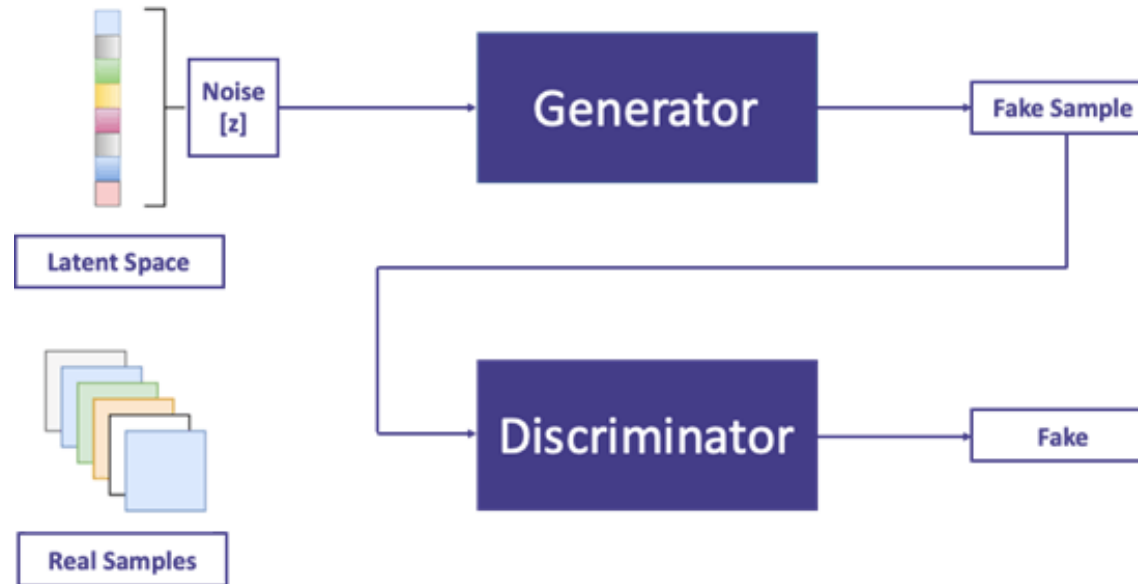
LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.



# Architectures

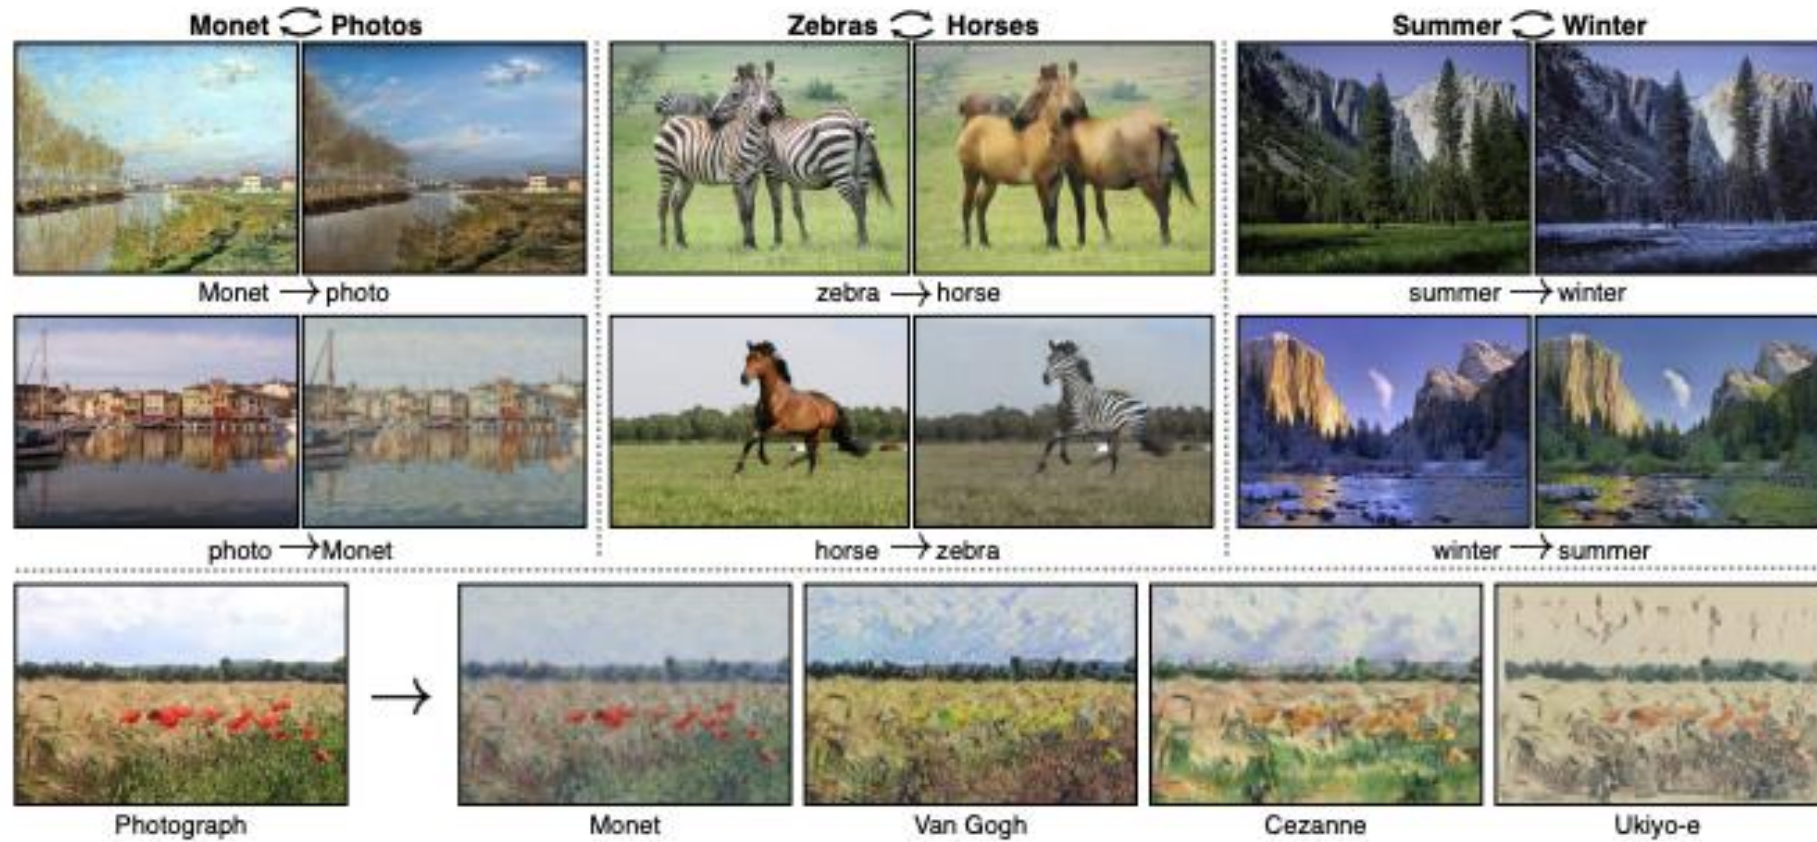
- CycleGAN
- StyleGAN
- ...

# GAN




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

# CycleGAN




# CycleGAN



Q1: content changed?



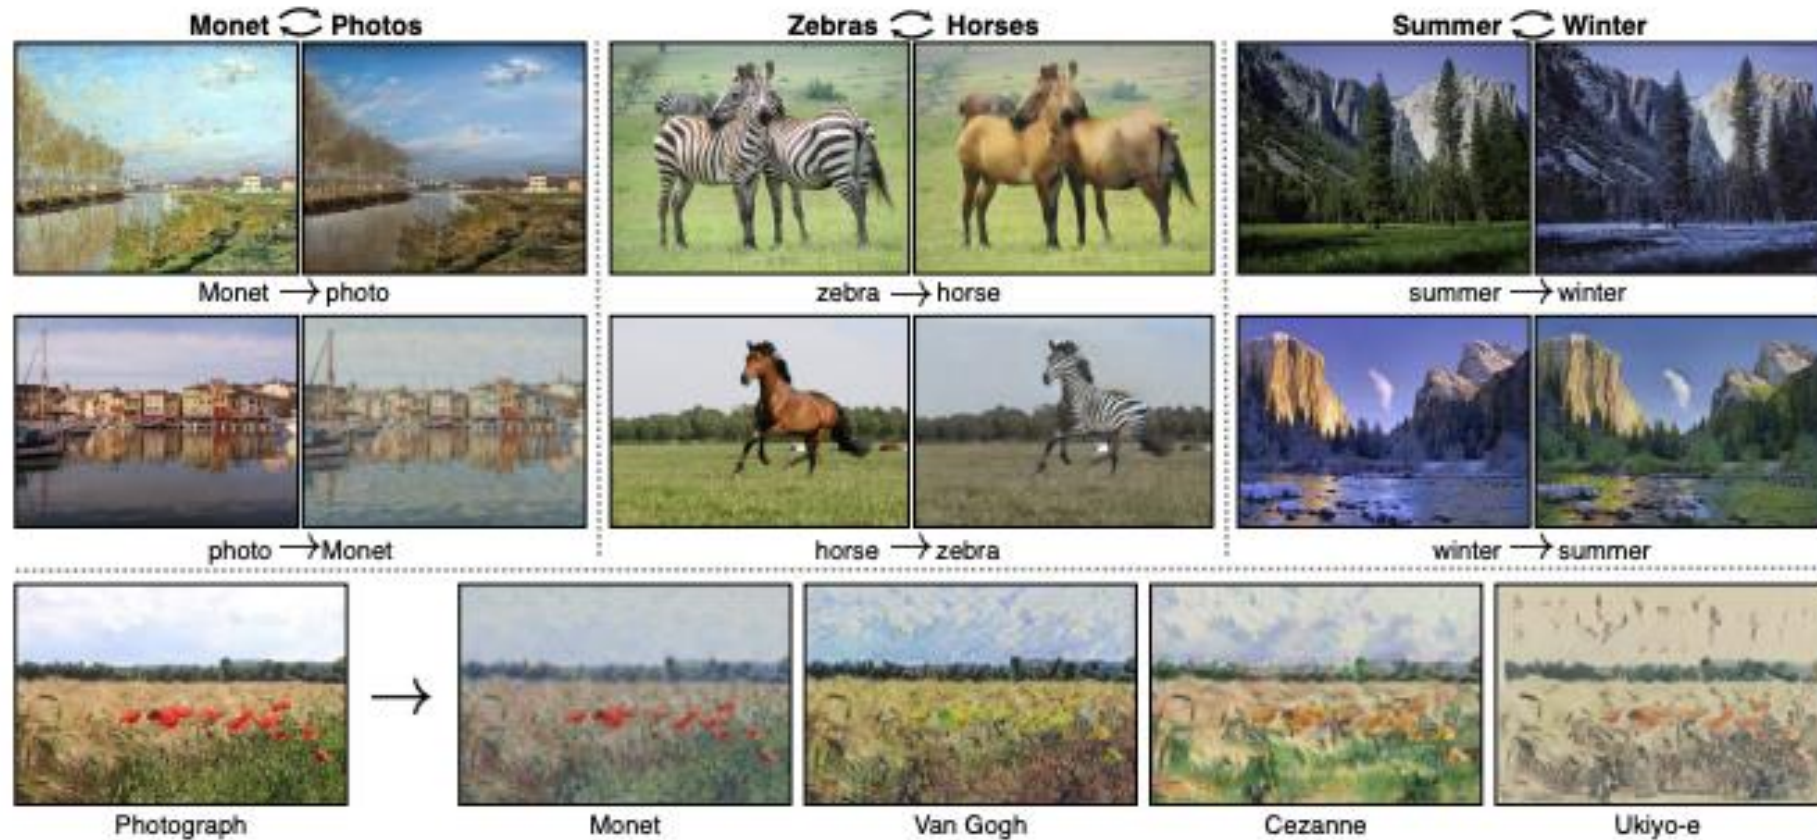
# CycleGAN



Q1: content changed?

Q2: what changed?


# CycleGAN




Q1: content changed?

Q2: what changed? → image-to-image translation

# CycleGAN





# CycleGAN






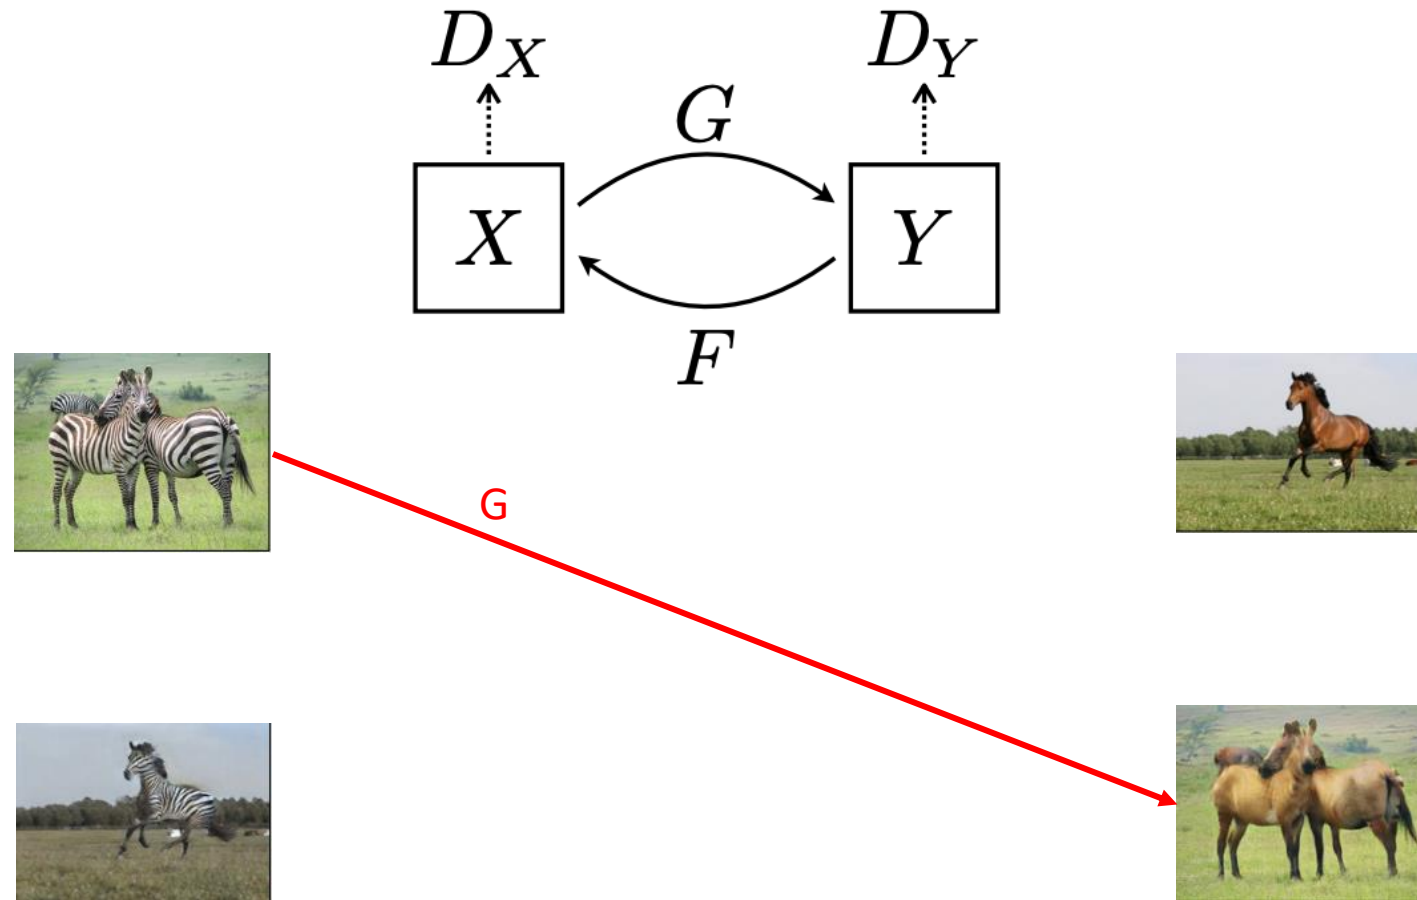
# CycleGAN




# CycleGAN




# CycleGAN




# CycleGAN



# CycleGAN




# CycleGAN








Q: how to ensure the content is not changed?

# CycleGAN

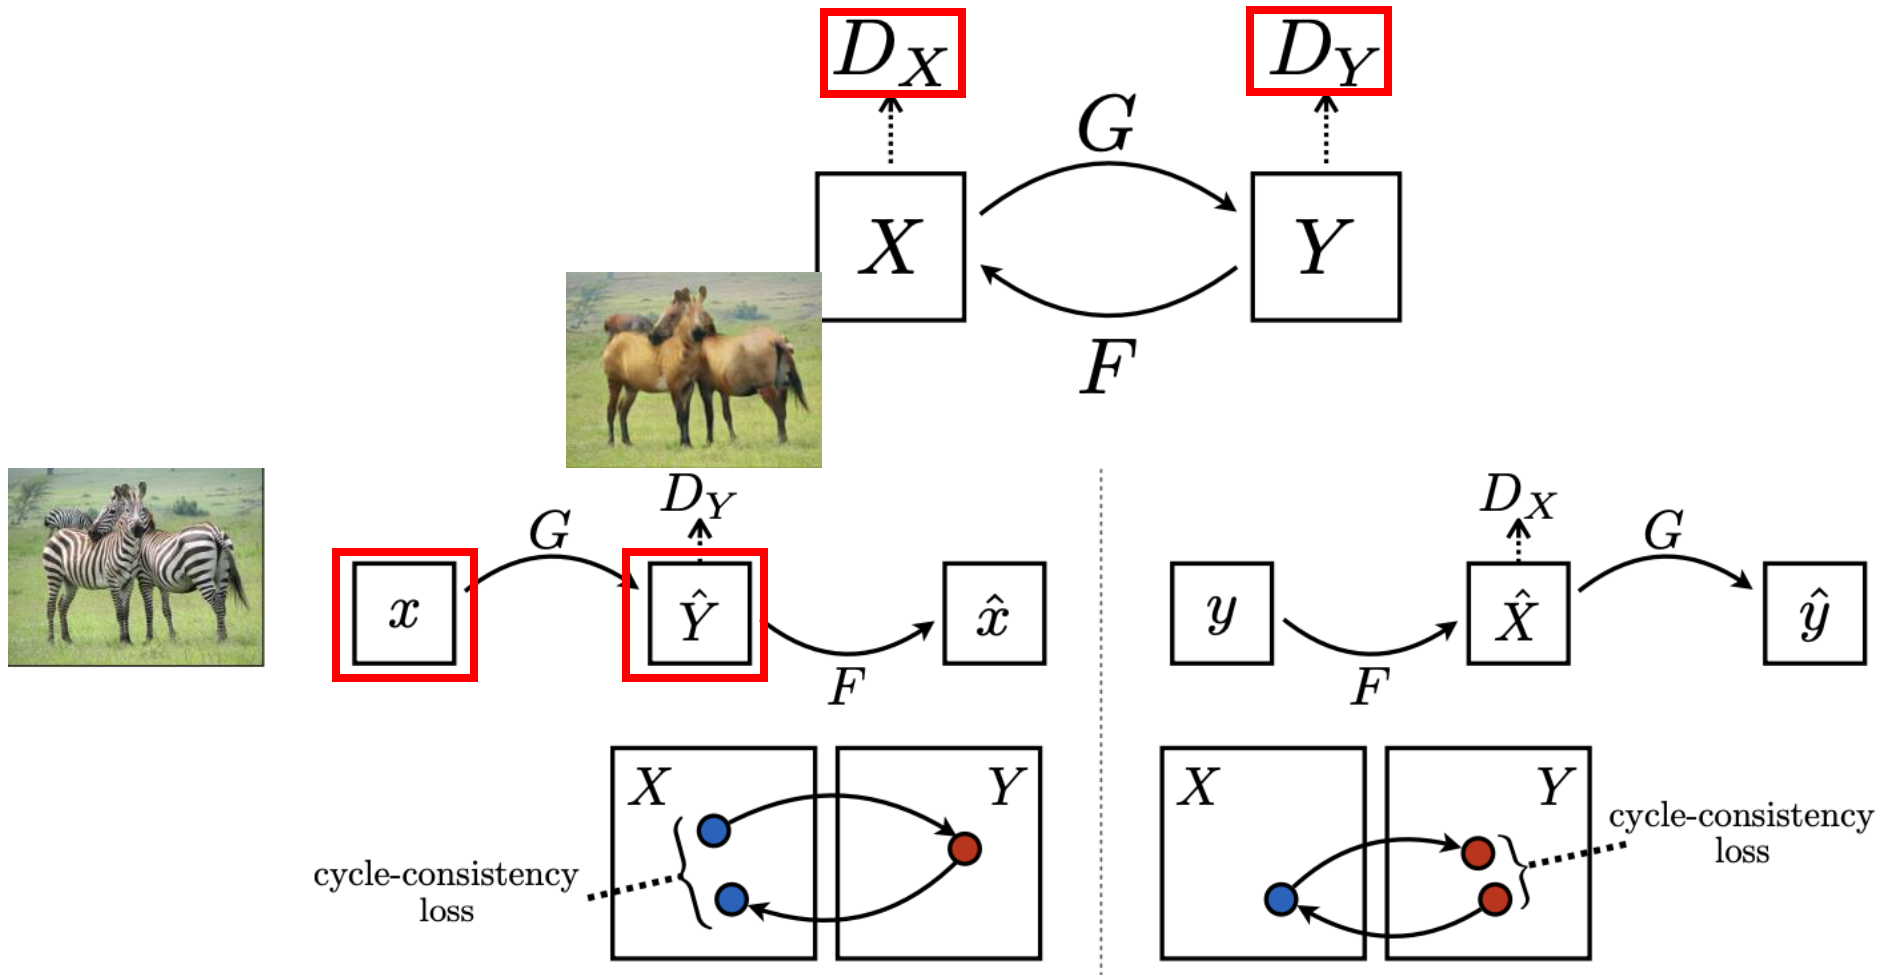


# CycleGAN

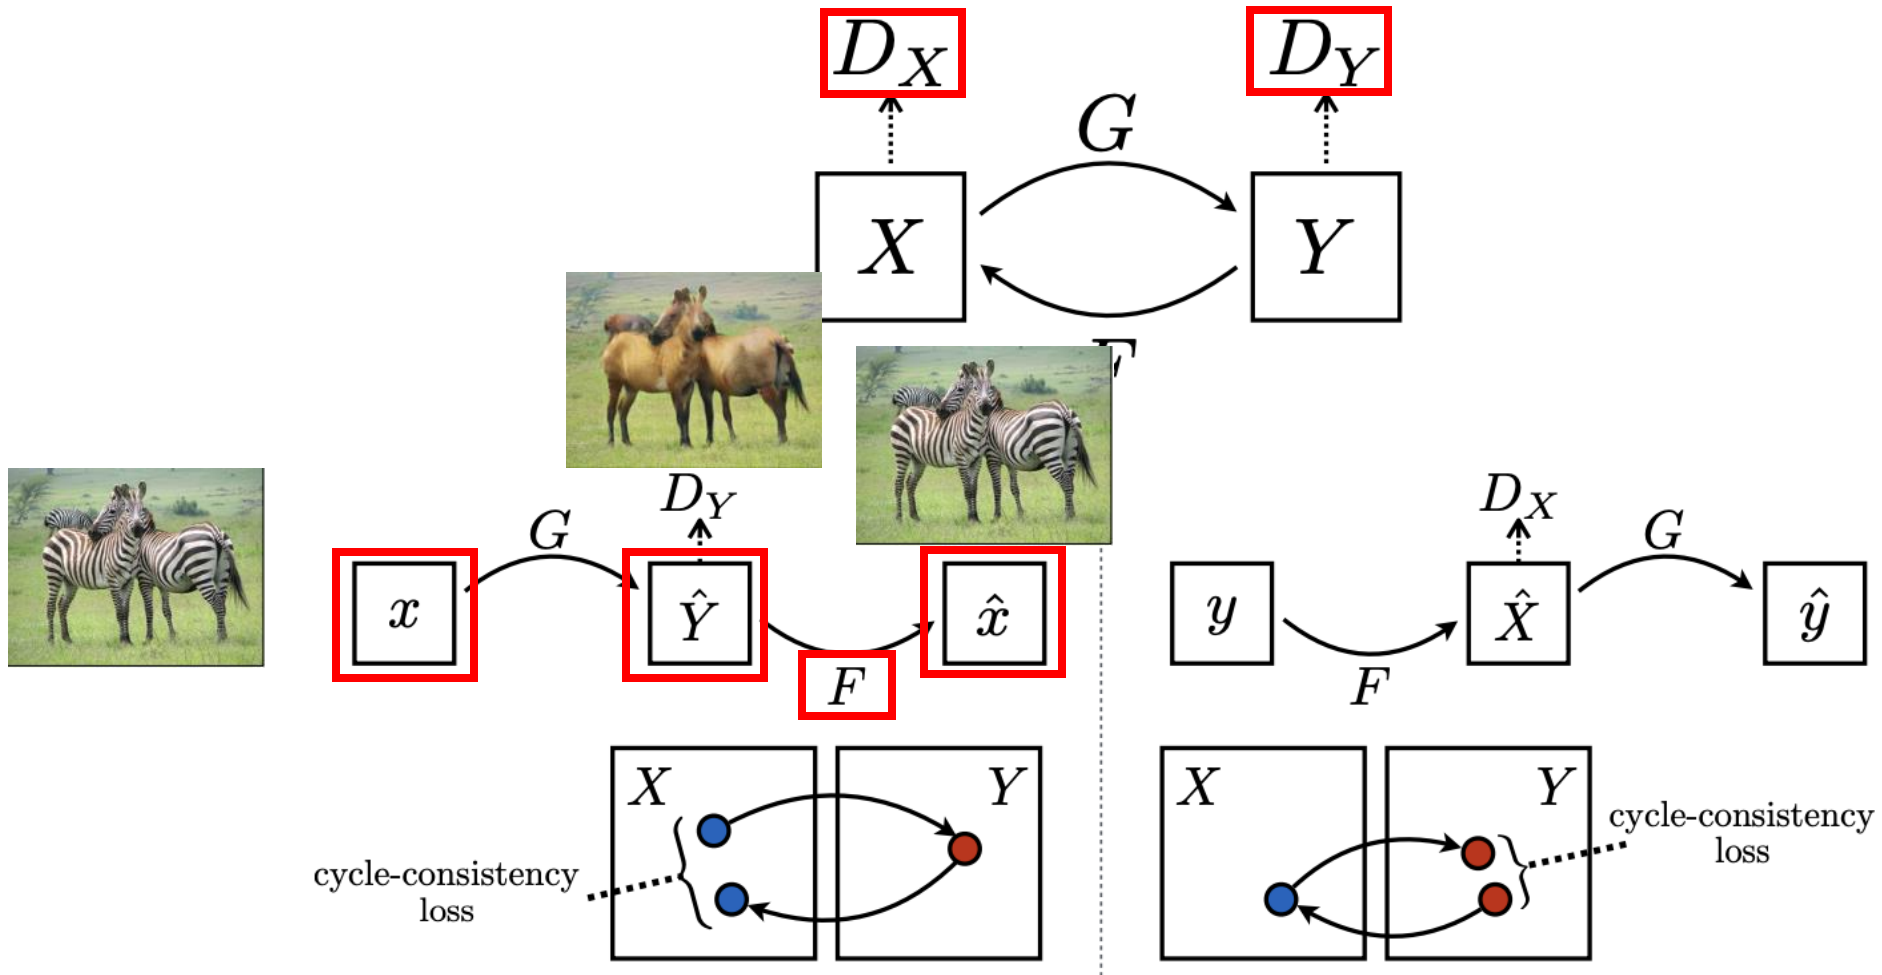





# CycleGAN



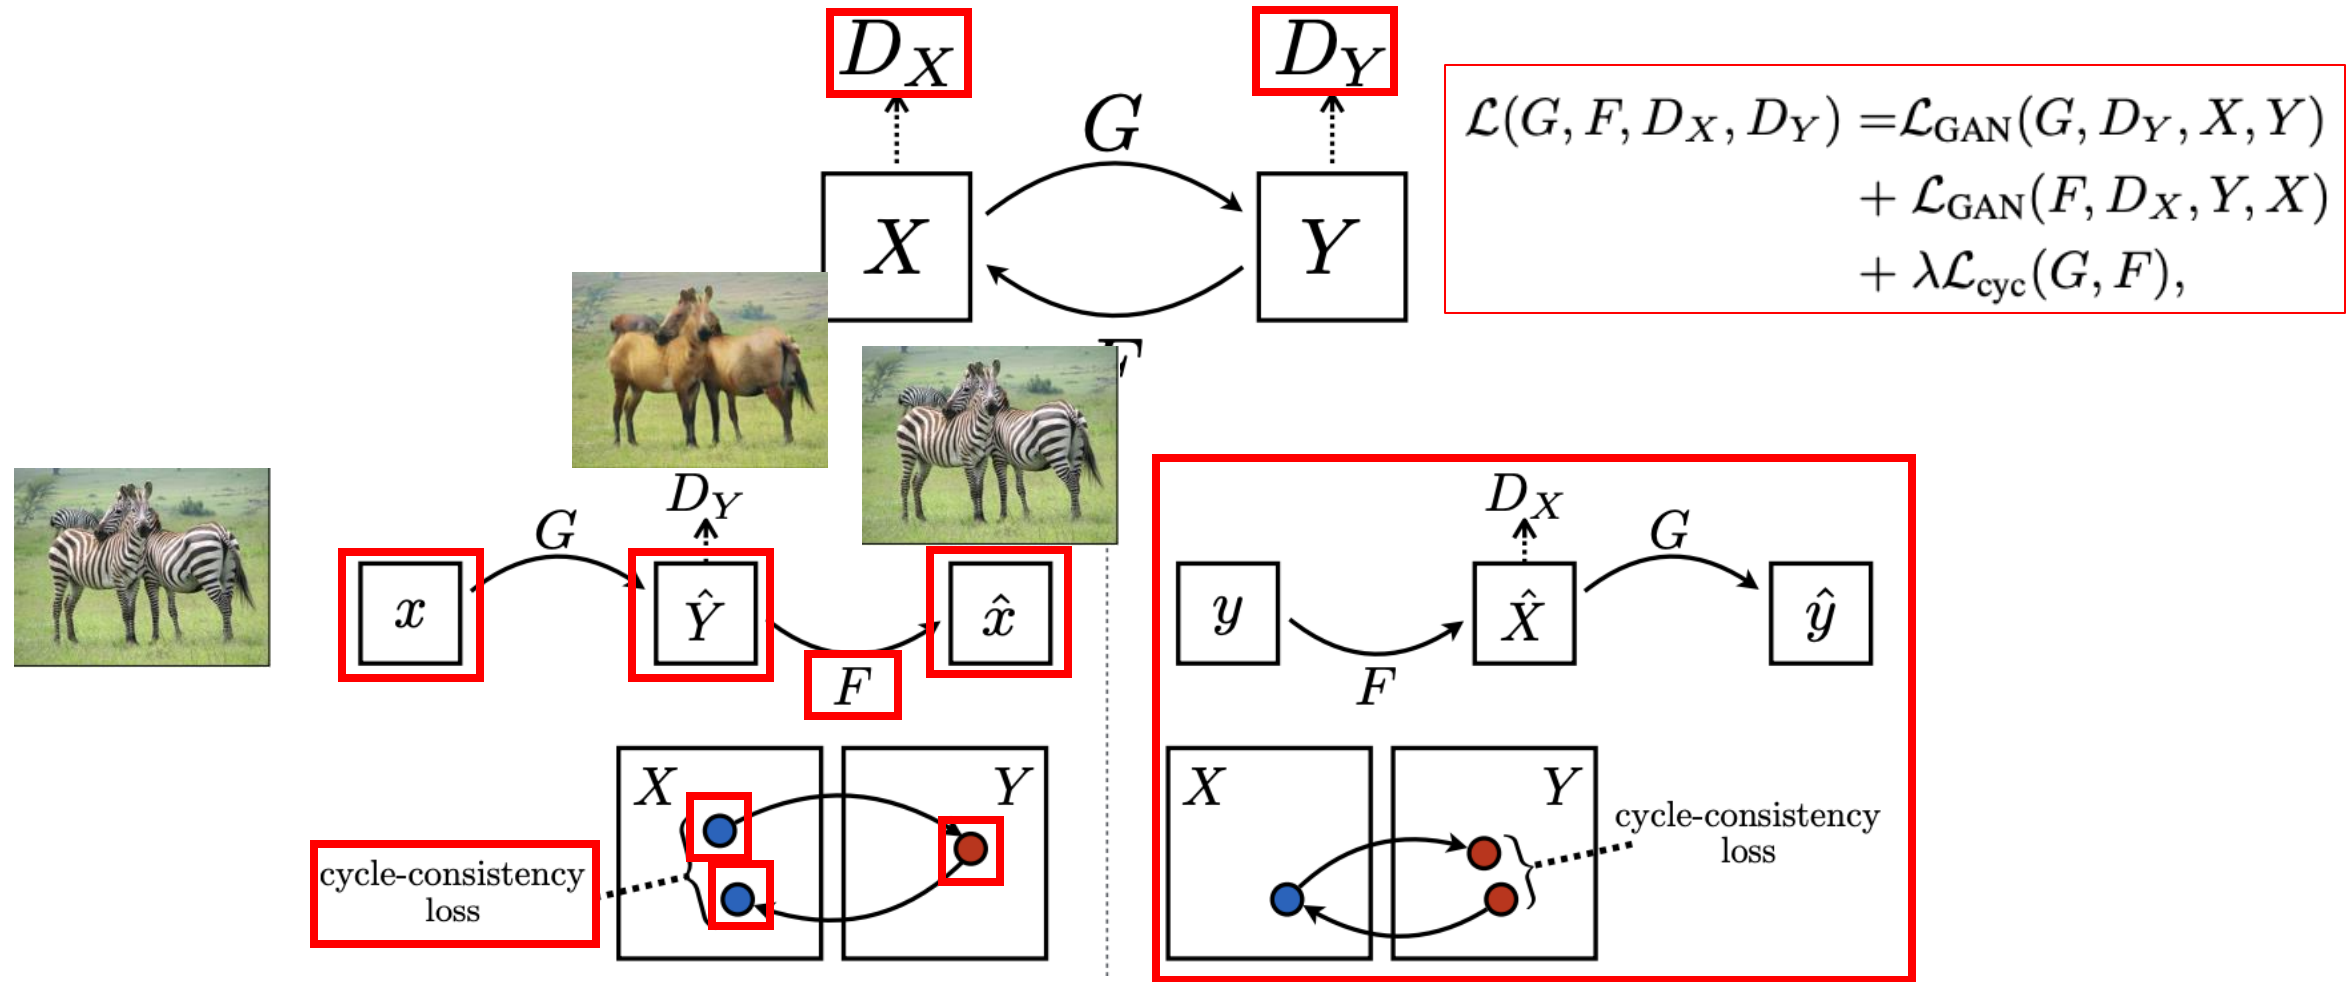
# CycleGAN



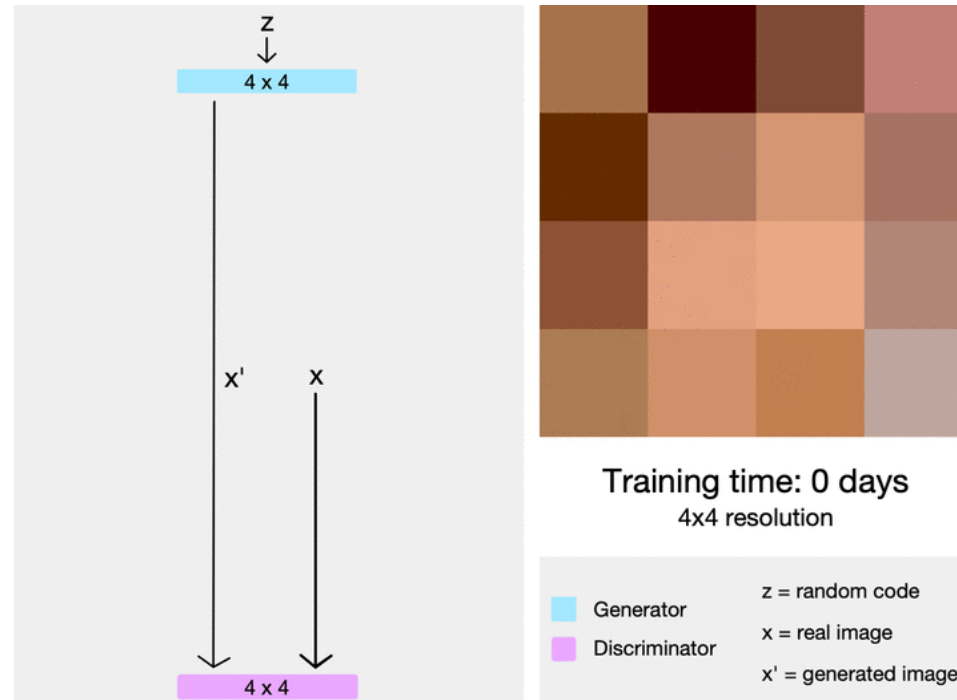
# CycleGAN



# CycleGAN

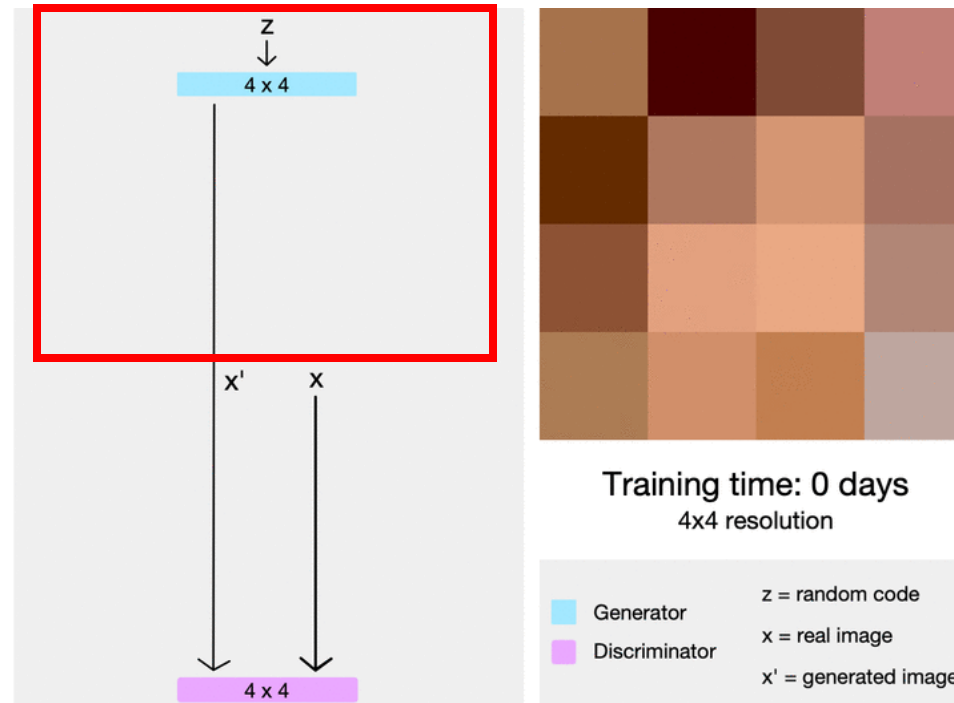


# StyleGAN



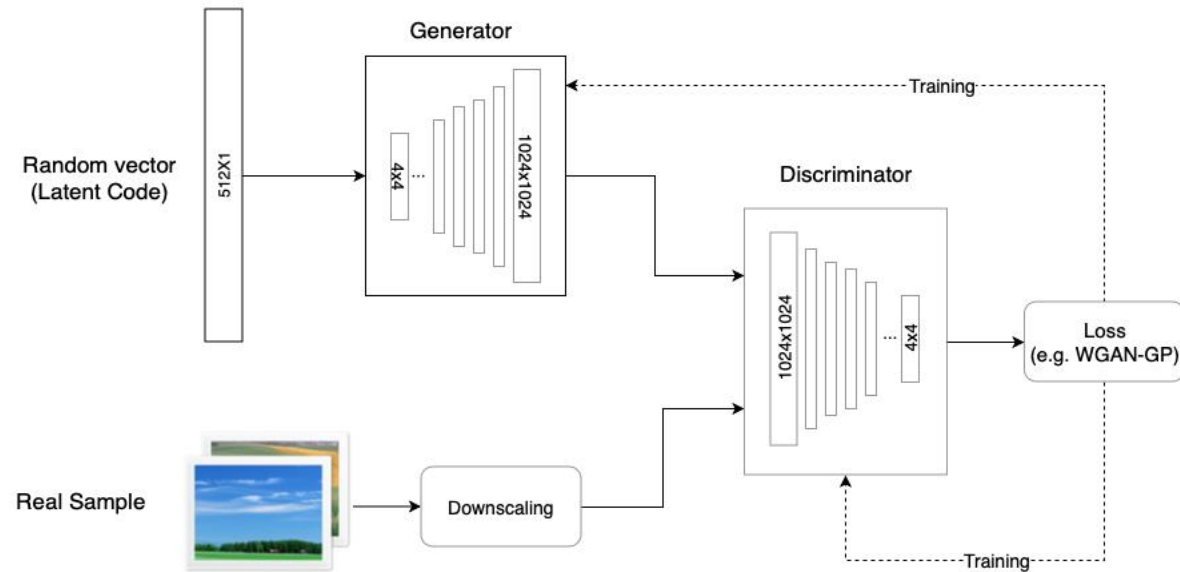
Progressive GAN <https://arxiv.org/pdf/1710.10196.pdf>

# StyleGAN



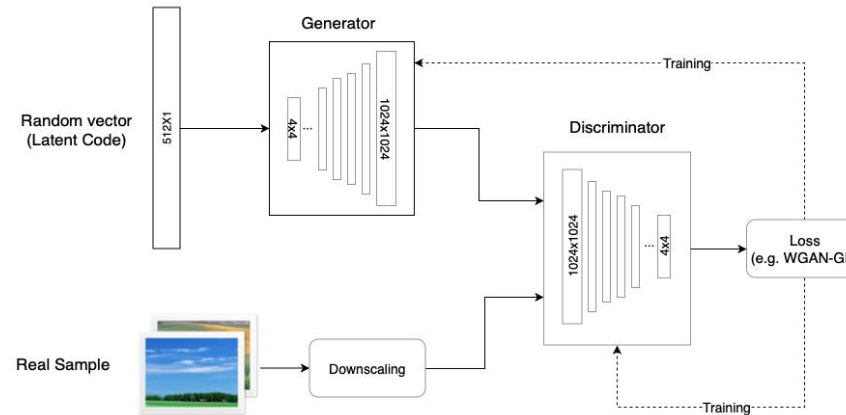
Progressive GAN <https://arxiv.org/pdf/1710.10196.pdf>

# StyleGAN



Progressive GAN <https://arxiv.org/pdf/1710.10196.pdf>

# StyleGAN




1. Coarse - resolution of up to 82 - affects pose, general hair style, face shape, etc
2. Middle - resolution of 162 to 322 - affects finer facial features, hair style, eyes open/closed, etc.
3. Fine - resolution of 642 to 10242 - affects color scheme (eye, hair and skin) and micro features.

Karras, Tero, Samuli Laine, and Timo Aila. "A style-based generator architecture for generative adversarial networks." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401-4410. 2019. <https://arxiv.org/pdf/1812.04948.pdf>

<https://towardsdatascience.com/explained-a-style-based-generator-architecture-for-gans-generating-and-tuning-realistic-6cb2be0f431>

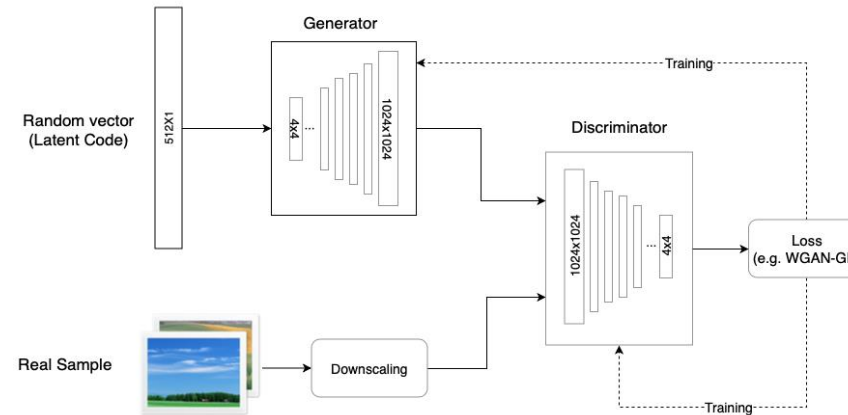


# StyleGAN




1. **Coarse** - resolution of **up to 82** - affects pose, general hair style, face shape, etc
2. **Middle** - resolution of **162 to 322** - affects finer facial features, hair style, eyes open/closed, etc.
3. **Fine** - resolution of **642 to 10242** - affects color scheme (eye, hair and skin) and micro features.

# StyleGAN



1. **Coarse** - resolution of **up to 82** - affects **pose, general hair style, face shape**, etc
2. **Middle** - resolution of **162 to 322** - affects **finer facial features, hair style, eyes open/closed**, etc.
3. **Fine** - resolution of **642 to 10242** - affects **color scheme** (eye, hair and skin) and micro features.

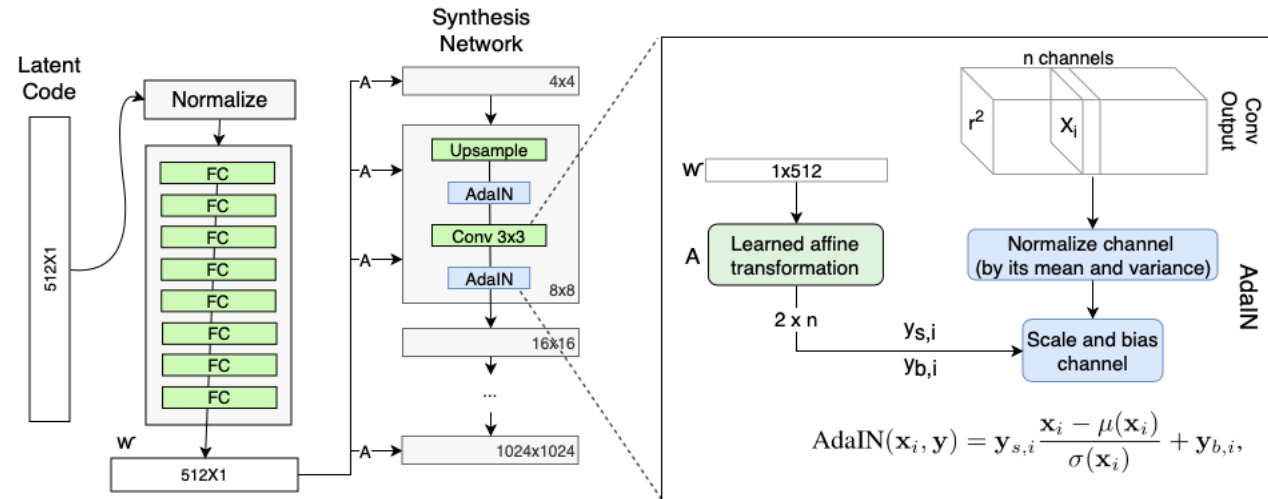
# StyleGAN



Q: can we control these styles?

1. **Coarse** - resolution of **up to 82** - affects **pose, general hair style, face shape, etc**
2. **Middle** - resolution of **162 to 322** - affects **finer facial features, hair style, eyes open/closed, etc.**
3. **Fine** - resolution of **642 to 10242** - affects **color scheme** (eye, hair and skin) and micro features.

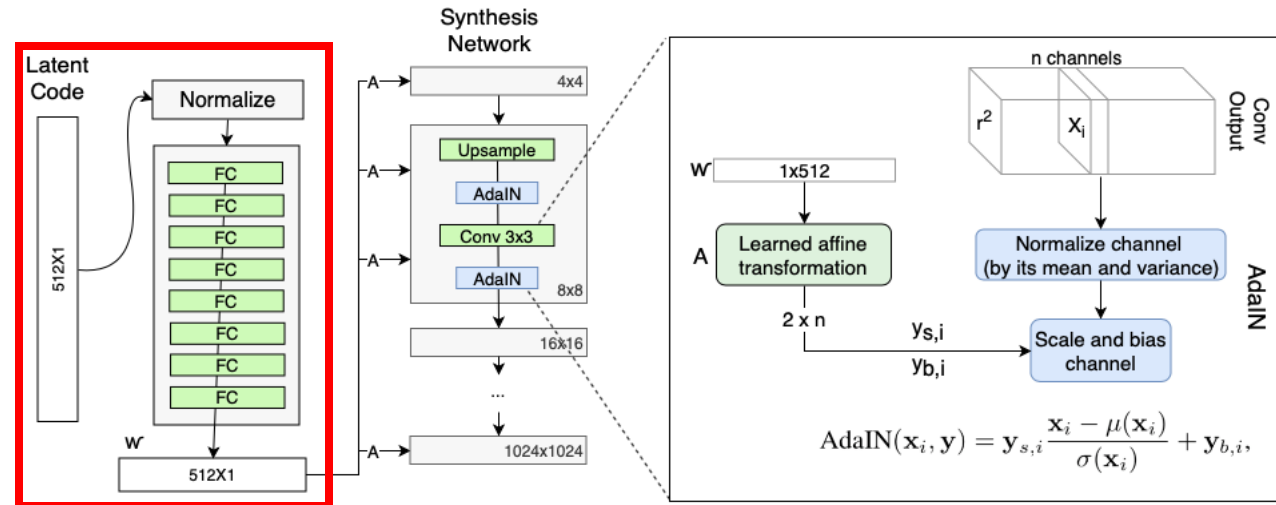
# StyleGAN



Karras, Tero, Samuli Laine, and Timo Aila. "A style-based generator architecture for generative adversarial networks." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401-4410. 2019. <https://arxiv.org/pdf/1812.04948.pdf>

<https://towardsdatascience.com/explained-a-style-based-generator-architecture-for-gans-generating-and-tuning-realistic-6cb2be0f431>

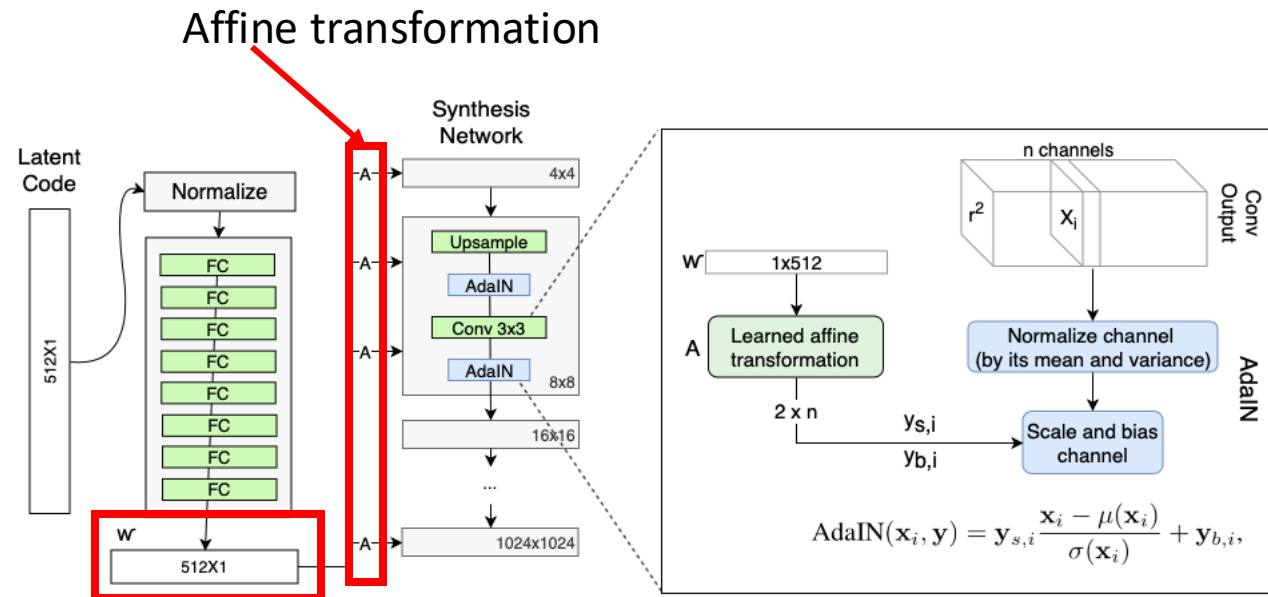
# StyleGAN



Karras, Tero, Samuli Laine, and Timo Aila. "A style-based generator architecture for generative adversarial networks." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401-4410. 2019. <https://arxiv.org/pdf/1812.04948.pdf>

<https://towardsdatascience.com/explained-a-style-based-generator-architecture-for-gans-generating-and-tuning-realistic-6cb2be0f431>

# StyleGAN



Karras, Tero, Samuli Laine, and Timo Aila. "A style-based generator architecture for generative adversarial networks." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401-4410. 2019. <https://arxiv.org/pdf/1812.04948.pdf>

<https://towardsdatascience.com/explained-a-style-based-generator-architecture-for-gans-generating-and-tuning-realistic-6cb2be0f431>

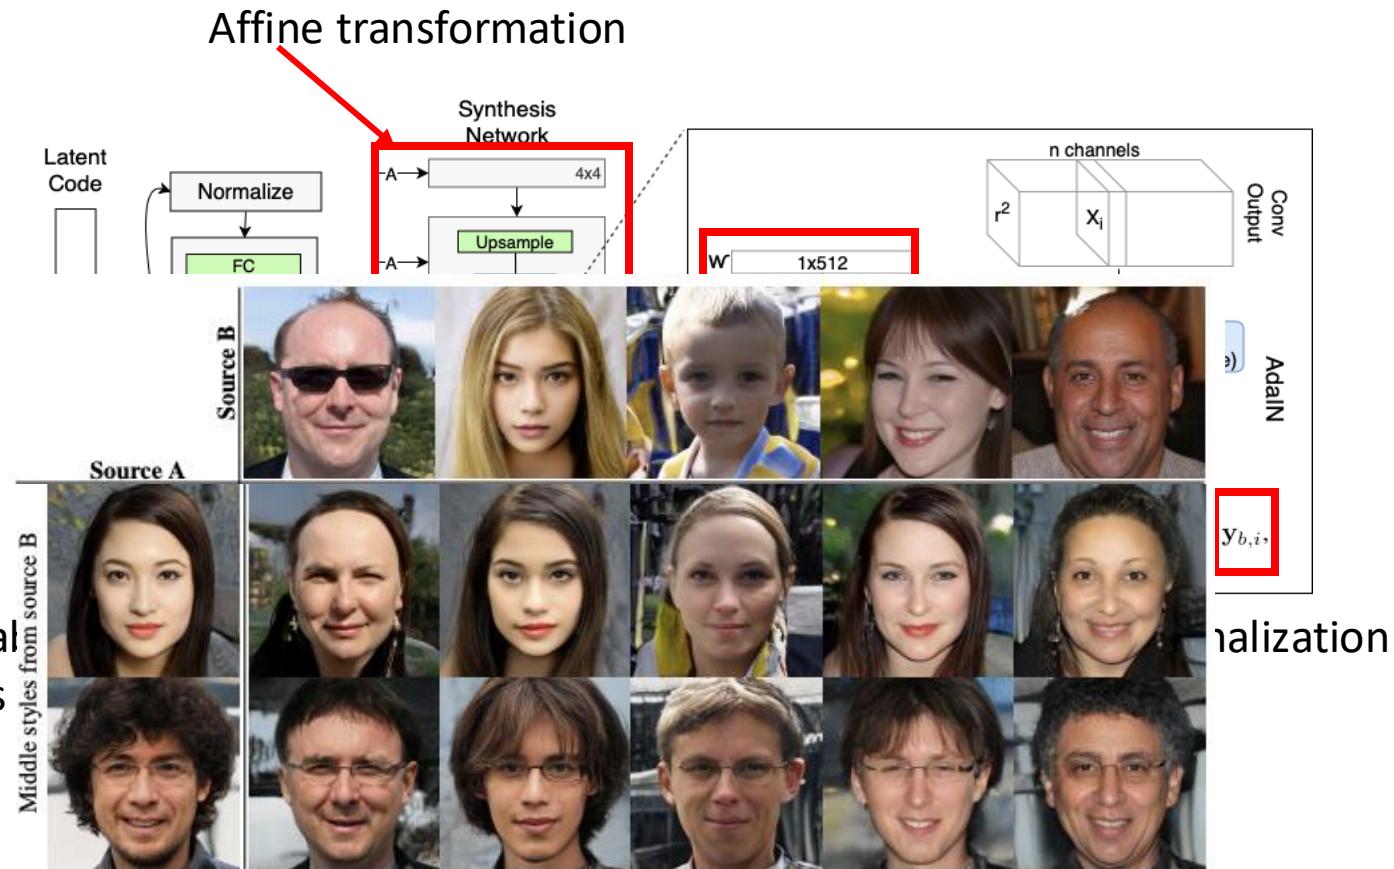






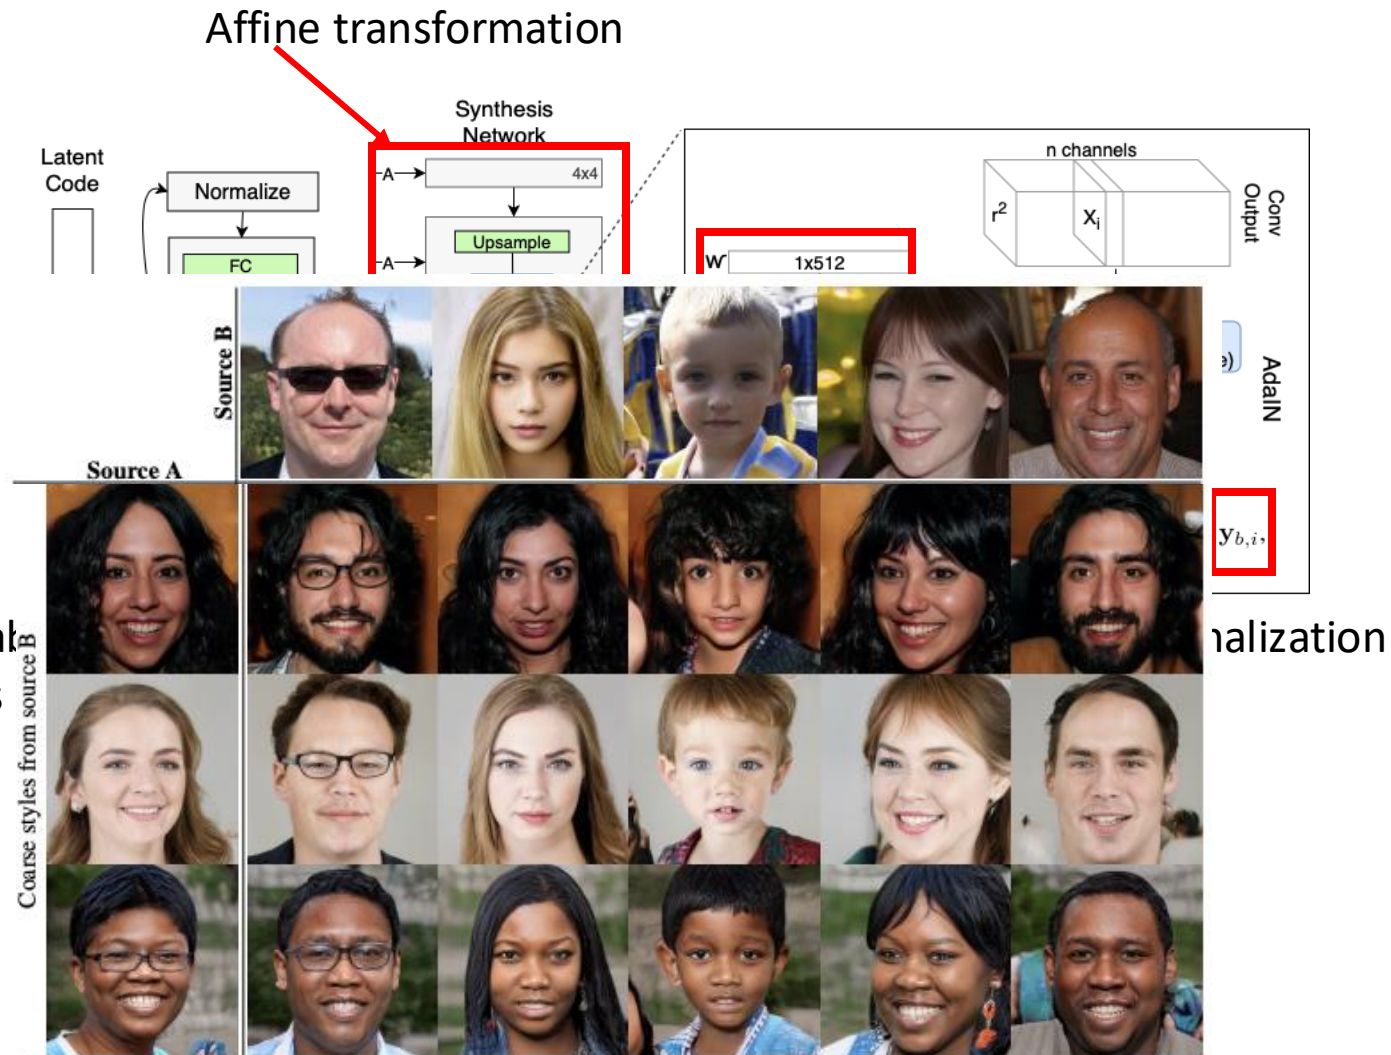


# StyleGAN



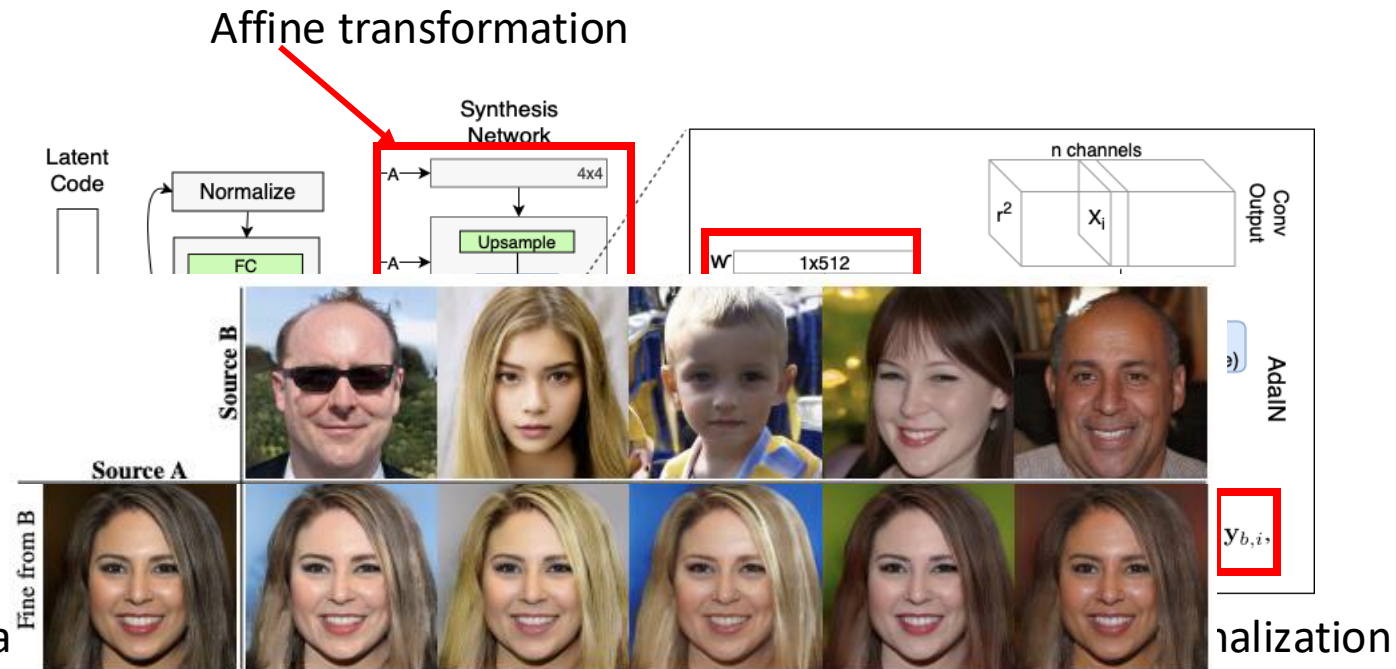
Style mixing: what all  
recourses at various

# StyleGAN



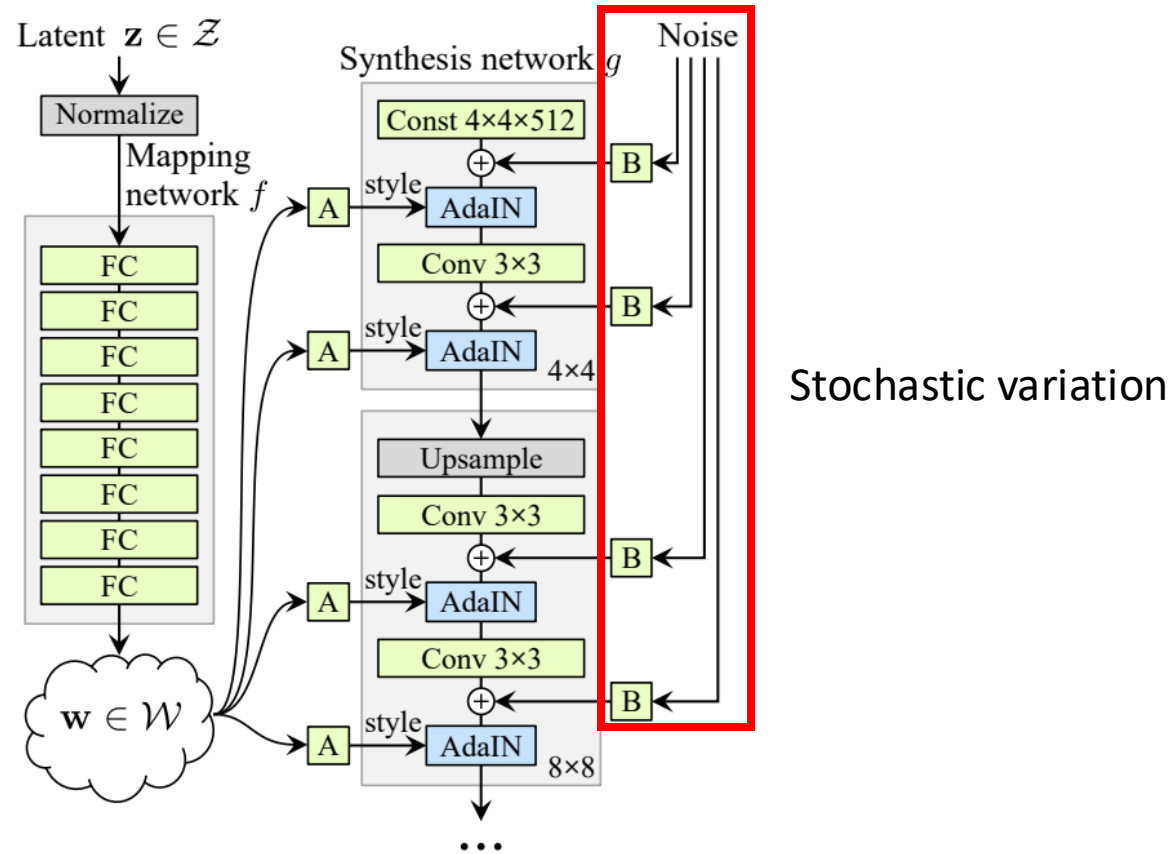
Style mixing: what all  
recourses at various

# StyleGAN



**Style mixing:** what a recoursees at various layers?

# StyleGAN



(b) Style-based generator



# StyleGAN



(a) Generated image (b) Stochastic variation (c) Standard deviation

...

(b) Style-based generator