Generative Adversarial Networks

Stephan Halbritter

June 13, 2017

Hamburg University of Applied Science

Outline

1. Basic idea
2. Generative Adversarial Nets (Goodfellow et al. 2014)
3. Deep Convolutional GAN (Radford et al. 2015)
4. Examples of current research
5. Future work and resources

Basic idea

Goal

Generate “realistic-looking” data

Intuition

Interpret images as samples from a high-dimensional probability distribution

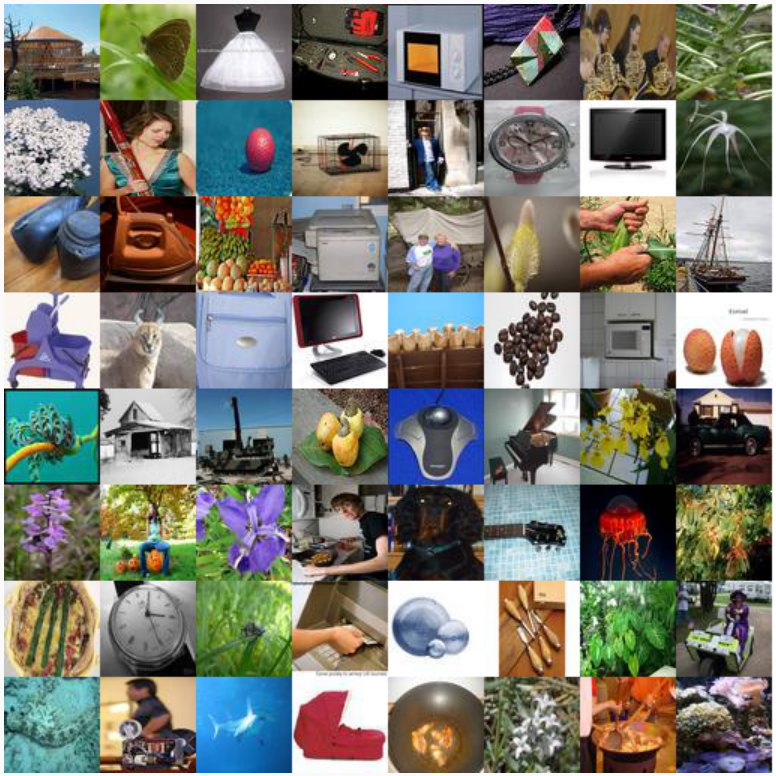

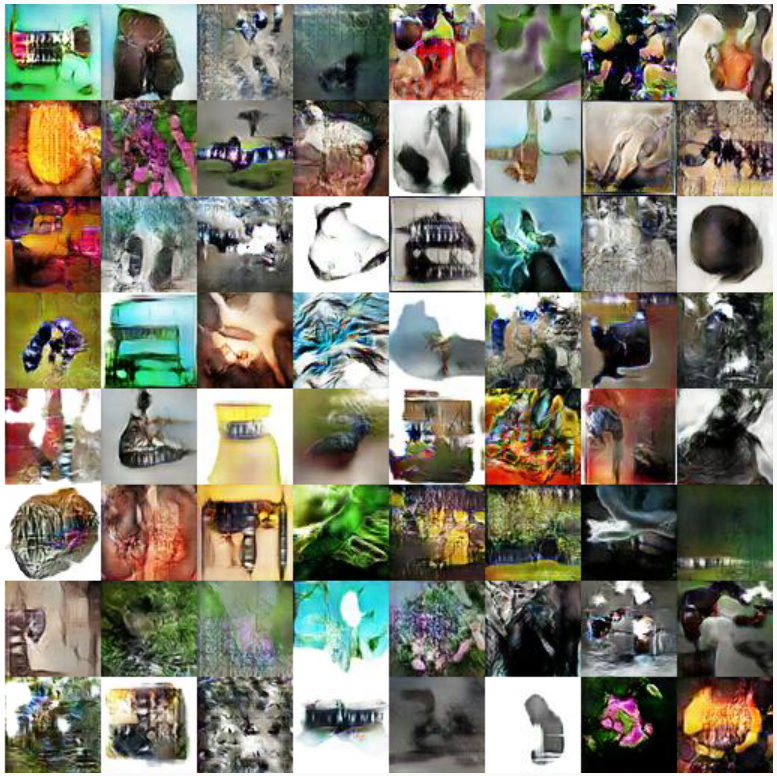

Source: kvfrans.com/generative-adversial-networks-explained/

Other approaches

- Restricted Bolzmann Machines *(RBM)*
- Fully visible belief networks *(FVBN)*
- Variational autoencoders *(VAE)*

Generative Adversarial Nets (Goodfellow et al. 2014)

Generative Adversarial Networks (Goodfellow et al. 2014)

- presented at the *Neural Information Processing Systems* (NIPS) conference 2014
- proposes the framework
- ground-breaking paper

Papers naming GAN variants

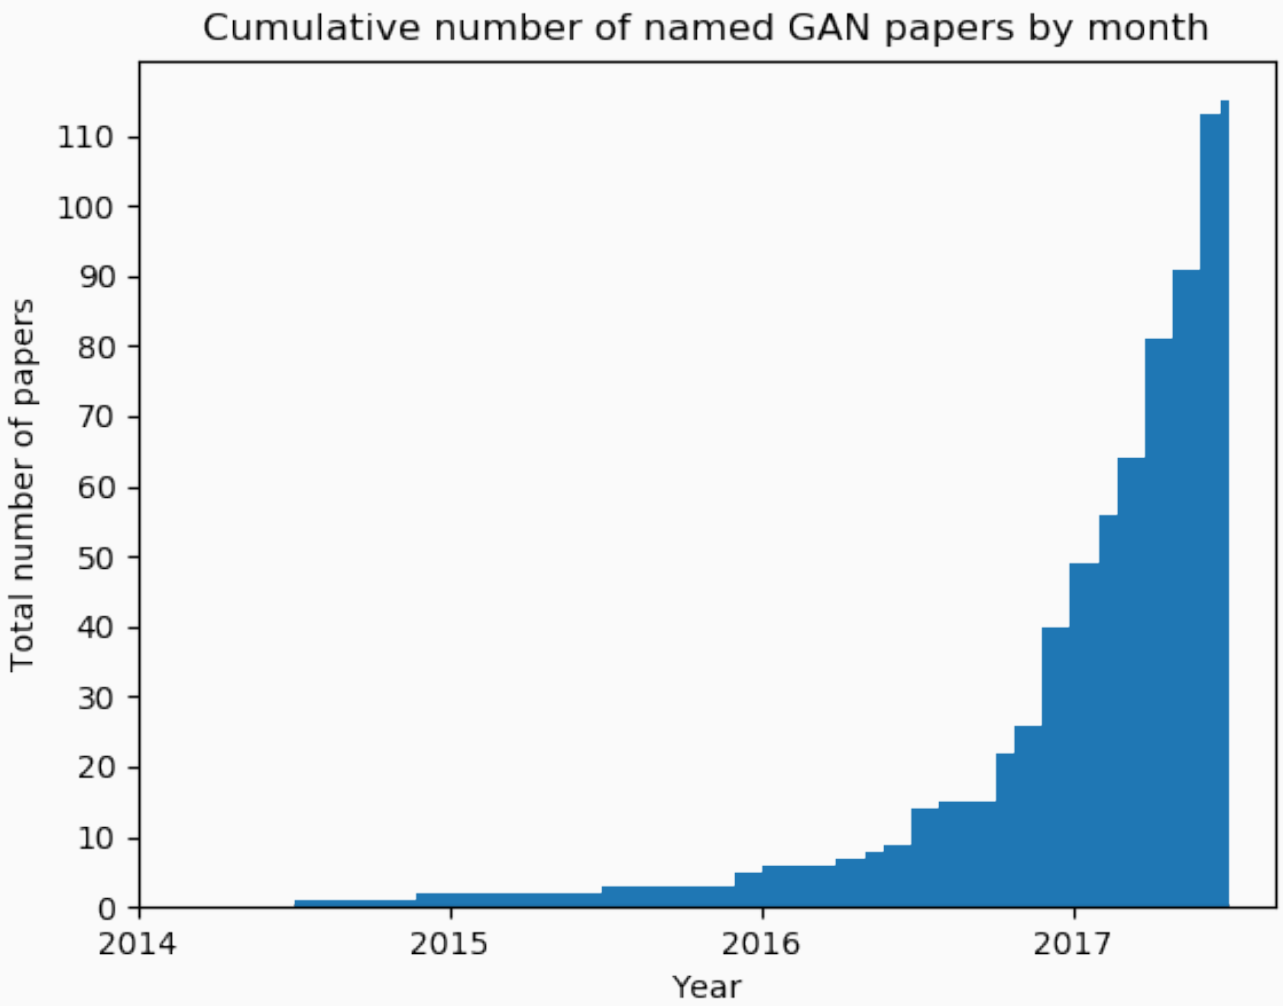

Source: github.com/hindupuravinash/the-gan-zoo as at June 10, 2017

Basic approach

Set up a game between two neural nets

- the generator creates samples
- the discriminator classifies these samples as *real* or *fake*
- both train each other

Overview

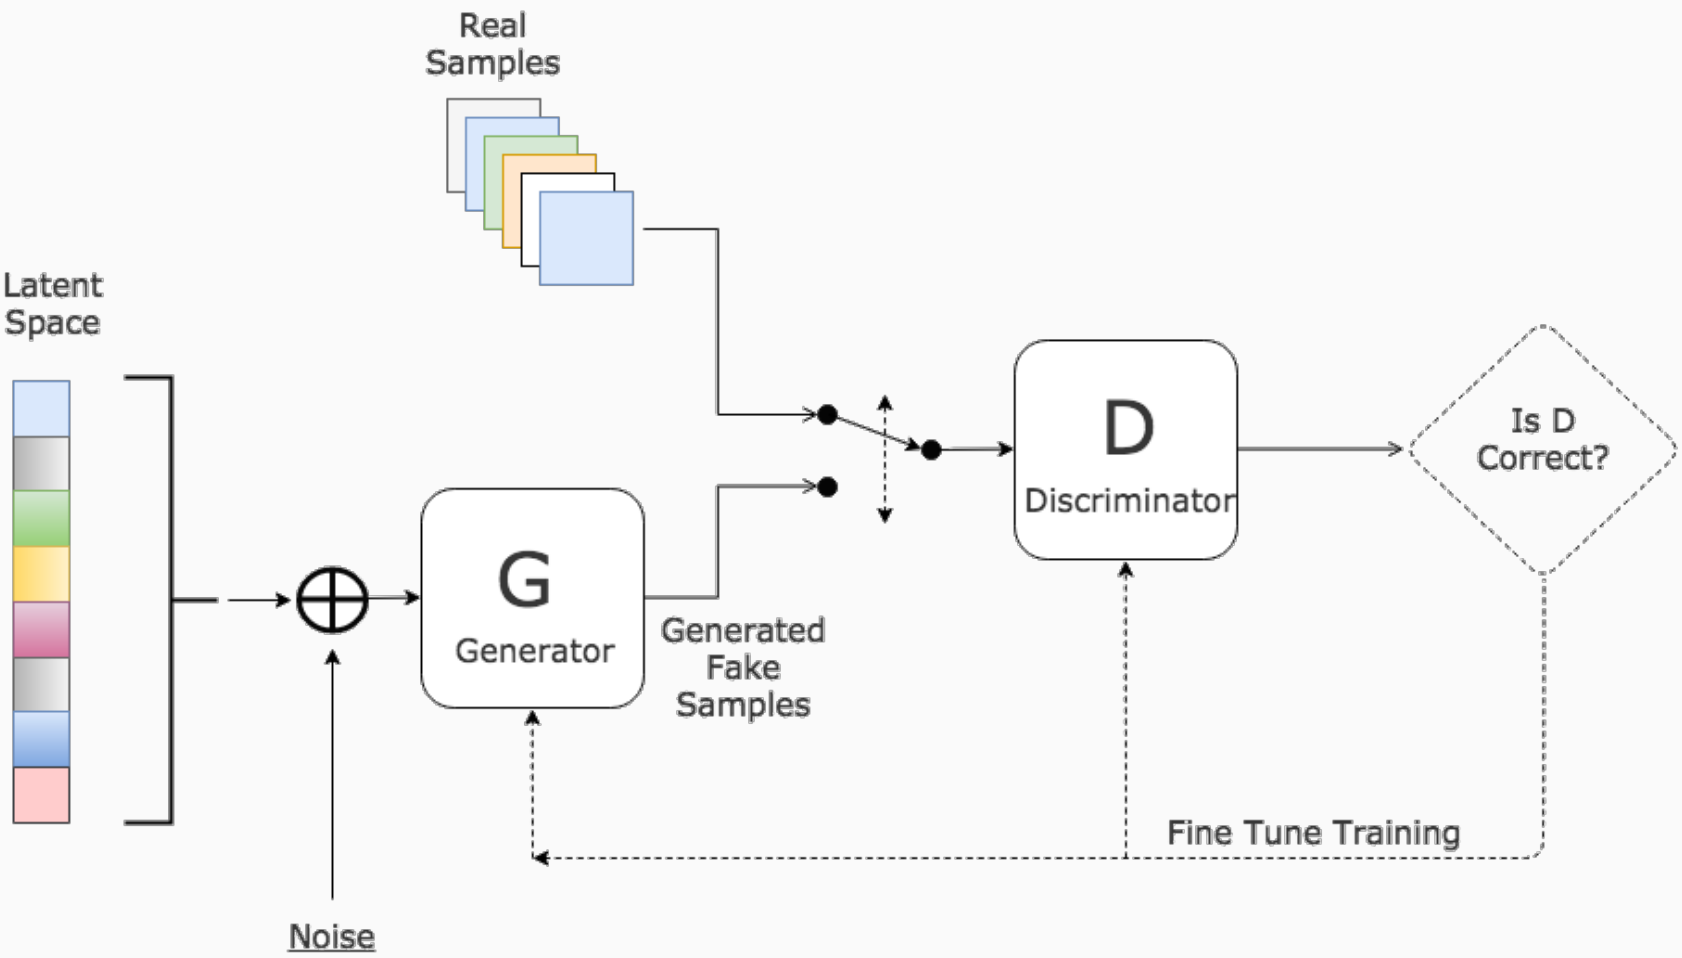

Source:

www.kdnuggets.com/2017/01/generative-adversarial-networks-hot-topic-machine-learning.html

Deep Convolutional GAN (Radford et al. 2015)

Deep Convolutional GAN (Radford et al. 2015)

Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks

- presented at ICLR 2016
- stabilizes GAN with architectural constraints
- combines CNNs and GANs

Discriminator

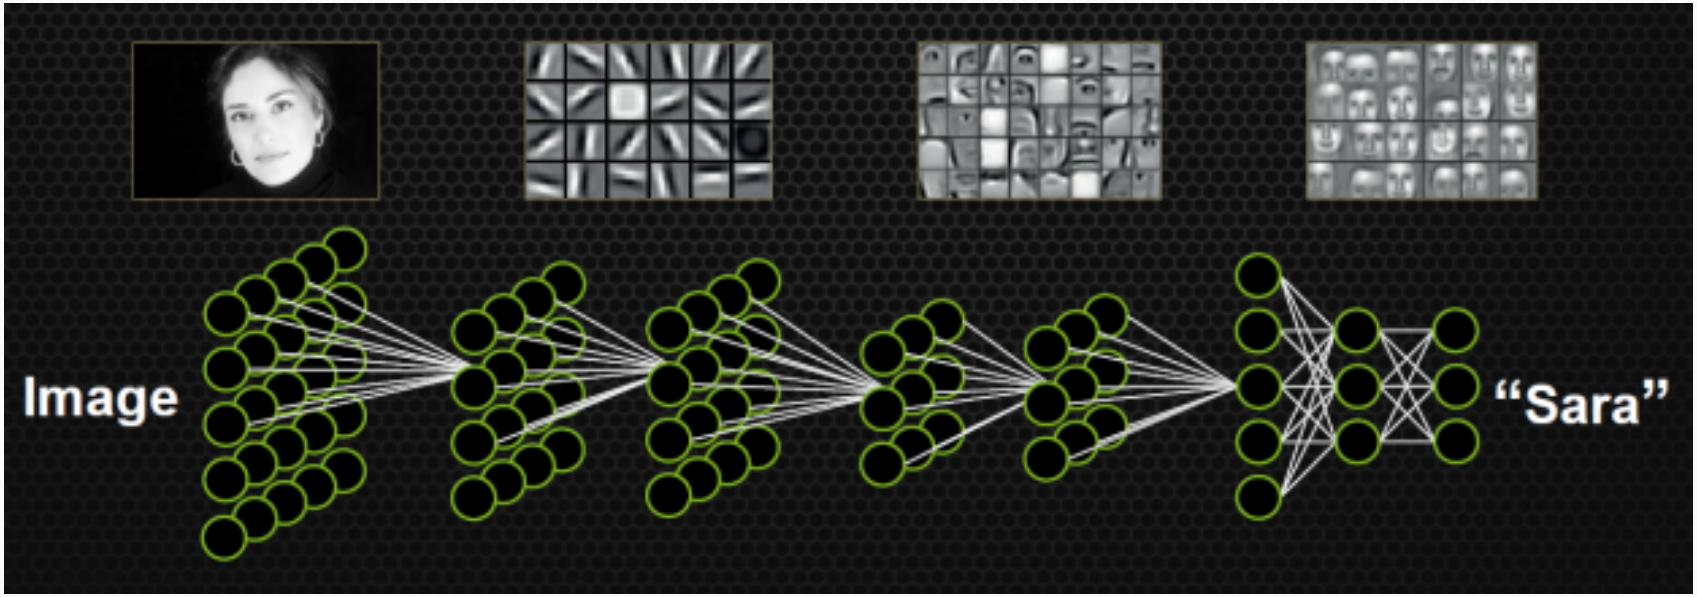

Source: https://devblogs.nvidia.com/parallelforall/accelerate-machine-learning-cudnn-deep-neural-network-library/

Generator

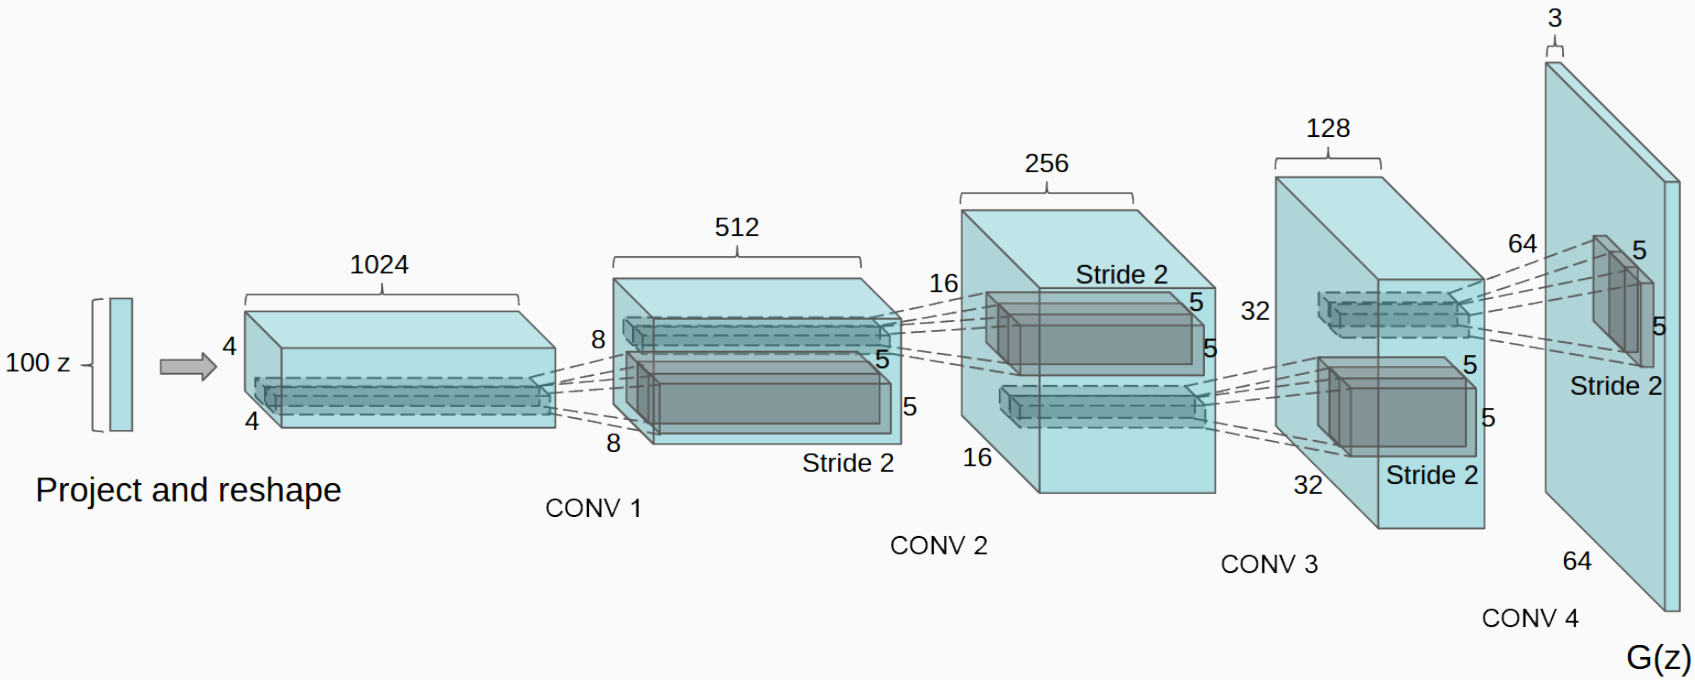

Source: Radford et al. (2015)

Upsampling with fractionally-strided convolution

Convolving a 3×3 kernel over a 3×3 input (with 1 zero inserted between inputs) padded with a 1×1 border of zeros using unit strides.

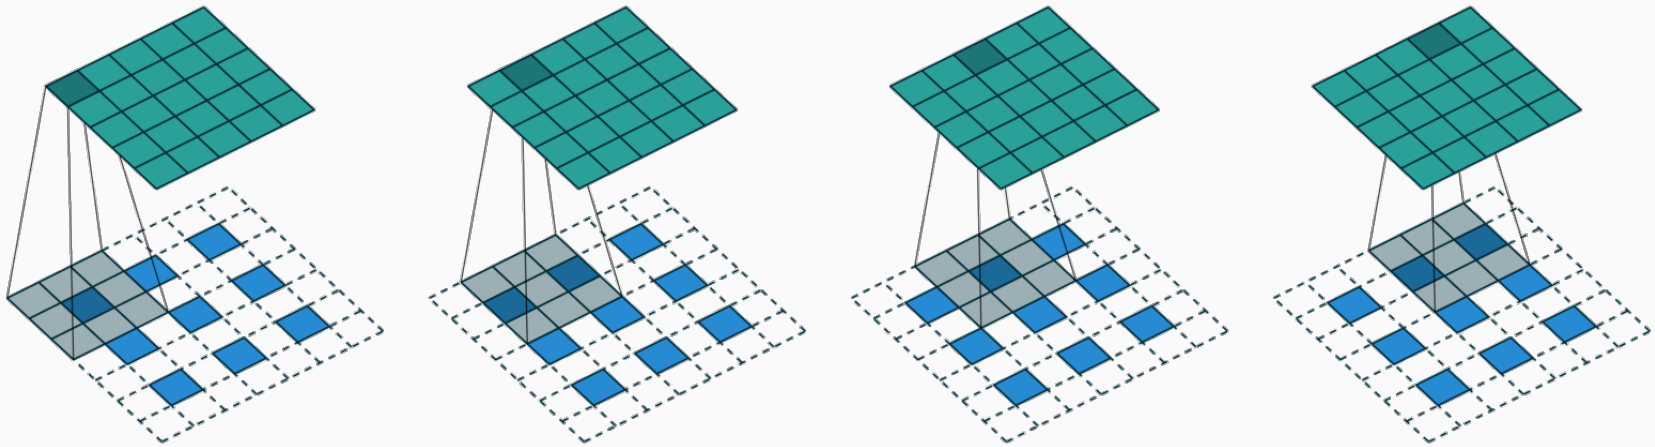

Source: Dumoulin and Visin (2016)

Training approach

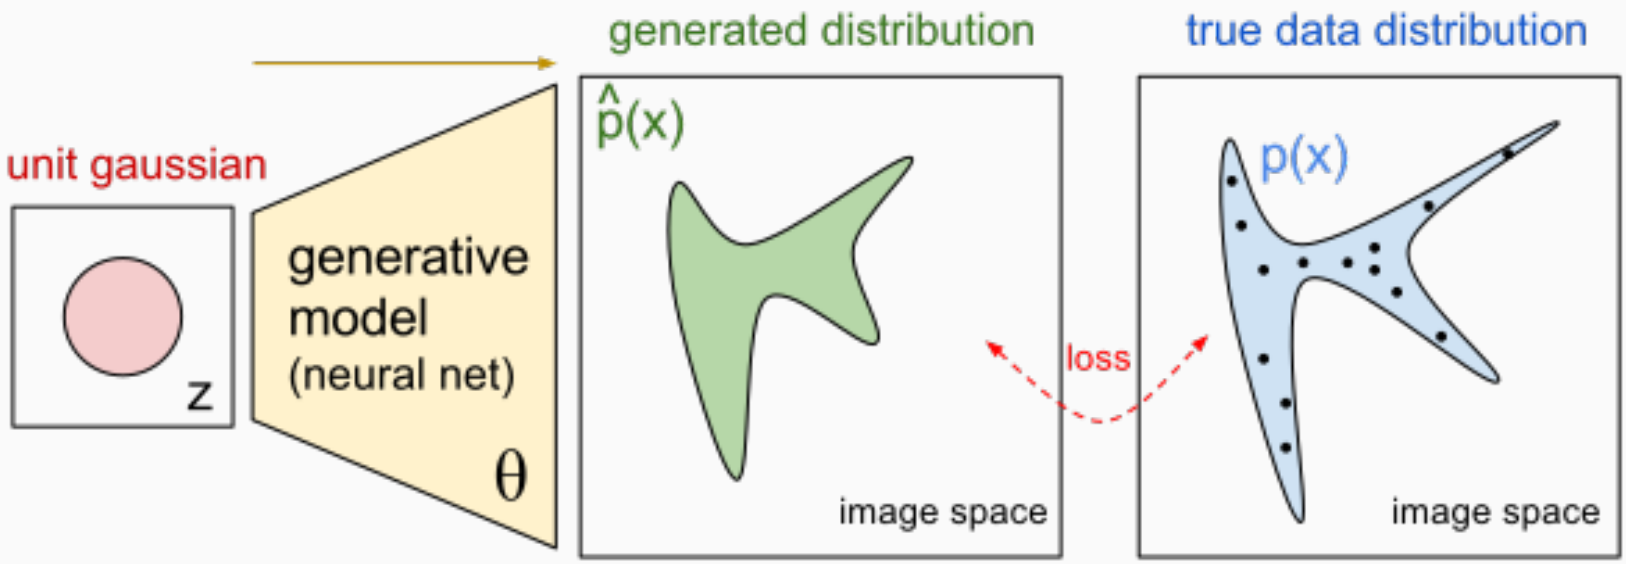

Source: blog.openai.com/generative-models/

Minimax game

$$\min_G \max_D V(D, G) = \overbrace{\mathbb{E}_{x \sim p_{\text{data}}}[\log D(x)] + \mathbb{E}_{z \sim p_z}[\log(1 - D(G(z)))]}^{\text{x is a real image}}$$

z is a noise vector

Training

for number of training iterations **do**

  for k steps **do**

  - Sample minibatch $\{z^{(1)}, \ldots, z^{(m)}\}$ from $p_g(z)$.
  - Sample minibatch $\{x^{(1)}, \ldots, x^{(m)}\}$ from $p_{\text{data}}(x)$.
  - Update D by ascending its stochastic gradient:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^{m} \left[\log D\left(x^{(i)}\right) + \log\left(1 - D\left(G\left(z^{(i)}\right)\right)\right) \right].$$

  end for

  - Sample minibatch $\{z^{(1)}, \ldots, z^{(m)}\}$ from $p_g(z)$.
  - Update G by descending its stochastic gradient:

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^{m} \log\left(1 - D\left(G\left(z^{(i)}\right)\right)\right).$$

end for

Examples of current research

DCGAN (Radford et al. 2015)

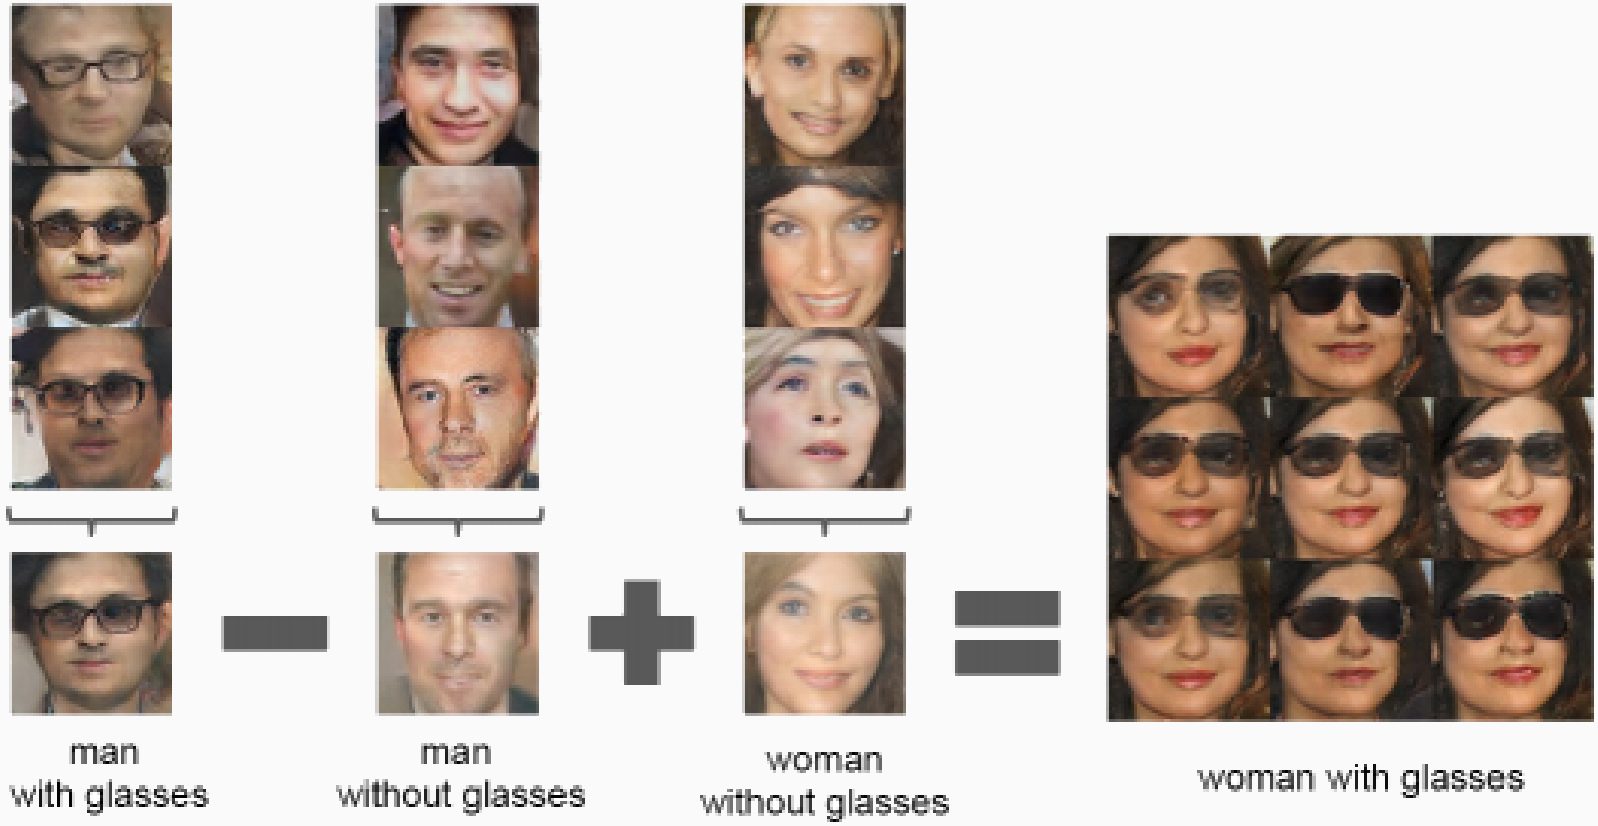

Source: Radford et al. (2015)

StackGAN (H. Zhang et al. 2016)

A small yellow bird with a black crown and a short black pointed beak

Source: H. Zhang et al. (2016)

Beyond Face Rotation (R. Huang et al. 2017)

Source: R. Huang et al. (2017)

Future work and resources

Future work

- learn more about deep learning
- better understand mathematical background
- learn a framework and create something

Helpful stuff

- arxiv-sanity.com
- blog.openai.com
- distill.pub
- offconvex.org

Conferences 2017

International Conference on Learning Representations

(ICLR) April 24 - 26 in Toulon, France.

International Conference on Machine Learning

(ICML) August 6 - 11 in Sidney, Australia.

Neural Information Processing System

(NIPS) December 4 - 9 in Long Beach, USA.

References

Dumoulin, Vincent, and **Francesco Visin**. 2016. "*A Guide to Convolution Arithmetic for Deep Learning*." *arXiv:1603.07285*.

Goodfellow, Ian, **Jean Pouget-Abadie**, **et al.** 2014. "*Generative Adversarial Networks*." *arXiv:1406.2661*.

Huang, Rui, **Shu Zhang**, **et al.** 2017. "*Beyond Face Rotation: Global and Local Perception GAN for Photorealistic and Identity Preserving Frontal View Synthesis*" arXiv:1704.04086.

Radford, Alec, **Luke Metz**, and **Soumith Chintala**. 2015. "*Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*." *CoRR* abs/1511.06434.

Zhang, Han, **Tao Xu**, **et al.** 2016. "*StackGAN: Text to Photo-Realistic Image Synthesis with Stacked Generative Adversarial Networks*." *arXiv:1612.03242*.

Any questions?