

# Deep Convolutional Generative Adversarial Networks (DCGAN)

David I. Inouye

Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.

# DCGAN: Randomly generated bedrooms show slightly odd but almost realistic bedrooms

Original GAN (CIFAR10)



DCGAN (bedrooms)



Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.

# DCGAN can show interpolation between imaginary hotel rooms



Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.

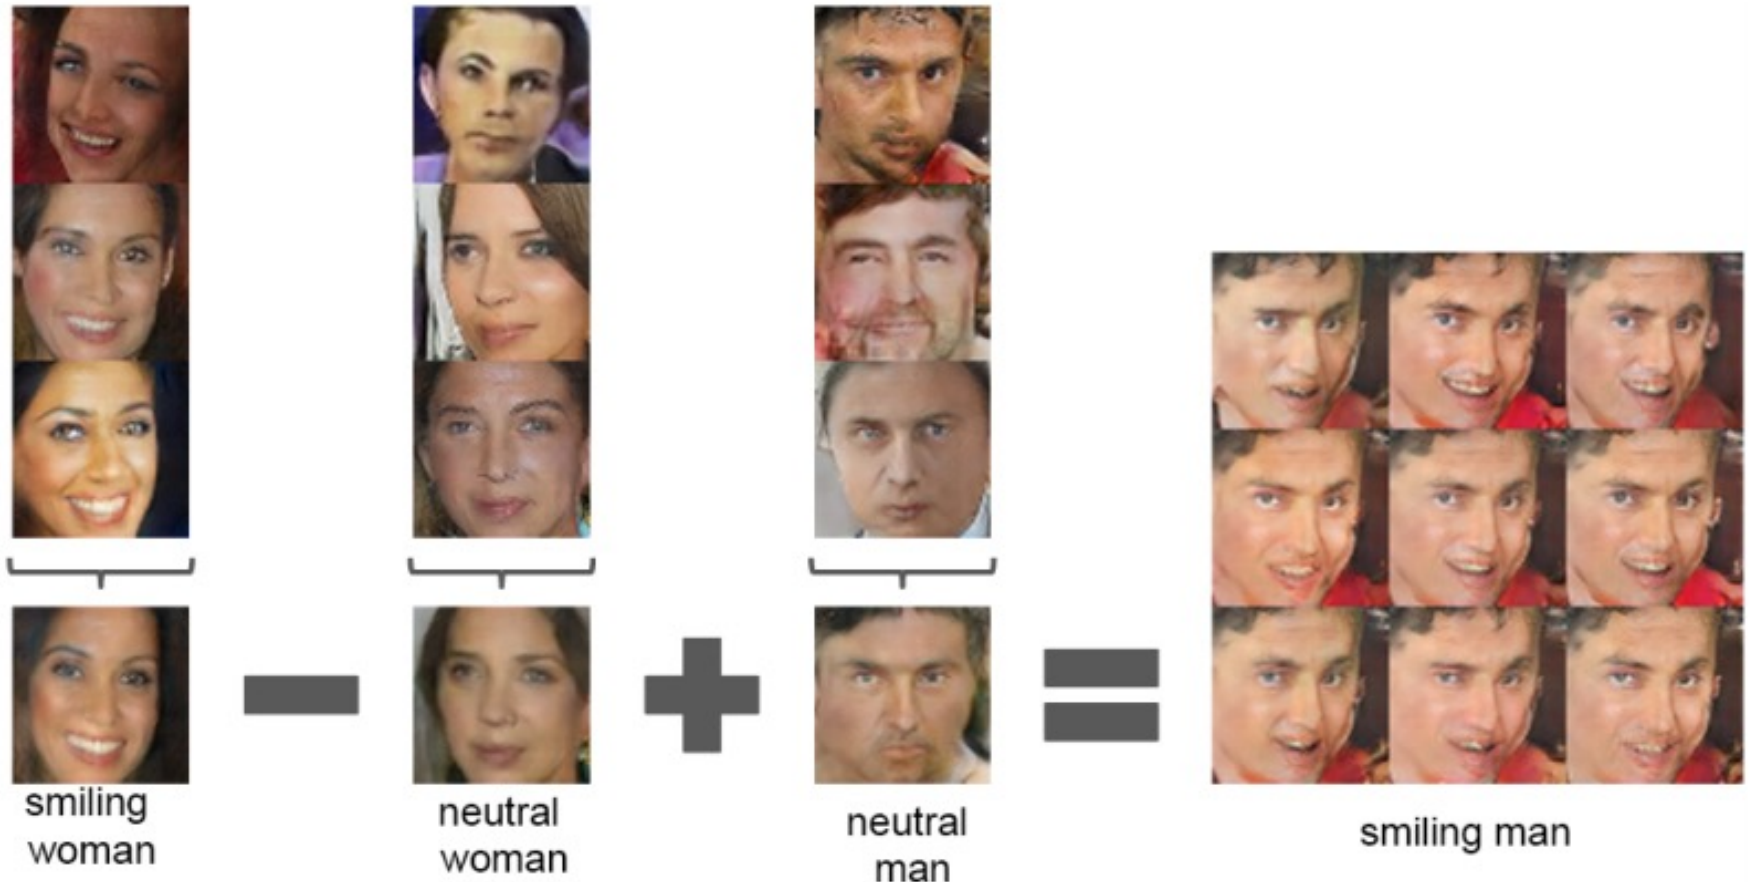
# Removing certain filters can modify the generated images (in this case, a “window” filter)



Figure 6: Top row: un-modified samples from model. Bottom row: the same samples generated with dropping out “window” filters. Some windows are removed, others are transformed into objects with similar visual appearance such as doors and mirrors. Although visual quality decreased, overall scene composition stayed similar, suggesting the generator has done a good job disentangling scene representation from object representation. Extended experiments could be done to remove other objects from the image and modify the objects the generator draws.

Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.

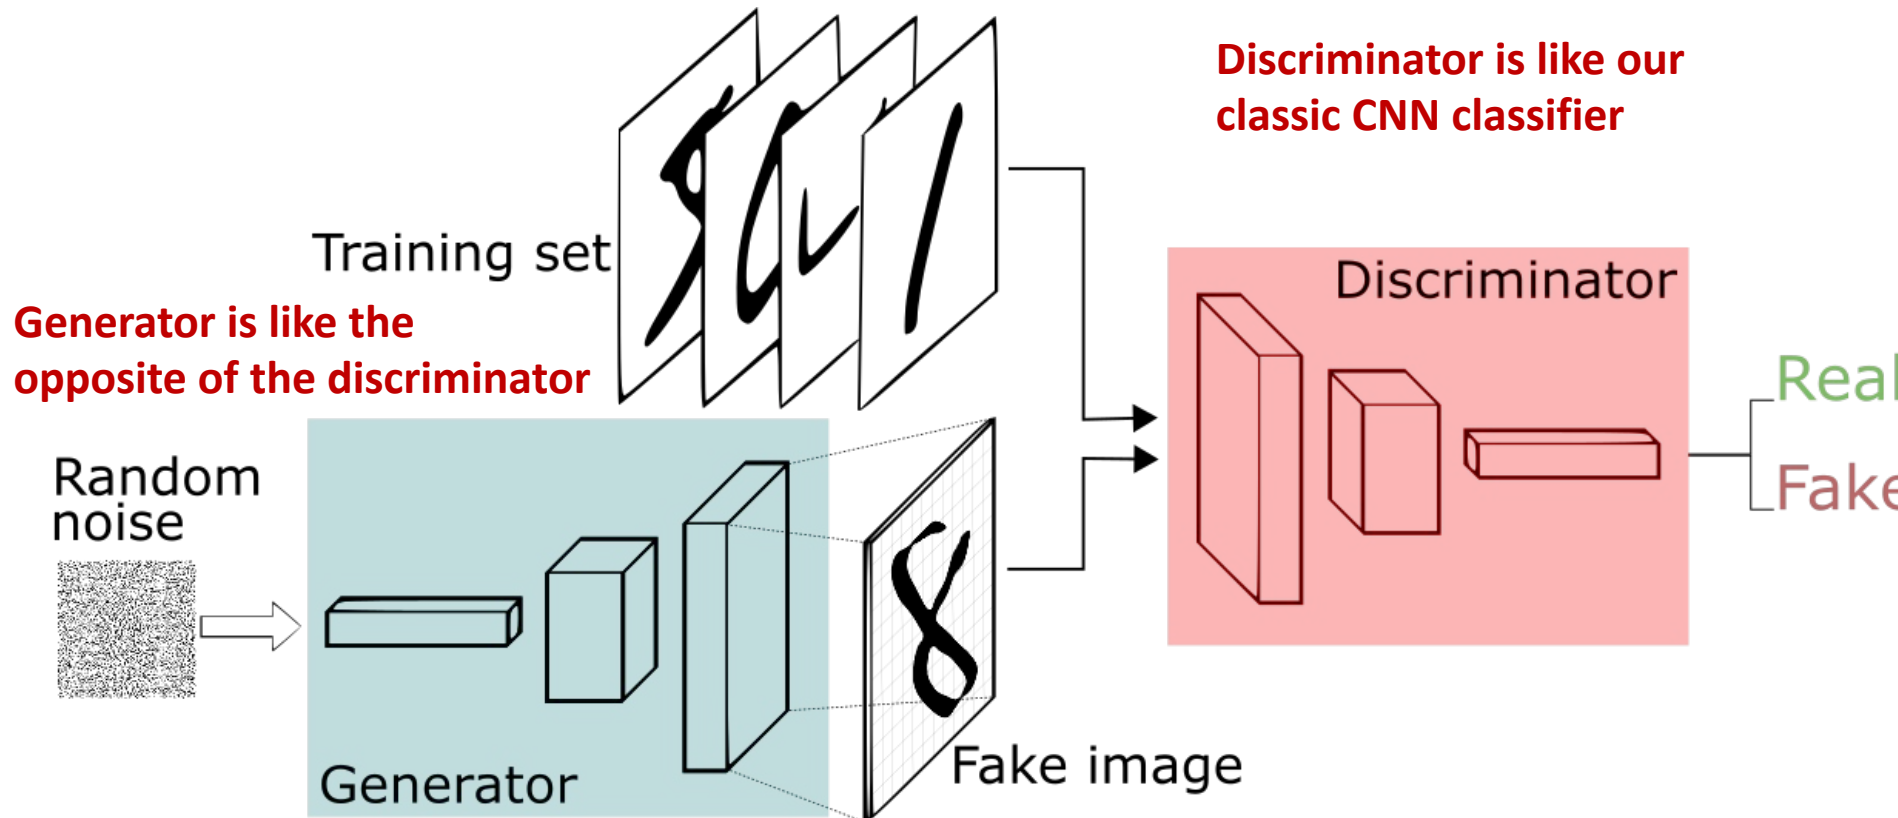
# Simple vector arithmetic in latent space of DCGAN can generate new faces



Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.

# DCGAN Architecture:

Two deep networks that look like opposites



<https://www.freecodecamp.org/news/an-intuitive-introduction-to-generative-adversarial-networks-gans-7a2264a81394/>

# DCGAN generator upsamples the size of the image in multiple stages




Figure 1: DCGAN generator used for LSUN scene modeling. A 100 dimensional uniform distribution  $Z$  is projected to a small spatial extent convolutional representation with many feature maps. A series of four fractionally-strided convolutions (in some recent papers, these are wrongly called deconvolutions) then convert this high level representation into a  $64 \times 64$  pixel image. Notably, no fully connected or pooling layers are used.

DCGAN requires transposed convolutions, BatchNorm, and a few other training tricks

- ▶ Transposed convolutions (for upsampling)
- ▶ BatchNorm (for stabilizing training)
- ▶ A few 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.





Transposed convolution can be used to **upsample** an tensor/image to have higher dimensions

- ▶ Also known as:
  - ▶ Fractionally-strided convolution
  - ▶ Improperly, deconvolution
- ▶ Remember: Convolution is like matrix multiplication

$$y = x (*) f \Leftrightarrow \text{vec}(y) = A_f \text{vec}(x)$$


- ▶ Transpose convolution is the transpose of  $A_f$ :
$$\text{vec}(y) = A_f^T \text{vec}(x)$$

# Convolution operator with corresponding matrix




[https://github.com/naokishibuya/deep-learning/blob/master/python/transposed\\_convolution.ipynb](https://github.com/naokishibuya/deep-learning/blob/master/python/transposed_convolution.ipynb)

# Transposed convolution operator with corresponding matrix



[https://github.com/naokishibuya/deep-learning/blob/master/python/transposed\\_convolution.ipynb](https://github.com/naokishibuya/deep-learning/blob/master/python/transposed_convolution.ipynb)


Transposed convolution can be **equivalent** to a simple convolution with zero rows/columns added (added zeros simulate fractional strides)



Original input

Inputs (2, 2)


Zero-padded input



Inputs (6, 6)

(\*)


Kernel



Kernel (3, 3)

=

Output



Output (4, 4)

# Two GAN training stability tricks

1. Do backwards (gradient calc) separately for fake and real batches because of BatchNorm layers

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

- ▶ Compute  $\nabla_D \mathbb{E}_{p_{data}} [\log D(x)]$  and  $\nabla_D \mathbb{E}_{\epsilon} \left[ \log \left( 1 - D(G(z)) \right) \right]$  **separately**

2. Use modified objective for generator training from original GAN paper

$$\nabla_G \mathbb{E}_{\epsilon} [-\log D(G(z))] \text{ instead of } \nabla_G \mathbb{E}_{\epsilon} \left[ \log \left( 1 - D(G(z)) \right) \right]$$

- ▶ See DCGAN demo on MNIST

# Resources for GANs

- ▶ DCGAN Tutorial

- ▶ [https://pytorch.org/tutorials/beginner/dcgan\\_faces\\_tutorial.html](https://pytorch.org/tutorials/beginner/dcgan_faces_tutorial.html)

- ▶ GAN training tips/hacks

- ▶ <https://github.com/soumith/ganhacks>

- ▶ GAN common problems

- ▶ <https://developers.google.com/machine-learning/gan/problems>