Introduction to GANs

Ian Goodfellow, Staff Research Scientist, Google Brain

CVPR Tutorial on GANs

Salt Lake City, 2018-06-22

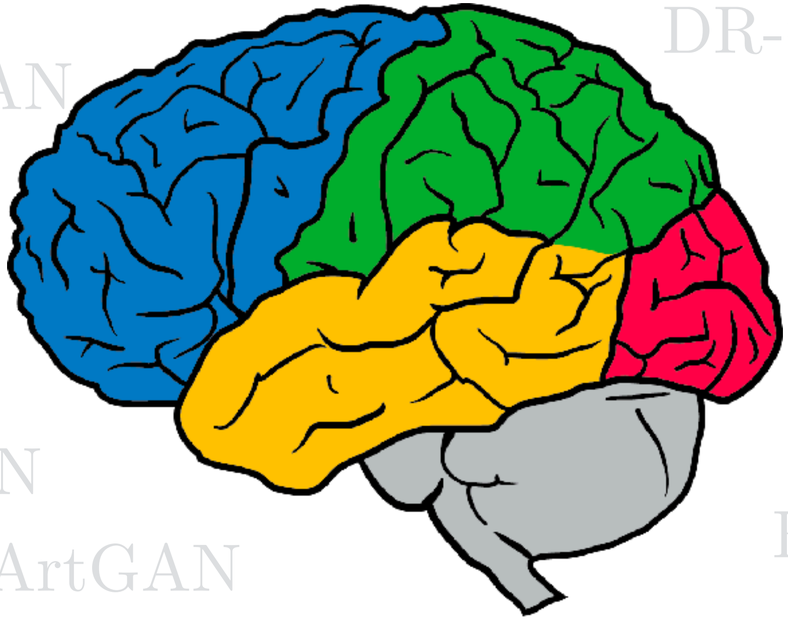

Generative Modeling: Density Estimation

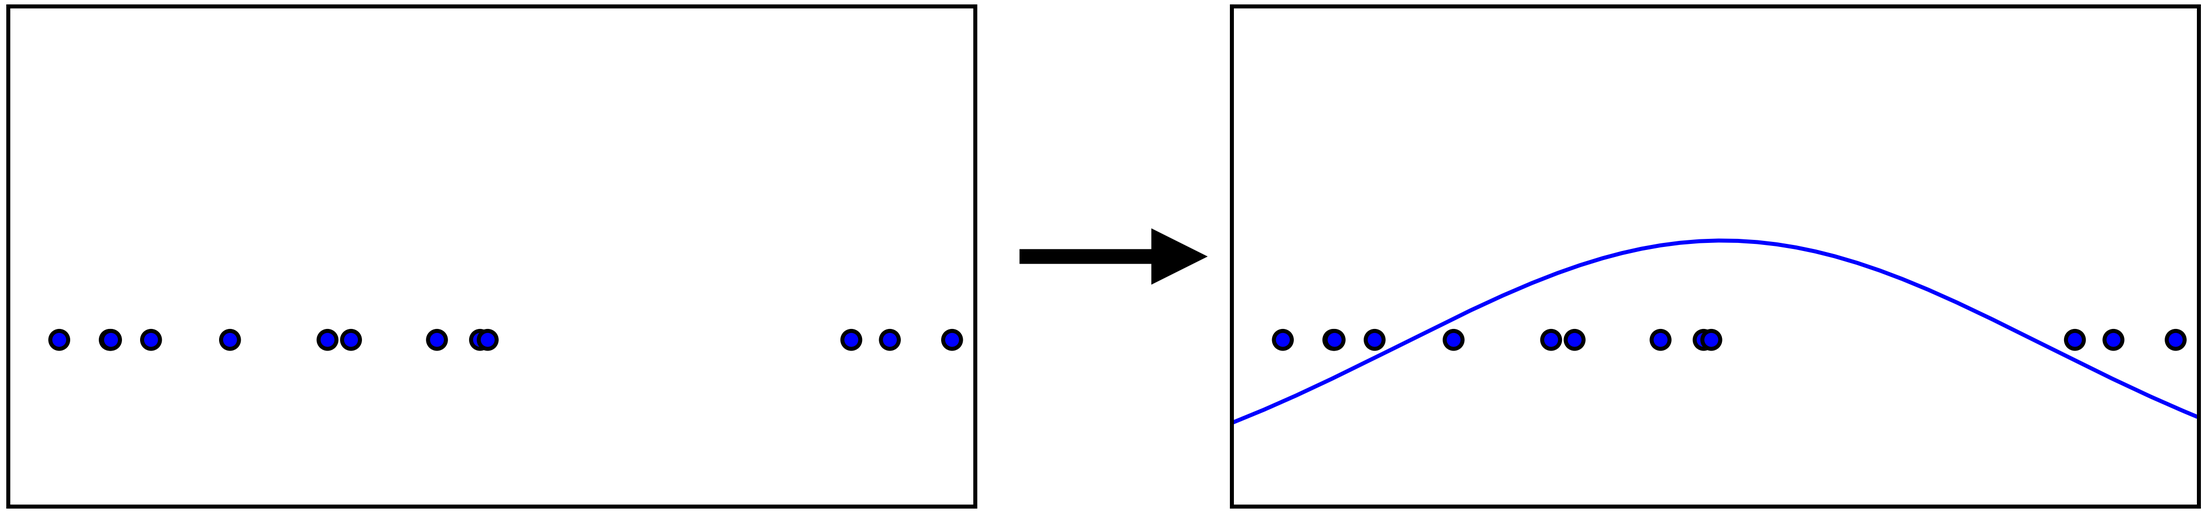

Training Data

Density Function

Generative Modeling: Sample Generation

Training Data
(CelebA)

Sample Generator
(Karras et al, 2017)

Adversarial Nets Framework

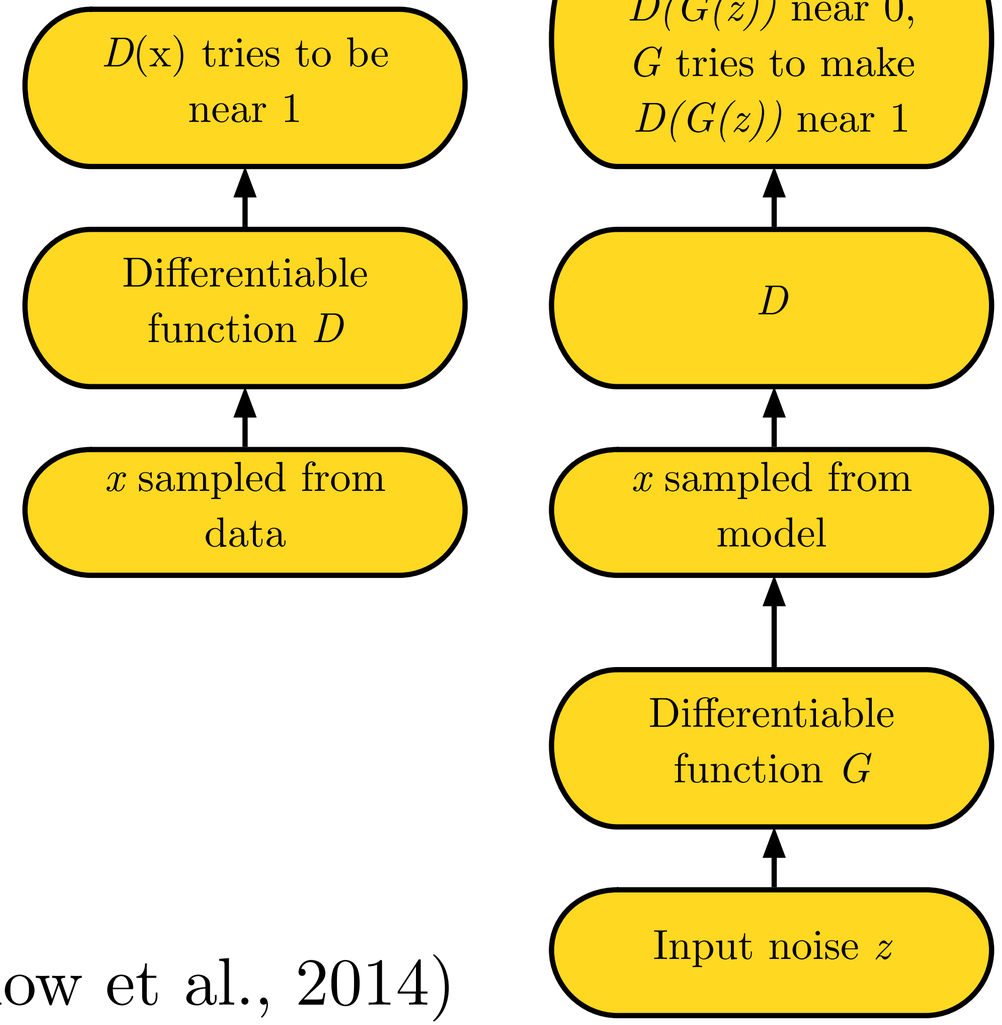

(Goodfellow et al., 2014)

Self-Play

1959: Arthur Samuel's checkers agent

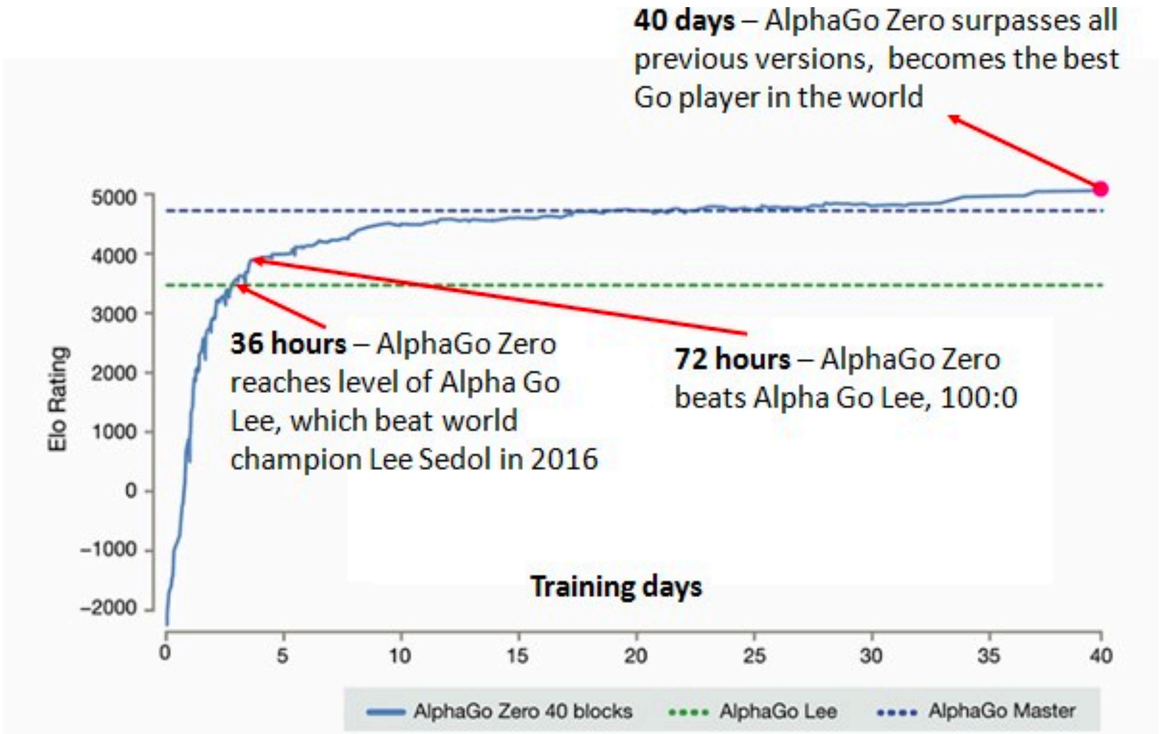

(Silver et al, 2017)

(OpenAI, 2017)

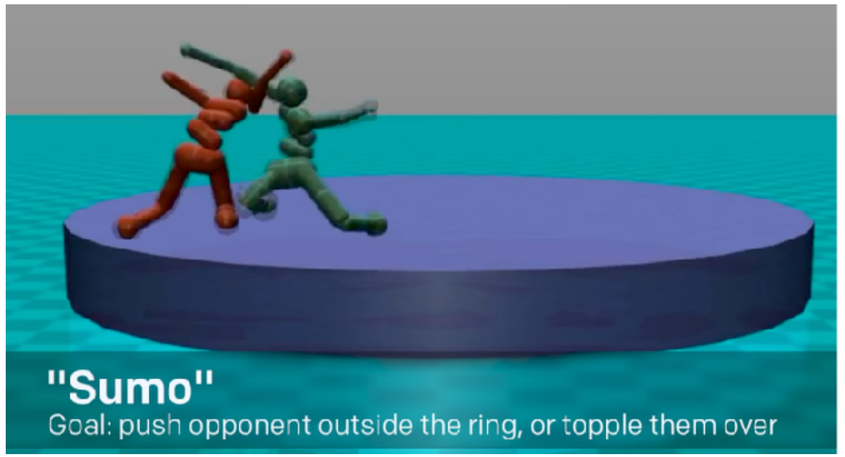

(Bansal et al, 2017)

3.5 Years of Progress on Faces

2014

2015

2016

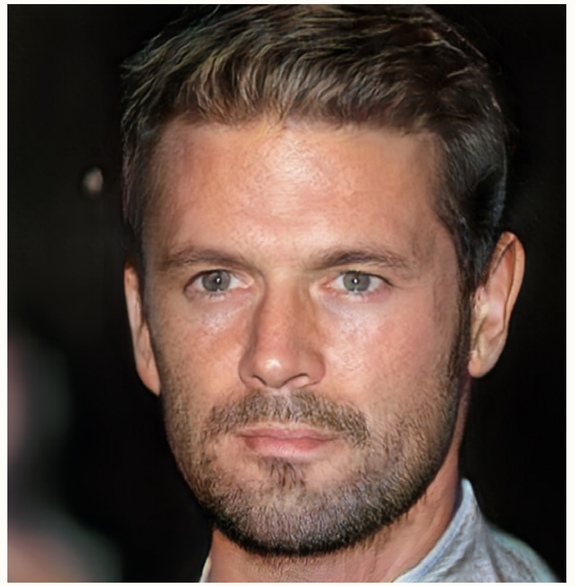

2017

(Brundage et al, 2018)

<2 Years of Progress on ImageNet

Odena et al 2016

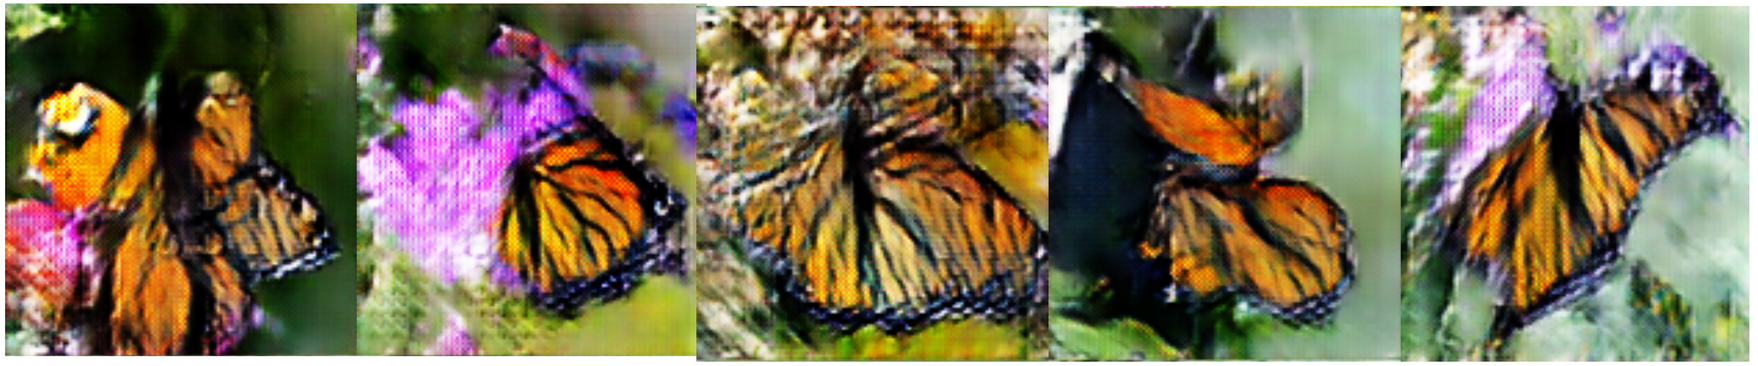

Miyato et al 2017

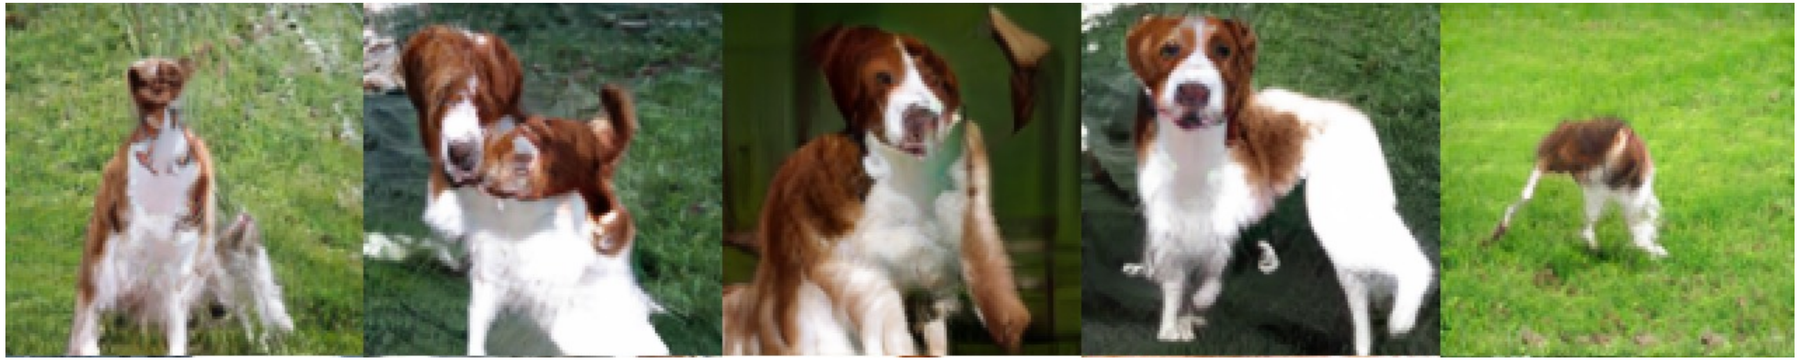

Zhang et al 2018

Self-Attention GAN

State of the art FID on ImageNet: 1000 categories, 128x128 pixels

Goldfish

Redshank

Broccoli

Tiger Cat

Geyser

Indigo Bunting

Stone Wall

(Zhang et al., 2018)

Saint Bernard

From GAN to SAGAN

- Depth and Convolution
- Class-conditional generation
- Spectral Normalization
- Hinge loss
- Two-timescale update rule
- Self-attention

No Convolution Needed to Solve Simple Tasks

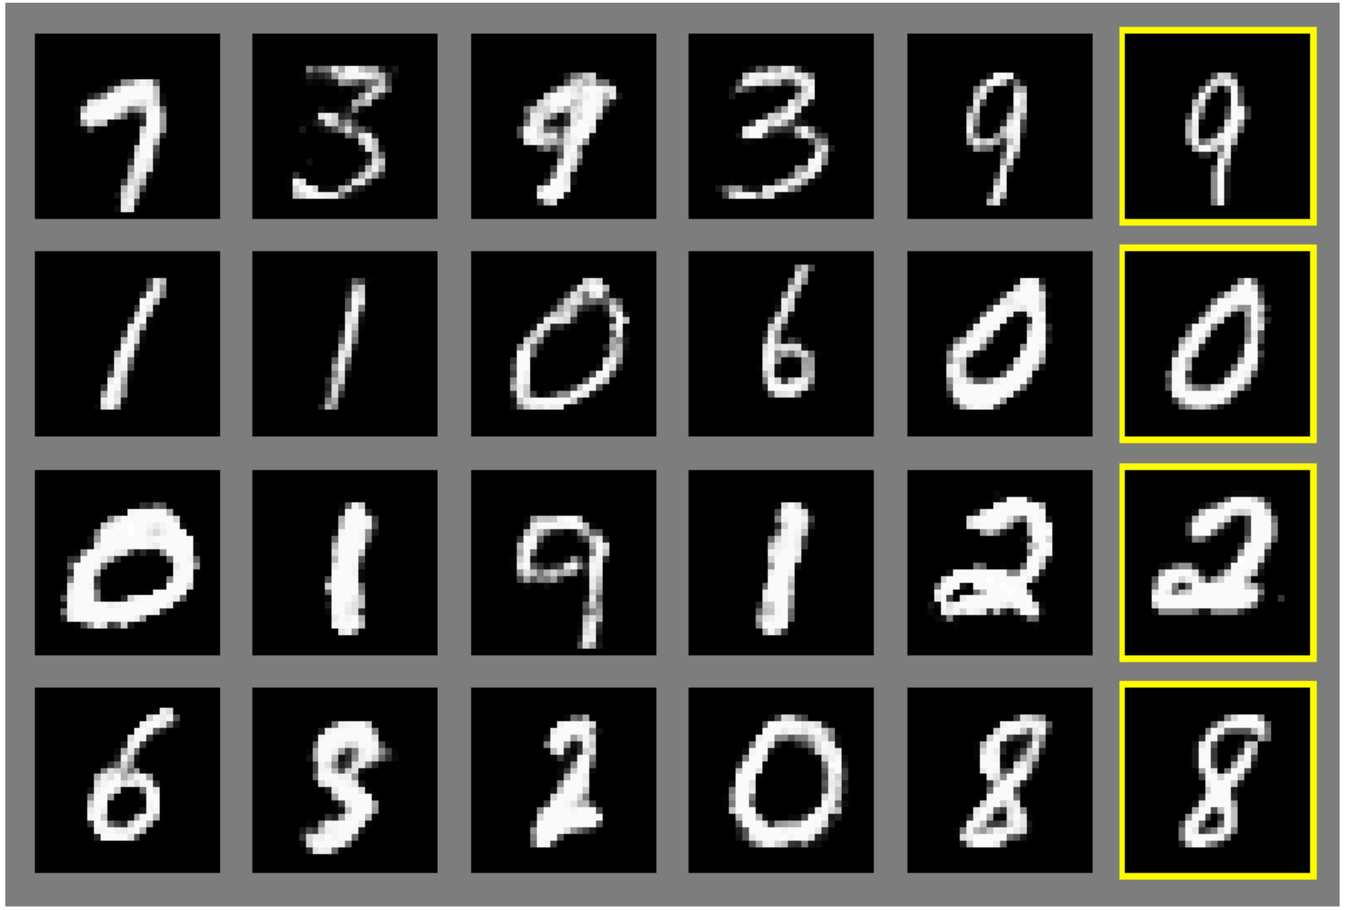

Original GAN, 2014

Depth and Convolution for Harder Tasks

Original GAN (CIFAR-10)

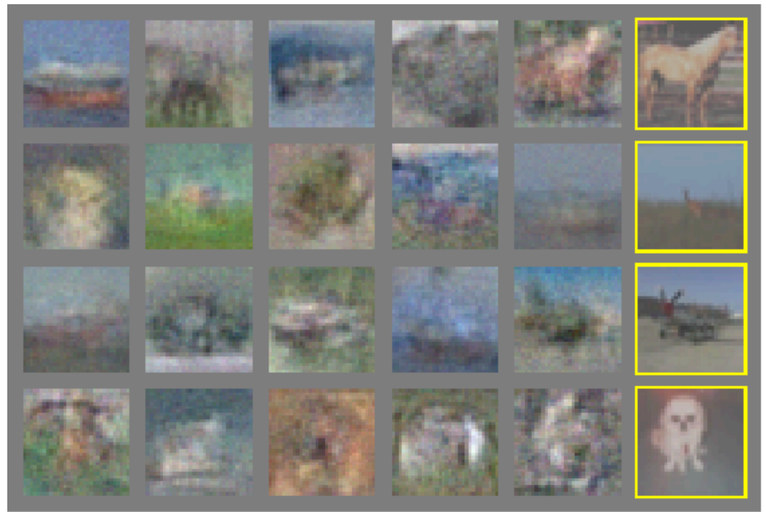

No convolution

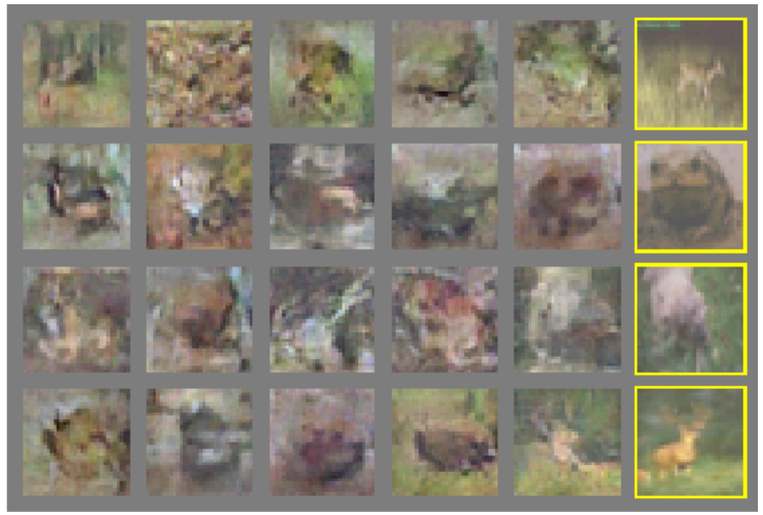

One convolutional layer

DCGAN (ImageNet)

Many convolutional layers
(Radford et al, 2015)

From GAN to SAGAN

- Depth and Convolution
- Class-conditional generation
- Spectral Normalization
- Hinge loss
- Two-timescale update rule
- Self-attention

Class-Conditional GANs

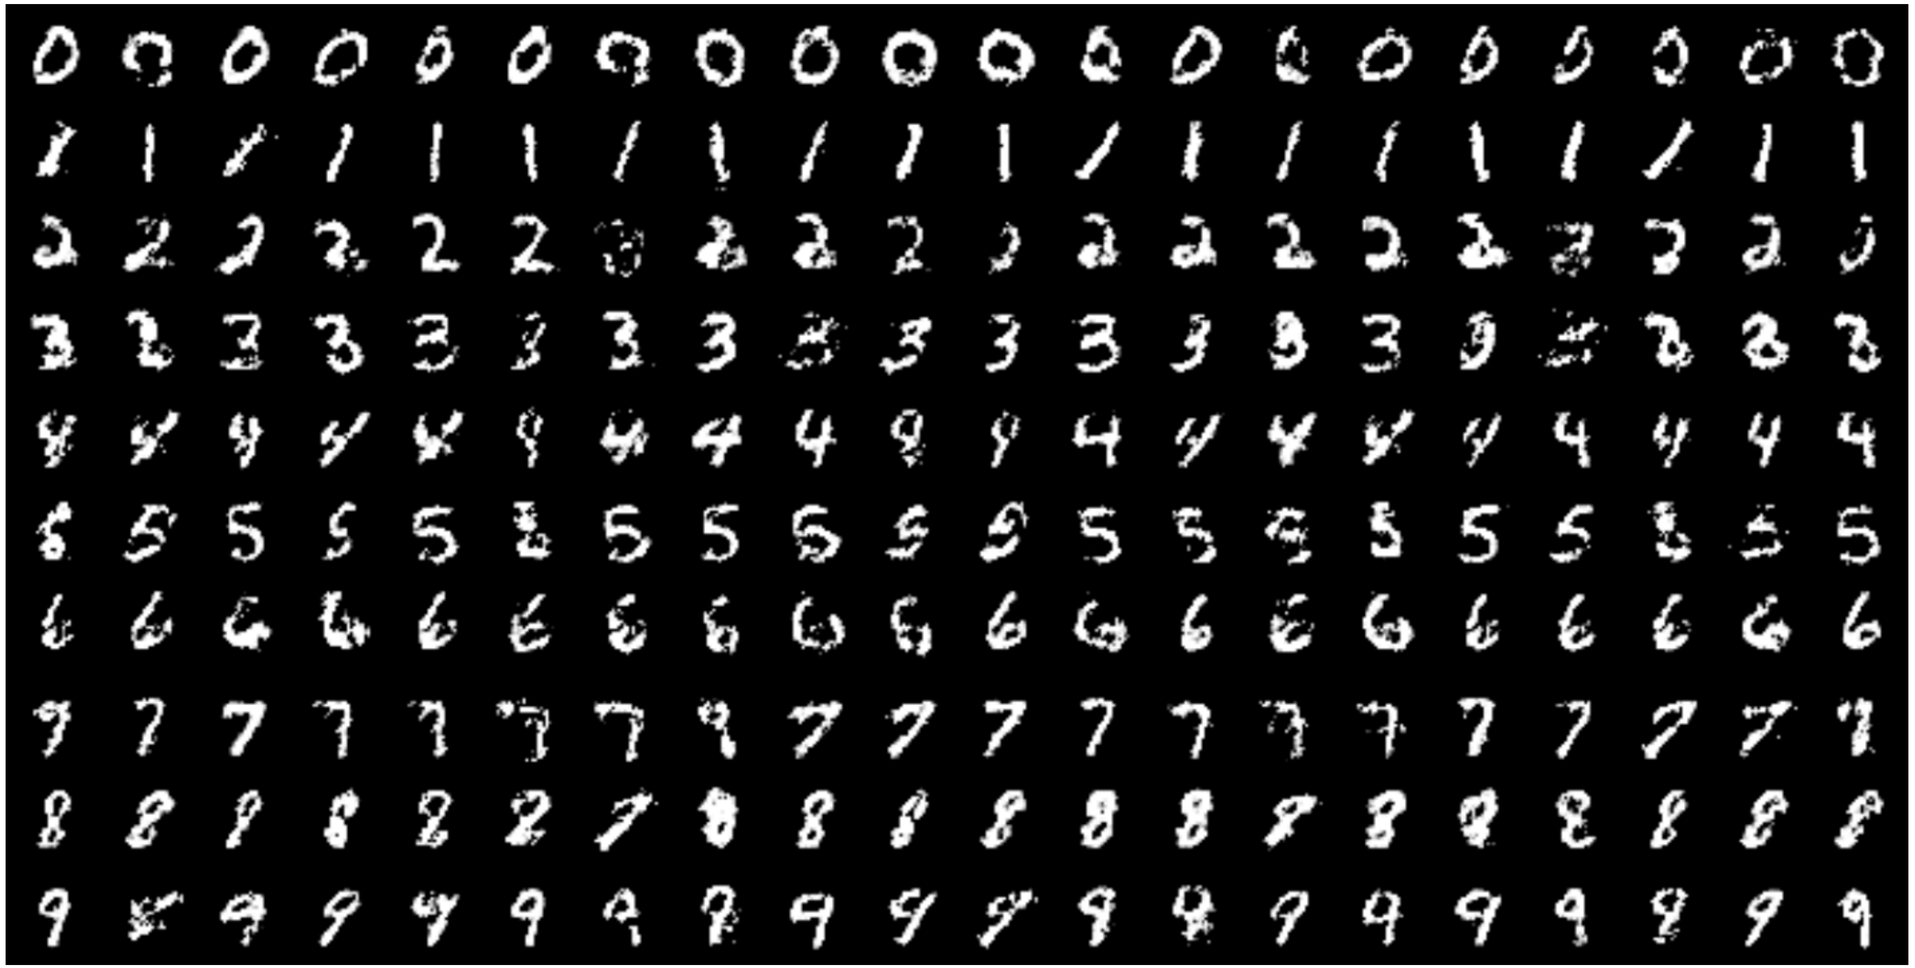

(Mirza and Osindero, 2014)

AC-GAN: Specialist Generators

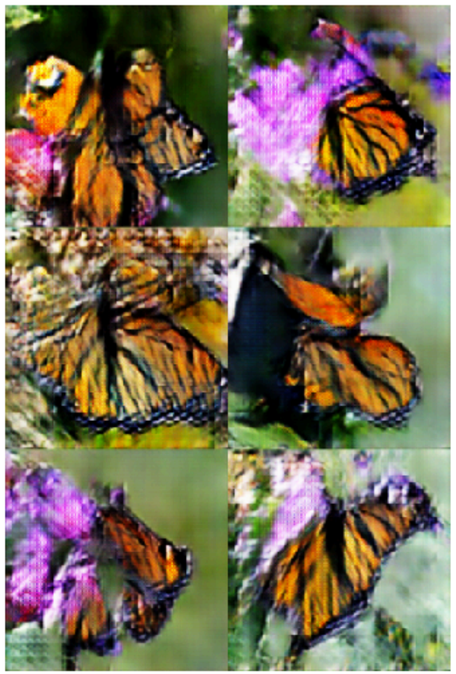

monarch butterfly

goldfinch

daisy

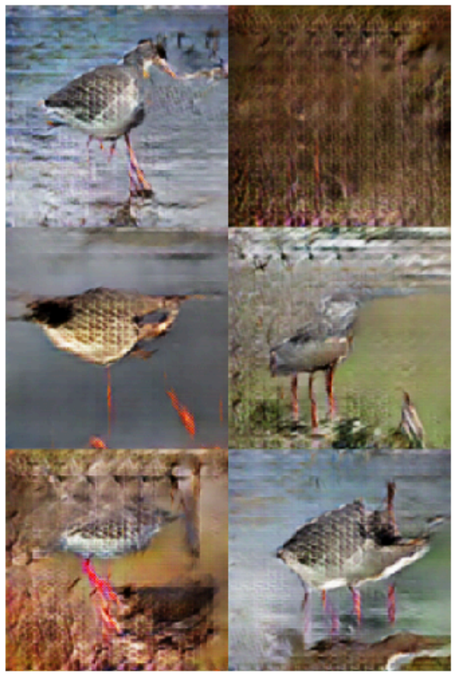

redshank

grey whale

(Odena et al, 2016)

SN-GAN: Shared Generator

Gray whale

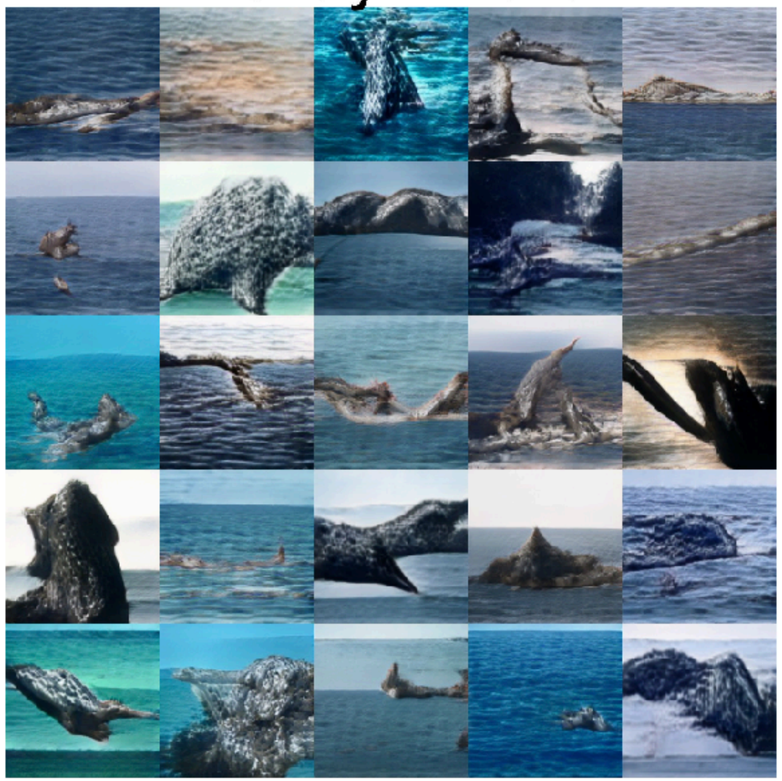

Welsh springer spaniel

Persian cat

(Miyato et al, 2017)

From GAN to SAGAN

- Depth and Convolution
- Class-conditional generation
- Spectral Normalization
- Hinge loss
- Two-timescale update rule
- Self-attention

Spectral Normalization

$$\sigma(A) := \max_{\boldsymbol{h}:\boldsymbol{h}\neq\mathbf{0}} \frac{\|A\boldsymbol{h}\|_2}{\|\boldsymbol{h}\|_2} = \max_{\|\boldsymbol{h}\|_2\leq 1} \|A\boldsymbol{h}\|_2$$

$$\|f\|_{\text{Lip}} \leq \prod_{l=1}^{L+1} \sigma(W^l)$$

$$\bar{W}_{\text{SN}}(W) := W/\sigma(W)$$

(Miyato et al, 2017)

From GAN to SAGAN

- Depth and Convolution
- Class-conditional generation
- Spectral Normalization
- Hinge loss
- Two-timescale update rule
- Self-attention

Hinge Loss

$$V_D(\hat{G}, D) = \mathop{\mathrm{E}}_{\boldsymbol{x} \sim q_{\mathrm{data}}(\boldsymbol{x})} \left[\min\left(0, -1 + D(\boldsymbol{x})\right)\right] + \mathop{\mathrm{E}}_{\boldsymbol{z} \sim p(\boldsymbol{z})} \left[\min\left(0, -1 - D\left(\hat{G}(\boldsymbol{z})\right)\right)\right] \tag{16}$$

$$V_G(G, \hat{D}) = -\mathop{\mathrm{E}}_{\boldsymbol{z} \sim p(\boldsymbol{z})} \left[\hat{D}\left(G(\boldsymbol{z})\right)\right], \tag{17}$$

(Miyato et al 2017, Lim and Ye 2017, Tran et al 2017)

From GAN to SAGAN

- Depth and Convolution
- Class-conditional generation
- Spectral Normalization
- Hinge loss
- Two-timescale update rule
- Self-attention

Two-Timescale Update Rule

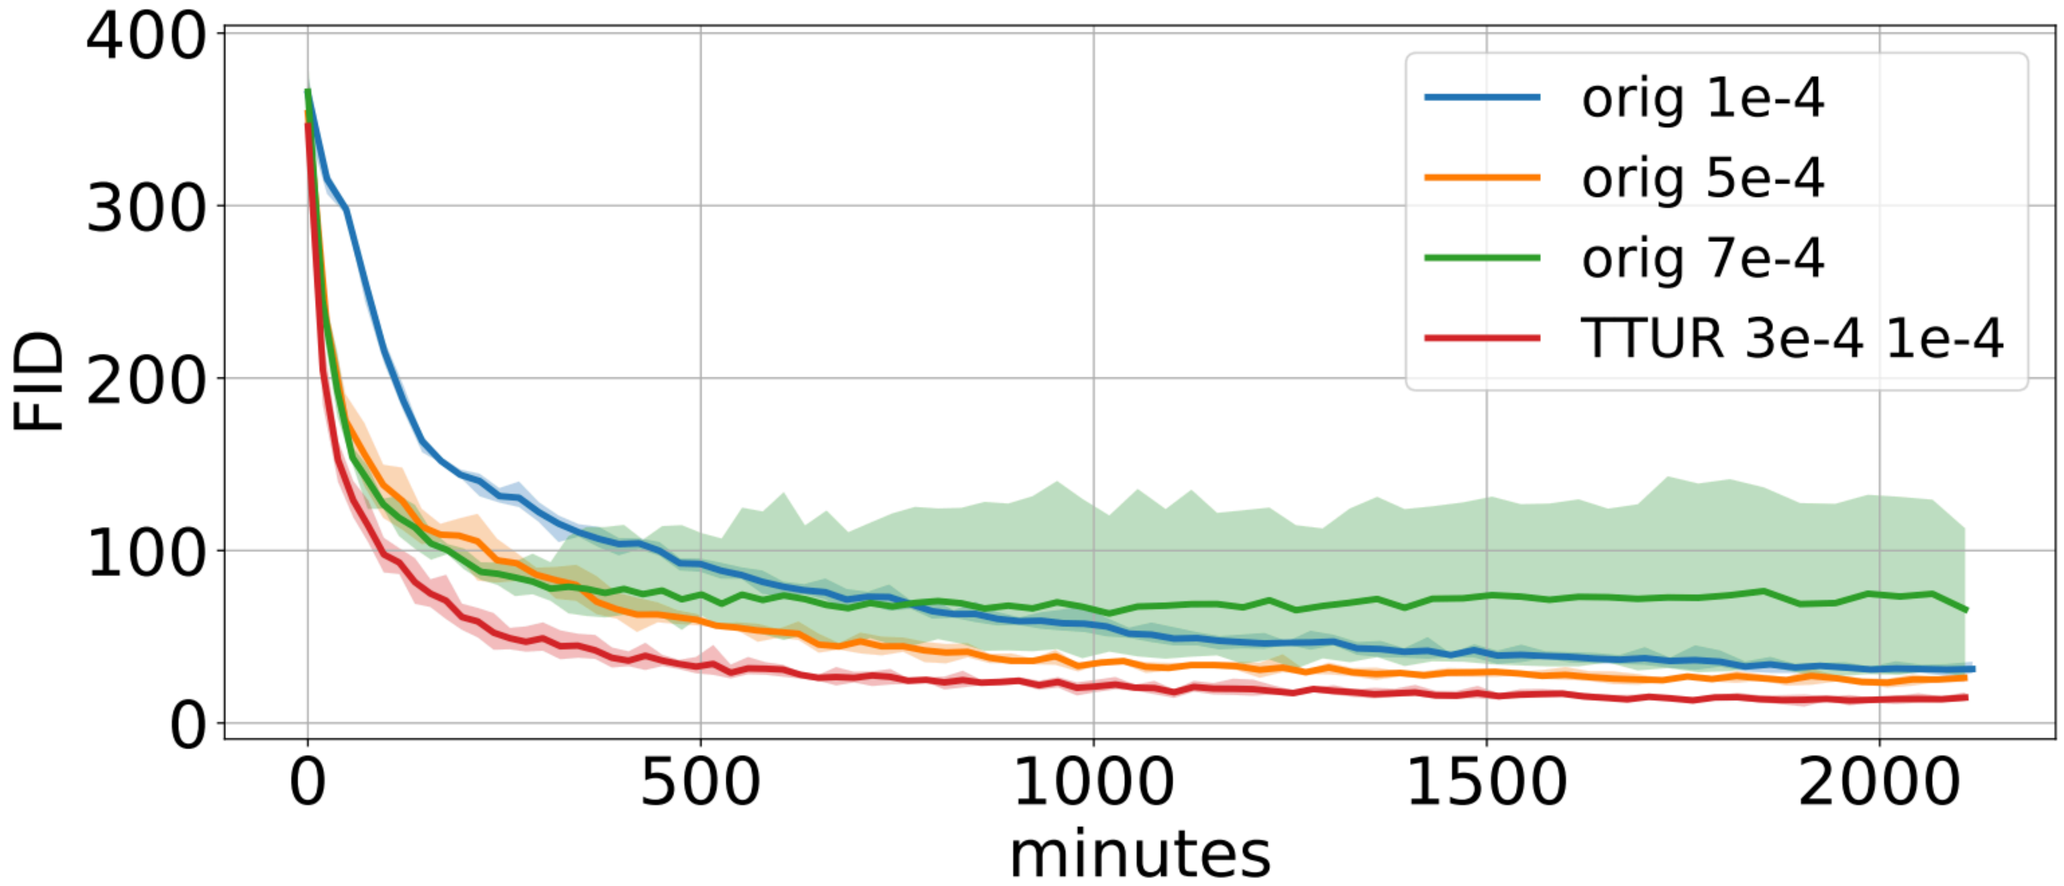

From GAN to SAGAN

- Depth and Convolution
- Class-conditional generation
- Spectral Normalization
- Hinge loss
- Two-timescale update rule
- Self-attention

Self-Attention

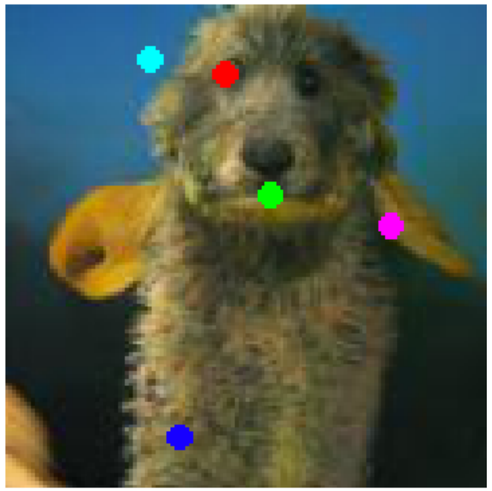


Use layers from
Wang et al 2018


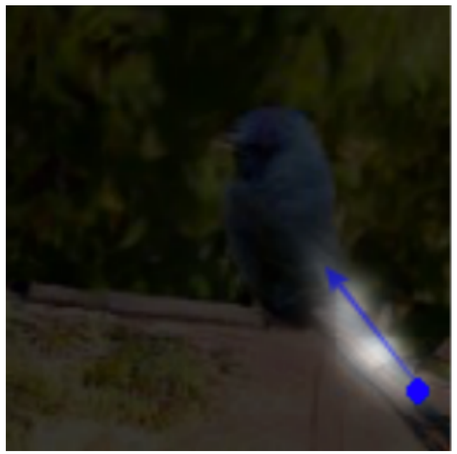

Applying GANs

- Semi-supervised Learning
- Model-based optimization
- Extreme personalization
- Program synthesis

Supervised Discriminator for Semi-Supervised Learning

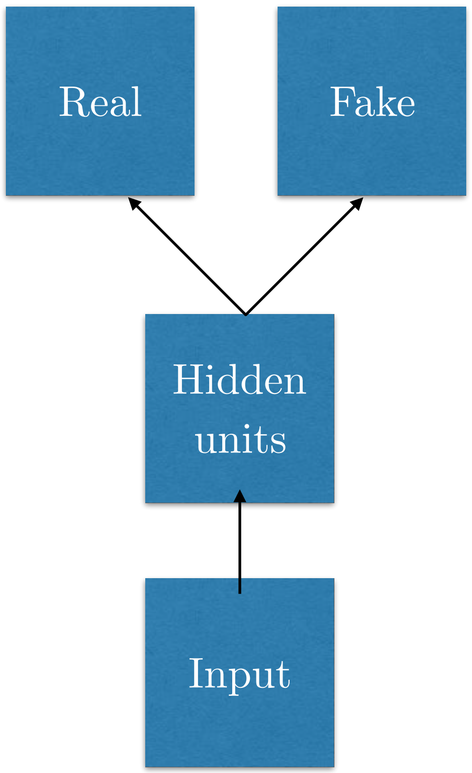

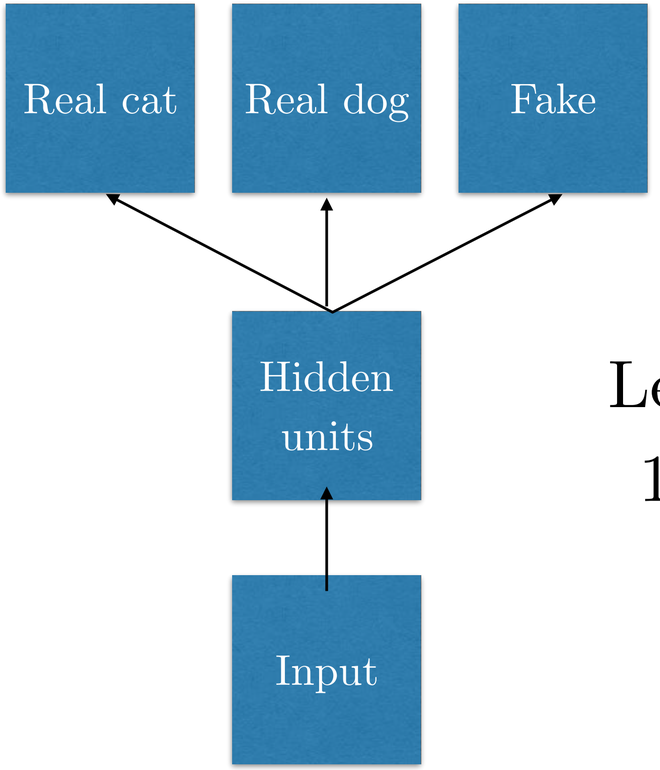

Learn to read with 100 labels rather than 60,000

(Odena 2016, Salimans et al 2016)

Semi-Supervised Classification

MNIST: 100 training labels -> 80 test mistakes

SVHN: 1,000 training labels -> 4.3% test error

CIFAR-10: 4,000 labels -> 14.4% test error

(Dai et al 2017)

Designing DNA to optimize protein binding

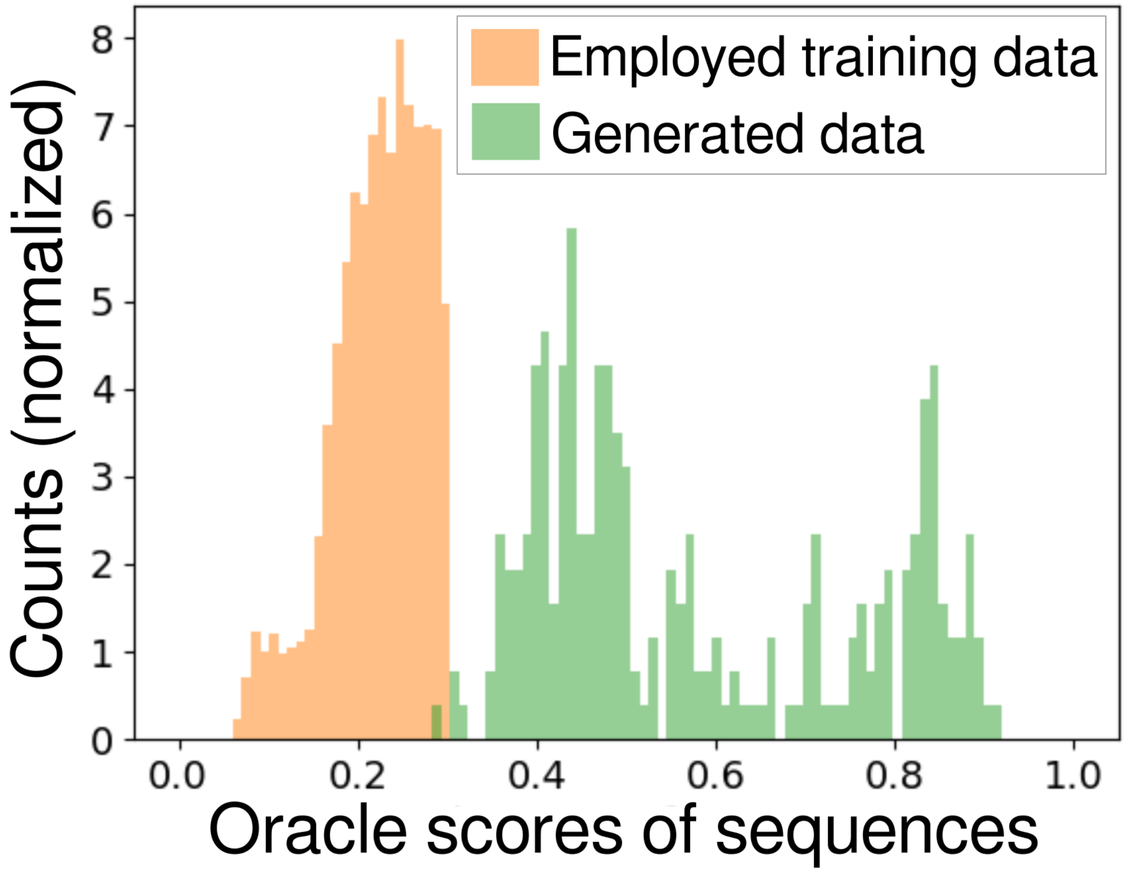

(Killoran et al, 2017)

Personalized GANufacturing

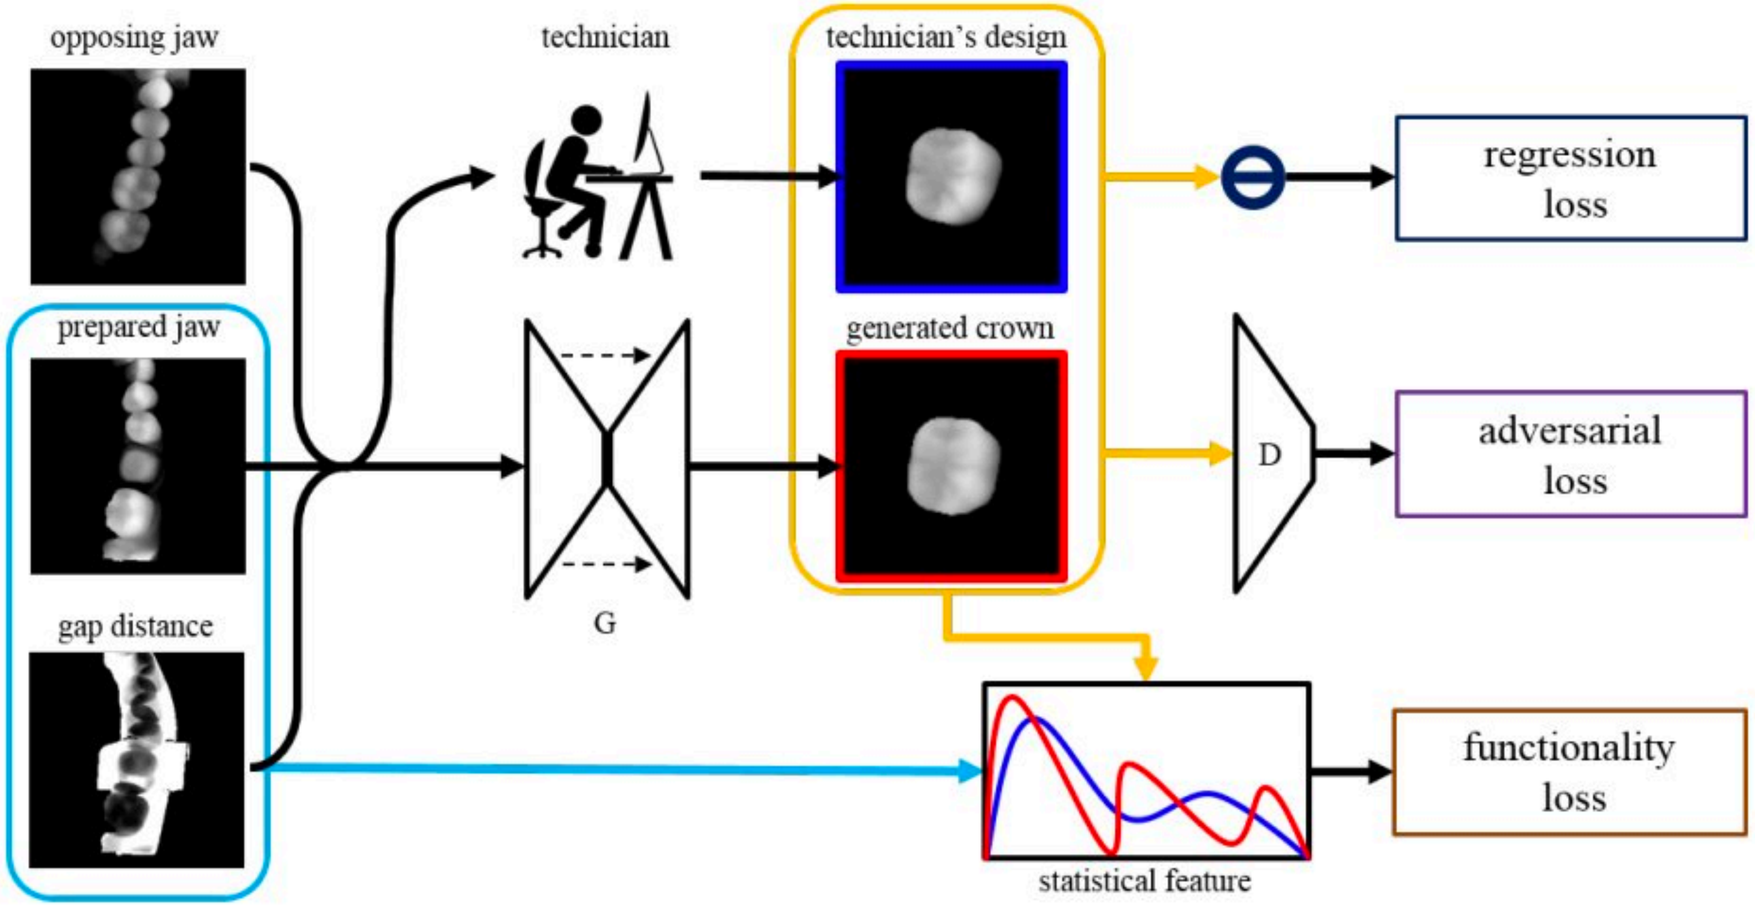

(Hwang et al 2018)

SPIRAL

Synthesizing Programs for Images Using Reinforced Adversarial Learning

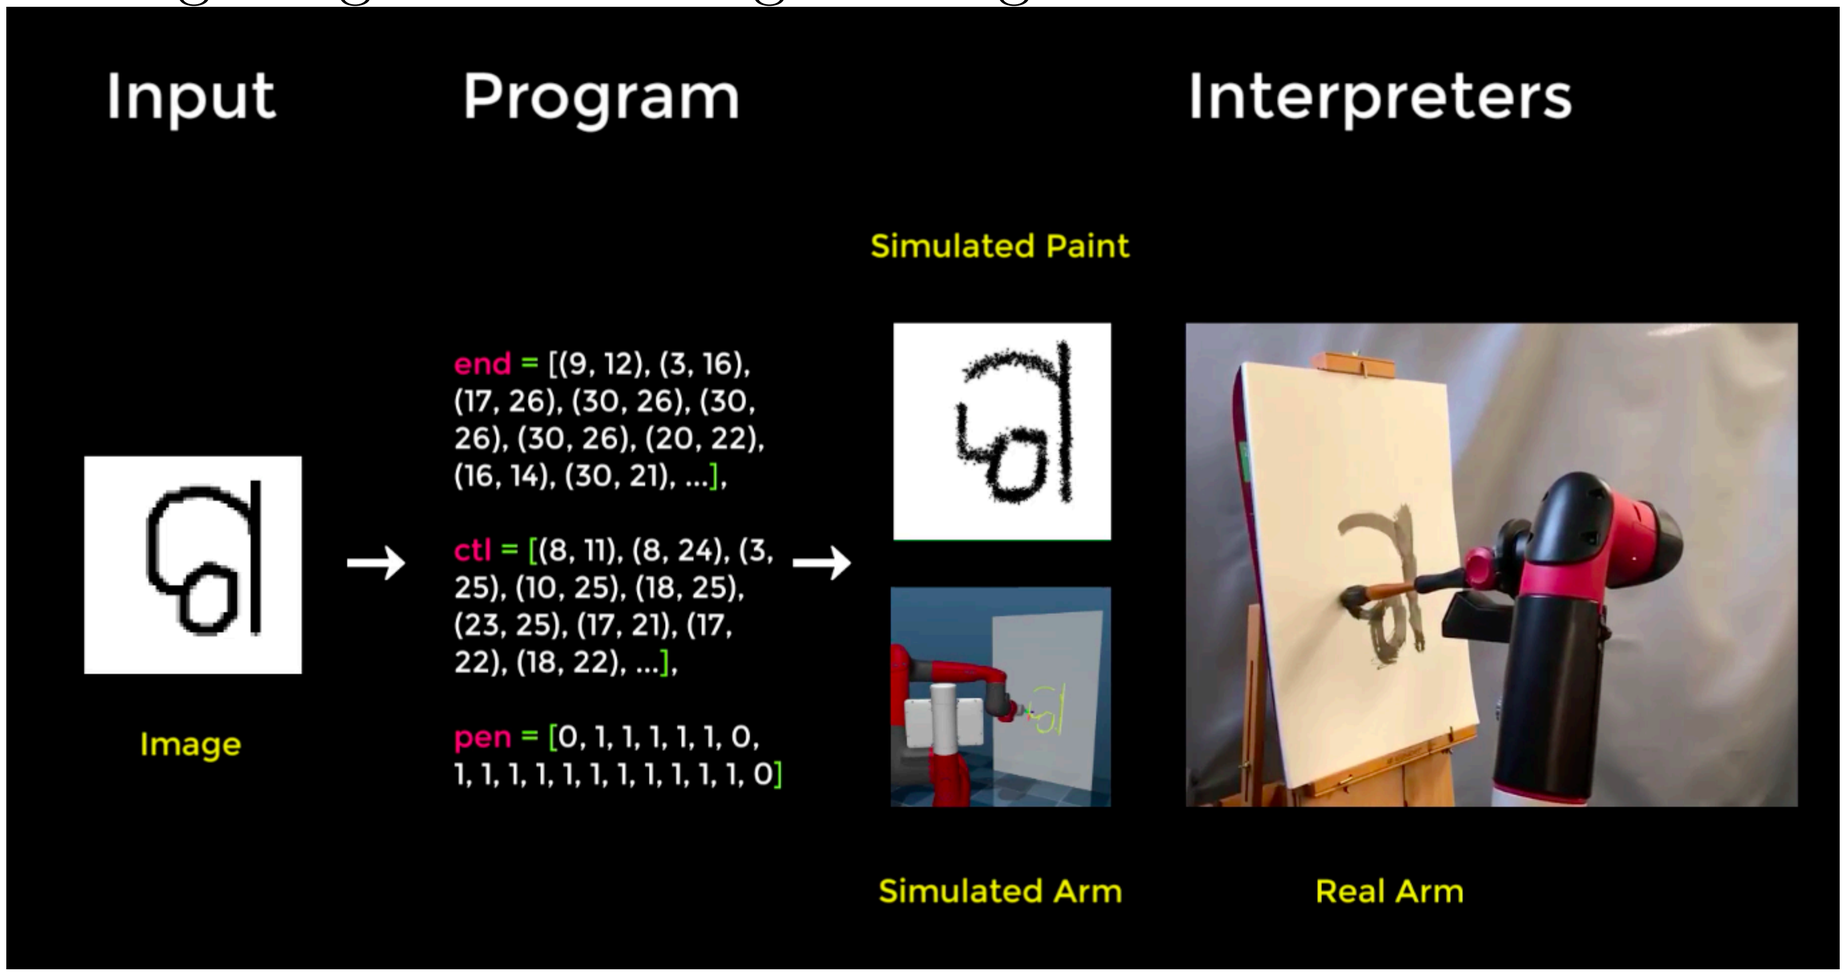

(Ganin et al, 2018)

Other applications

- Planning
- World Models for RL agents
- Fairness and Privacy
- Missing data
- Topics covered at workshop:
  - Training data for other agents (Philip Isola, Taesung Park, Jun-Yan Zhu)
  - Inference in other probabilistic models (Mihaela Rosca)
  - Domain adaptation (Judy Hoffman)
  - Imitation Learning (Stefano Ermon)

Track updates at the GAN Zoo

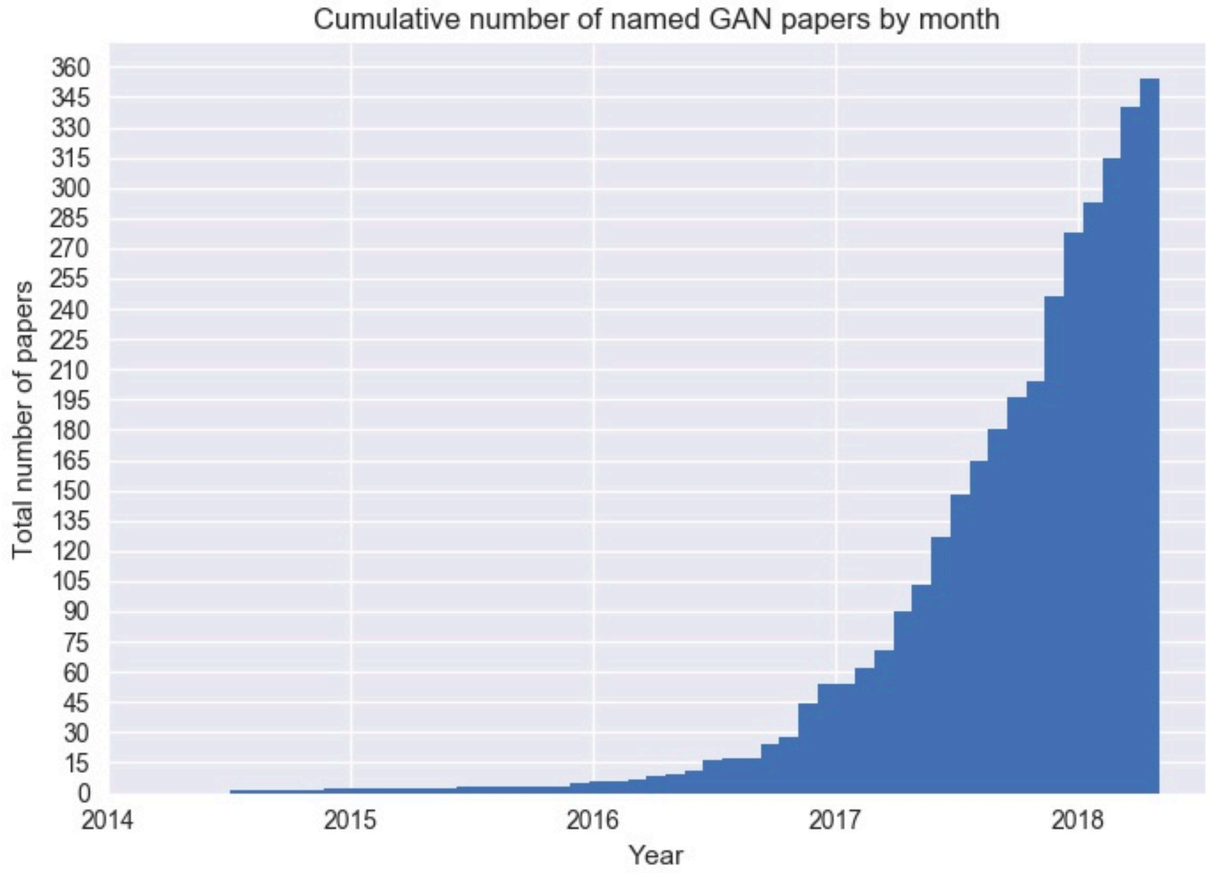

https://github.com/hindupuravinash/the-gan-zoo

Questions