

GANocracy Tutorial

<https://bit.ly/2MknJfU>

Please check that you have completed the updated setup instructions at the link above

GANocracy

Workshop on Theory, Practice
and Artistry of Deep Generative Modeling

Event made possible by our generous sponsors:



MIT-IBM Watson AI Lab



<https://bit.ly/2MknJfU>

Democratizing GANs



Image credit: goodfellow_ian (2019) "4.5 years of GAN progress on face generation." [Tweet]

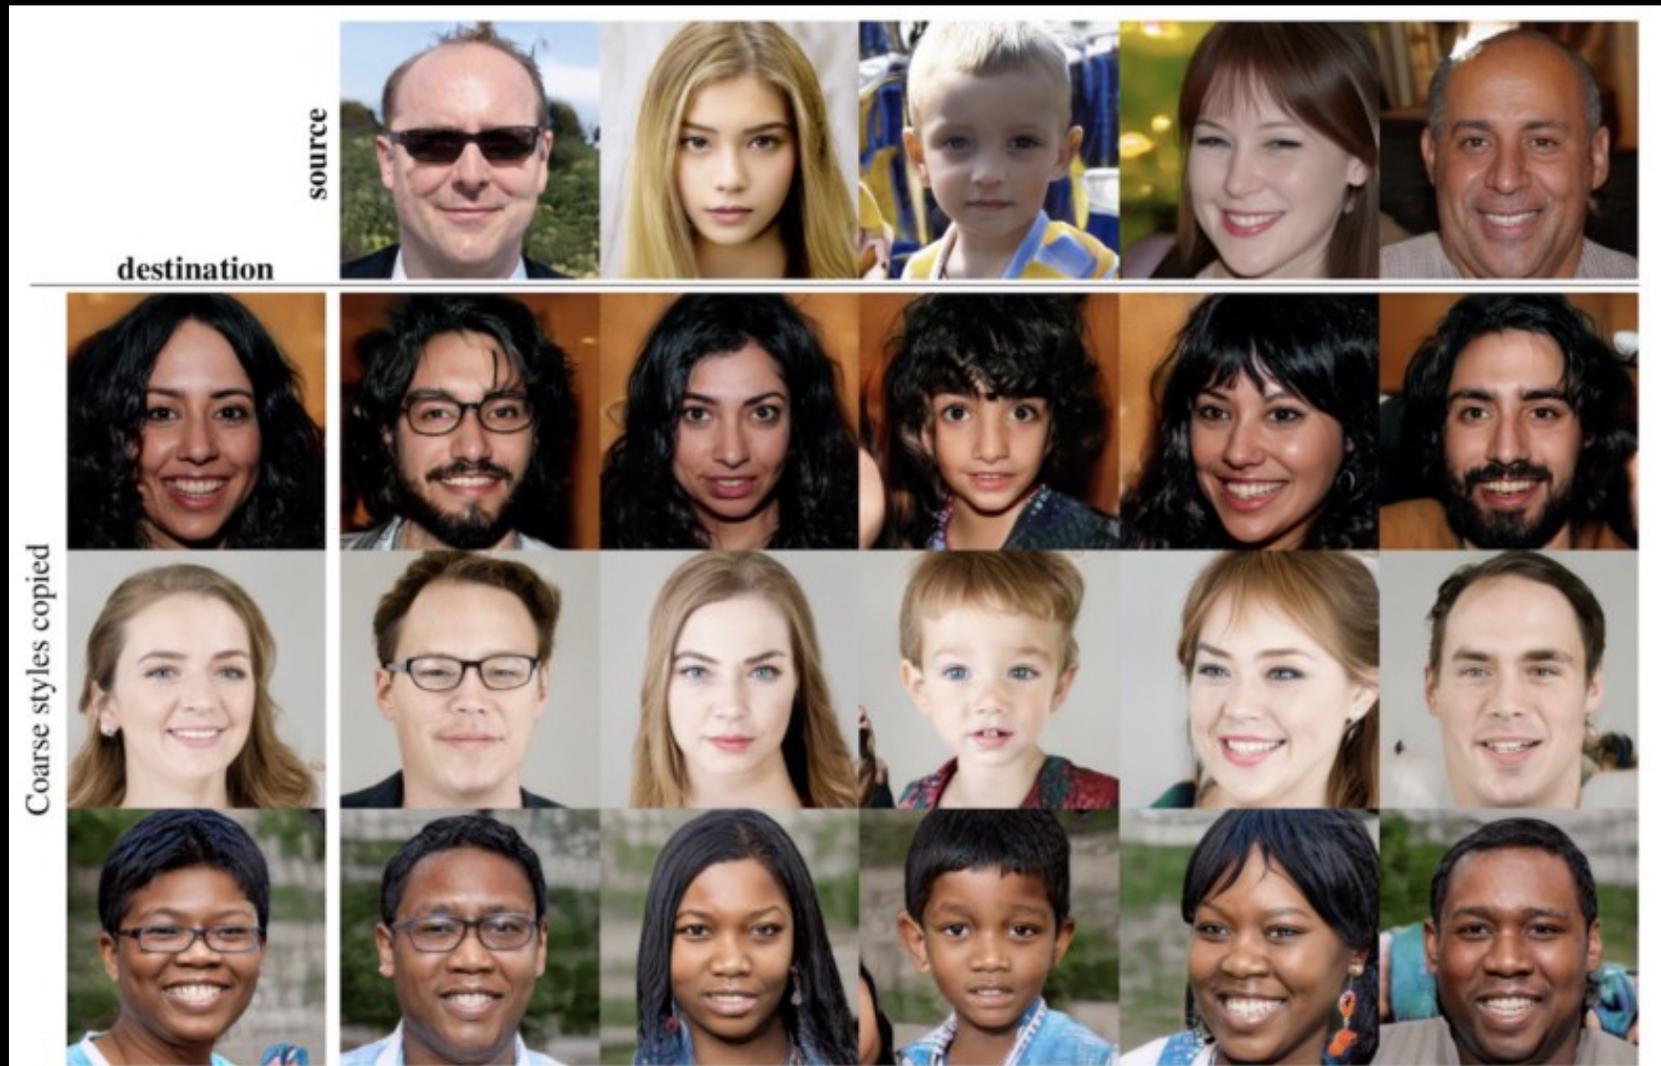


Image credit: Karras, T., Laine, S., & Aila, T. (2018). A style-based generator architecture for generative adversarial networks.

<https://bit.ly/2MknJfU>



Image credit: compiled from [Egor Zakharov].(21 May 2019) "Few-Shot Adversarial Learning of Realistic Neural Talking Head Models".

It's not just visual...

Audio

- style transfer
- sound generation

Cyber Security

- steganography
- password guessing

Data Augmentation

- signal boosting

Data Synthesis

- generate training samples

So how did we get here?

Given an observable variable X and a target variable Y :

Generative model

$$P(X, Y)$$

The joint probability
distribution on $X \times Y$

Discriminative model

$$P(Y|X = x)$$

The conditional probability
of Y , given an observation x

So how did we get here?

- 1) Generative models with a parametric specification of a probability distribution function, trained by maximizing the log likelihood [e.g. **Boltzmann machines**]



Generative models that don't need to explicitly represent the likelihood, trained by backpropagation [e.g. **Generative Stochastic Networks**] or via discriminative criterion [e.g. **Noise-contrastive Estimation**]

- 2) Pairing generator networks with other networks [e.g. **VAEs**]
- 3) Two competing neural networks [e.g. **predictability minimization**]

Generative Adversarial Nets

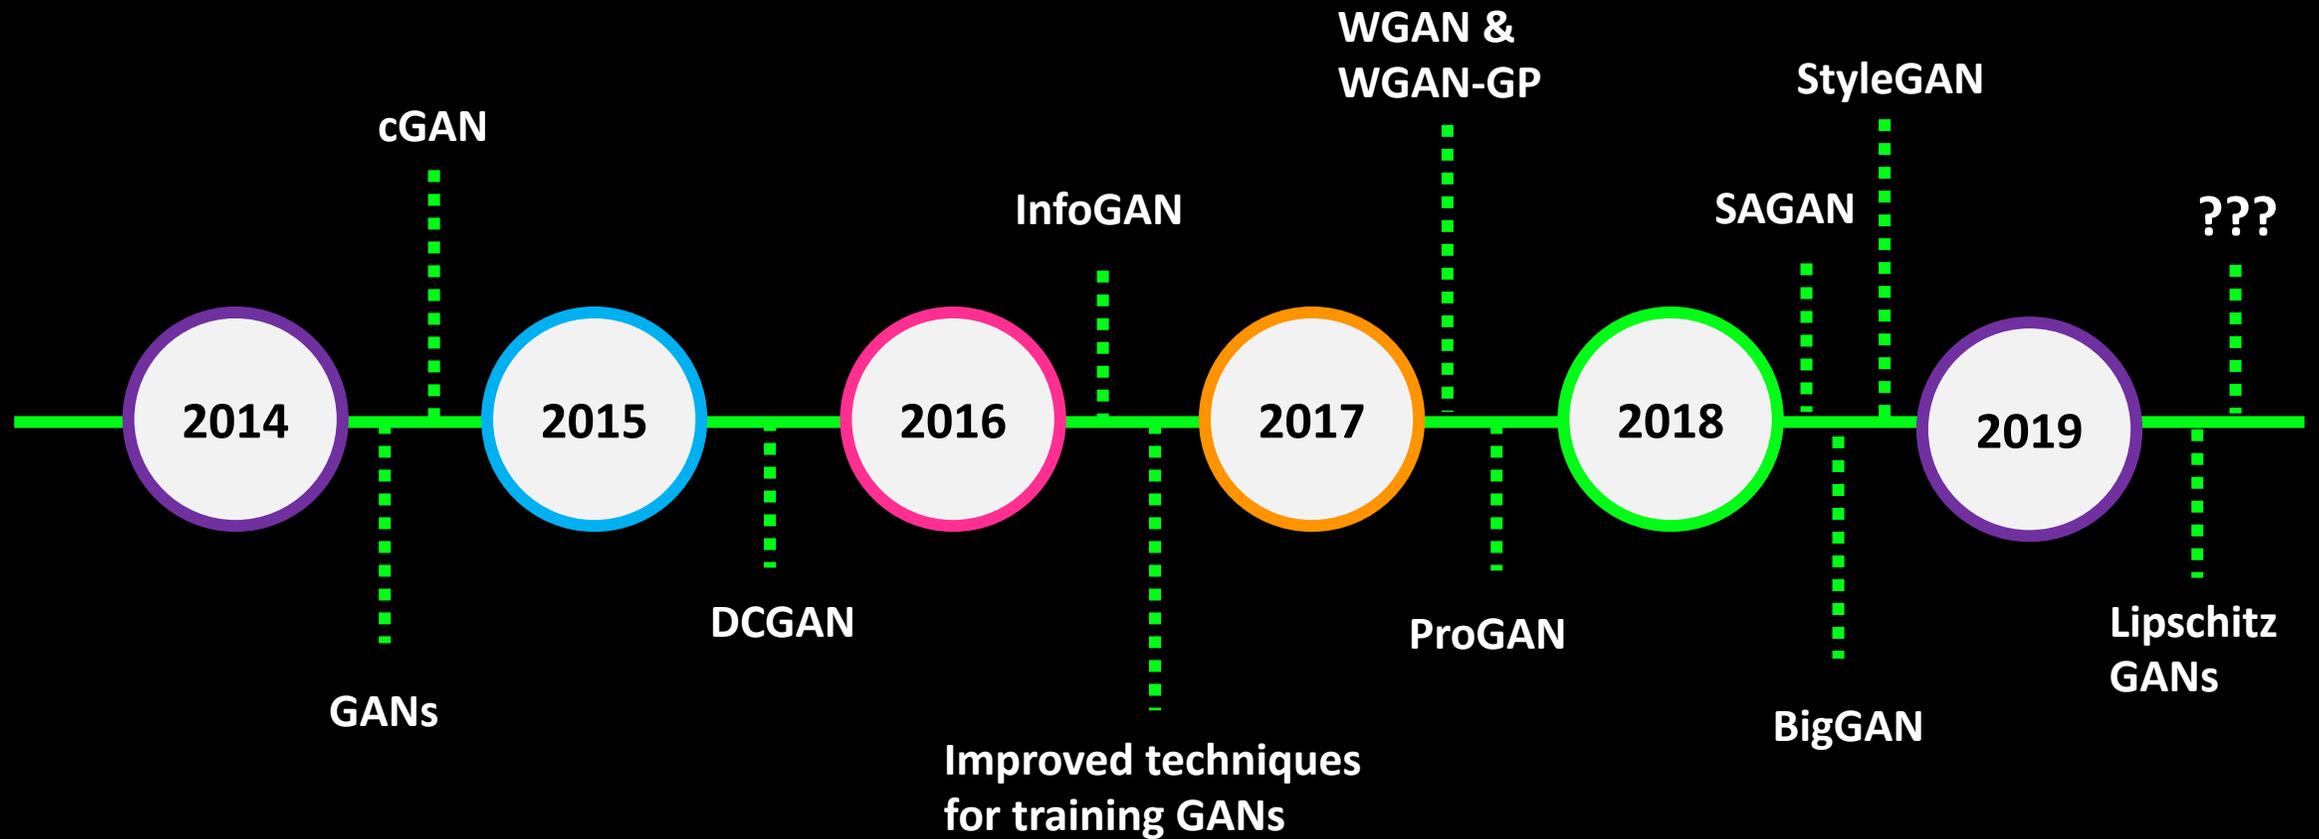


"The GANfather", Ian Goodfellow

"We propose a new framework for estimating generative models via an adversarial process, in which we simultaneously train two models: a generative model G that captures the data distribution, and a discriminative model D that estimates the probability that a sample came from the training data rather than G . The training procedure for G is to maximize the probability of D making a mistake. This framework corresponds to a minimax two-player game."

– [Generative Adversarial Nets, Goodfellow I. et al \(2014\)](#)

Generative Adversarial Nets

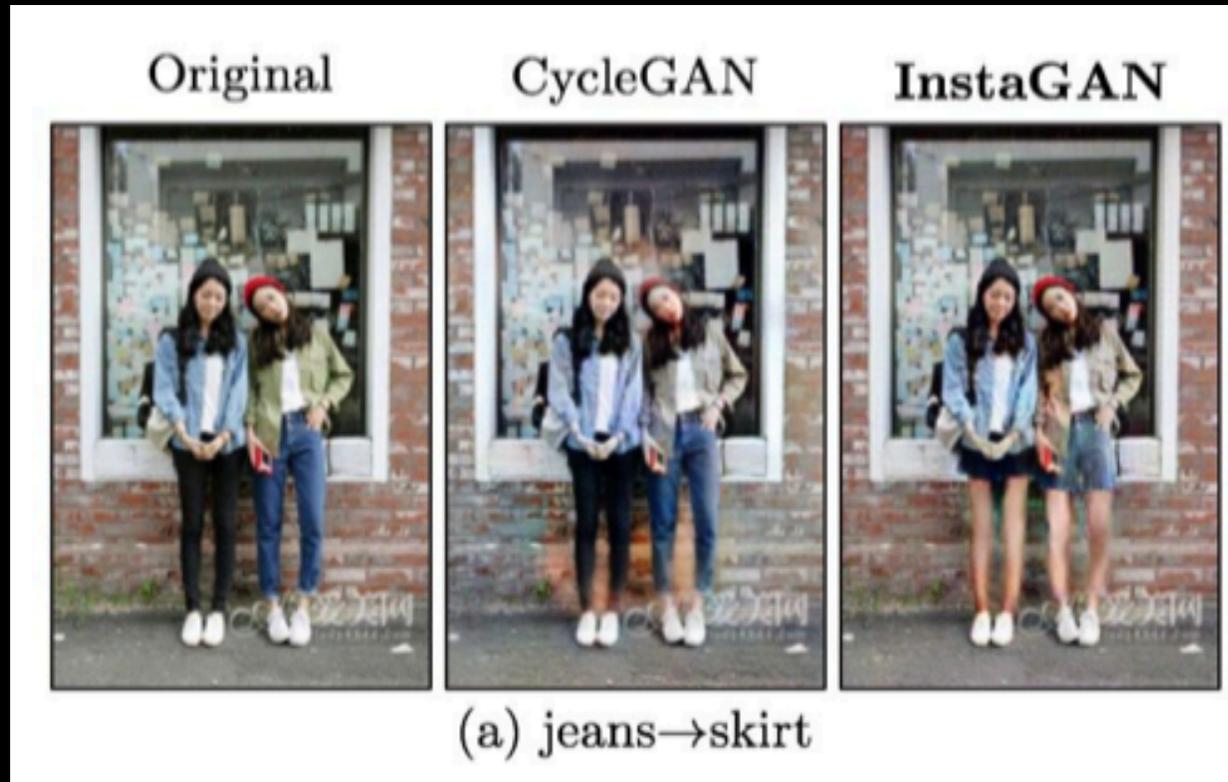


Where are we going?

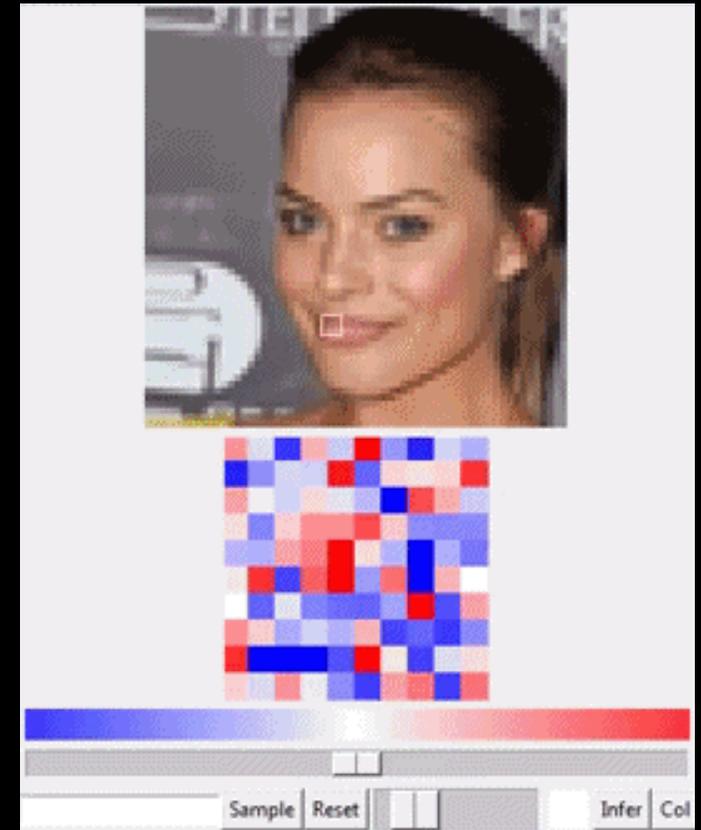


Image credit: goodfellow_ian (2019) "4.5 years of GAN progress on face generation." [Tweet]

Industry Applications



Mo, S., Cho, M., & Shin, J. (2018). InstaGAN: Instance-aware Image-to-Image Translation. *arXiv preprint arXiv:1812.10889*.



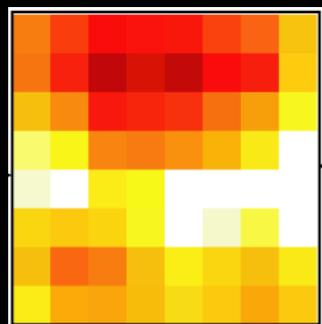
<https://github.com/ajbrock/Neural-Photo-Editor>

Tutorial Outline

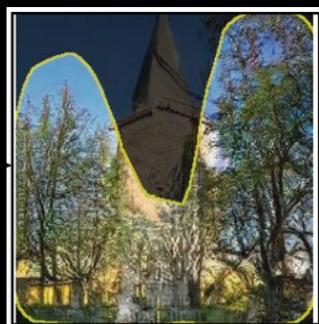
Exploring a generator

Presented by: David Bau

What can GANdissect show us about the internal representations of a generator?



feature map

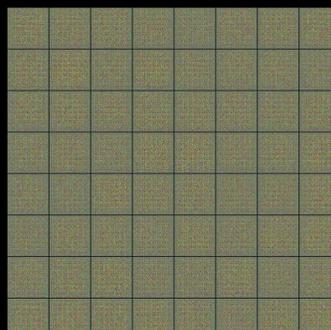


thresholded

Training DCGAN

Presented by: Alex Andonian

A how-to walkthrough on training GANs using Jupyter notebooks



DCGAN training progress

GANtidotes

Presented by: Hendrik Strobelt

Discouraging and mitigating antagonistic use of GANs



which one exists? (answer: neither)