Introduction to GANs

Ian Goodfellow, Staff Research Scientist, Google Brain

IEEE Workshop on Perception Beyond the Visible Spectrum

Salt Lake City, 2018-06-18

MedGAN ID-CGAN CoGAN LR-GAN CGAN IcGAN b-GAN LS-GAN LAPGAN DiscoGAN MPM-GAN AdaGAN AMGAN iGAN InfoGAN CatGAN IAN LSGAN SAGAN McGAN MIX+GAN MGAN FF-GAN BS-GAN GoGAN C-VAE-GAN C-RNN-GAN DR-GAN DCGAN CCGAN AC-GAN MAGAN 3D-GAN BiGAN GAWWN DualGAN CycleGAN GP-GAN Bayesian GAN AnoGAN EBGAN DTN MAD-GAN Context-RNN-GAN ALI BEGAN AL-CGAN MARTA-GAN f-GAN ArtGAN MalGAN

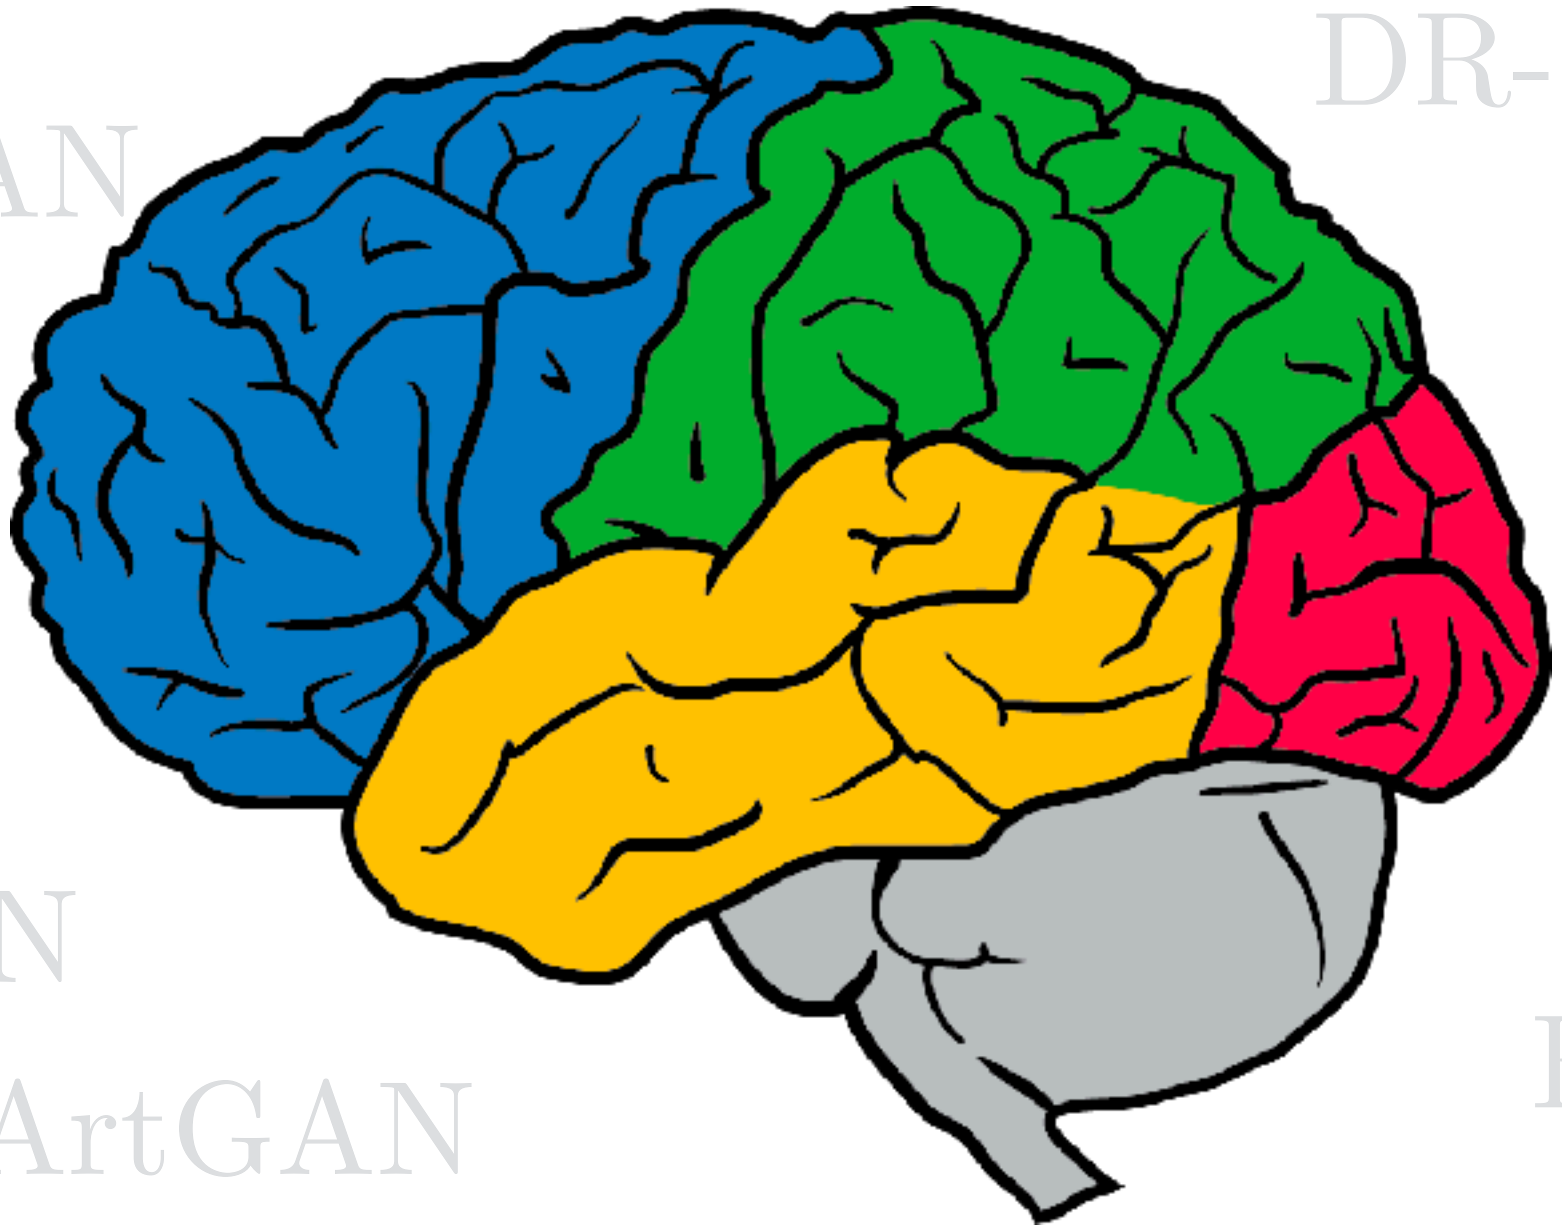

Generative Modeling: Density Estimation

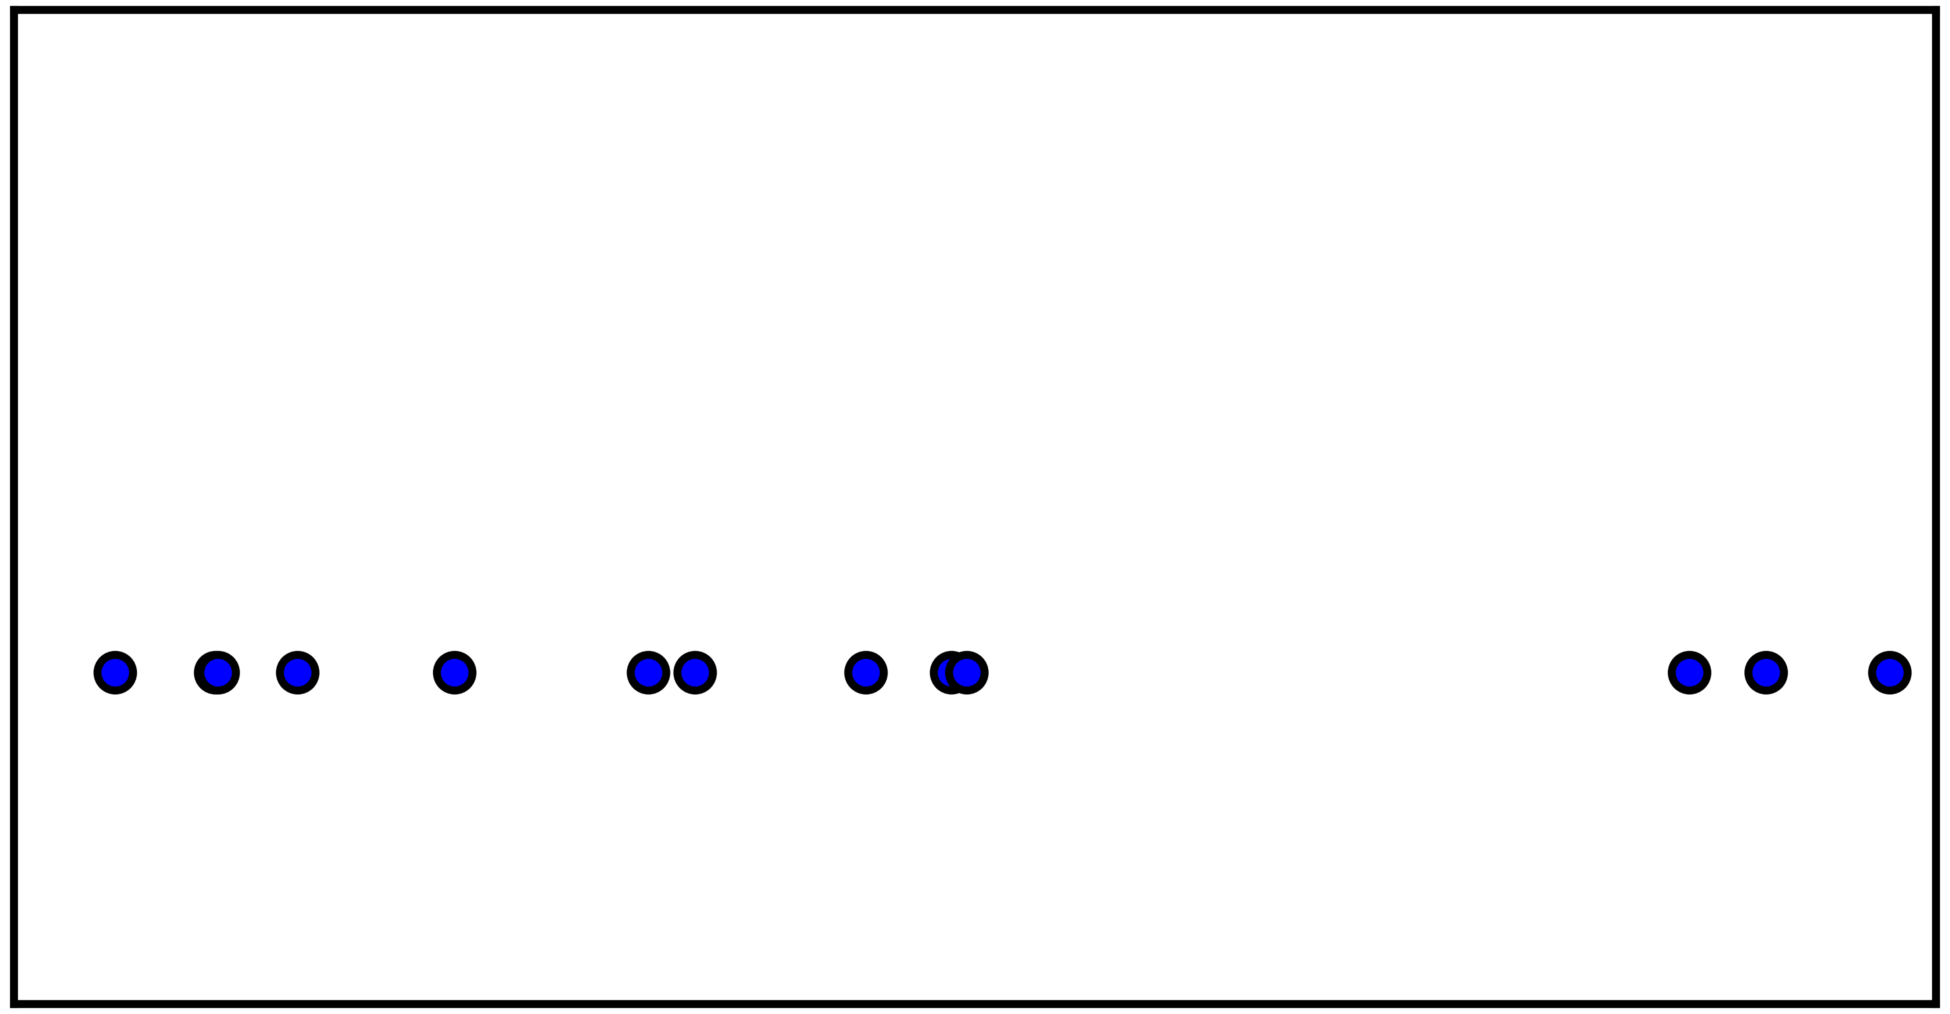

Training Data

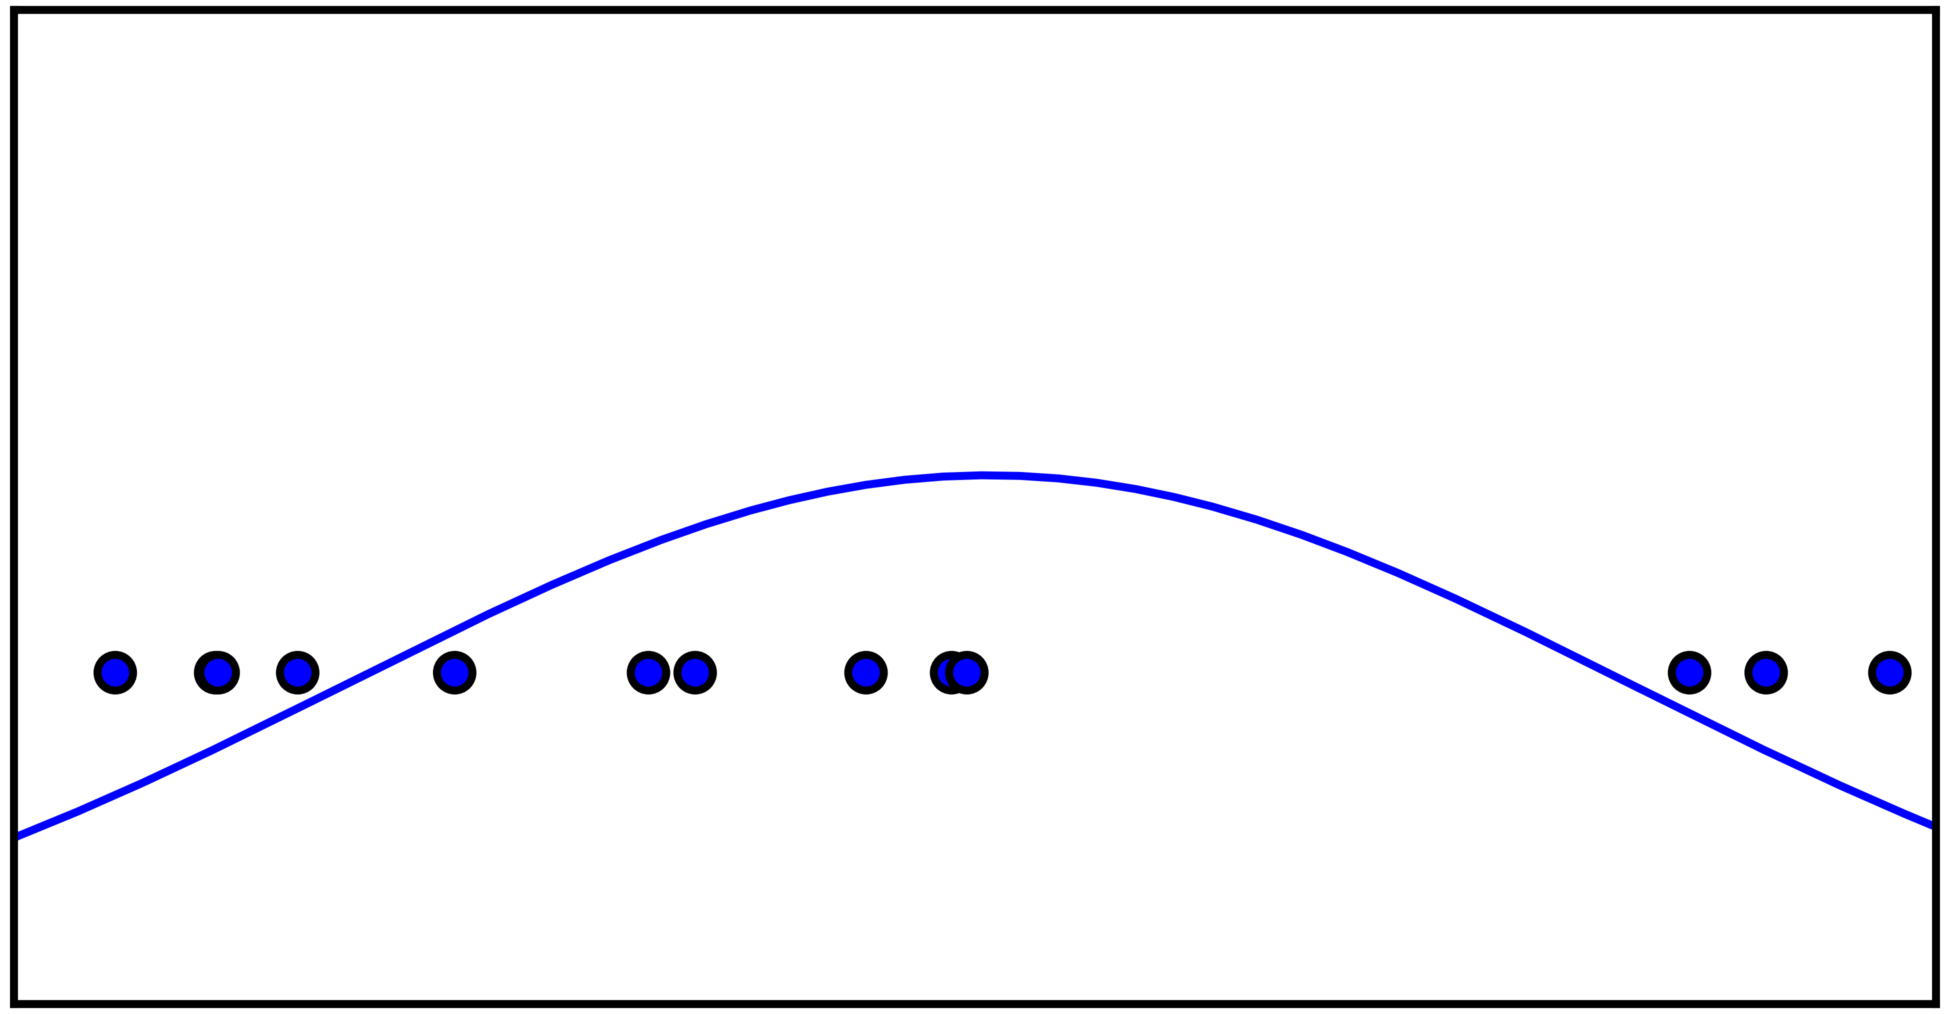

Density Function

(Goodfellow 2018)

Generative Modeling: Sample Generation

Training Data
(CelebA)

Sample Generator
(Karras et al, 2017)

Adversarial Nets Framework

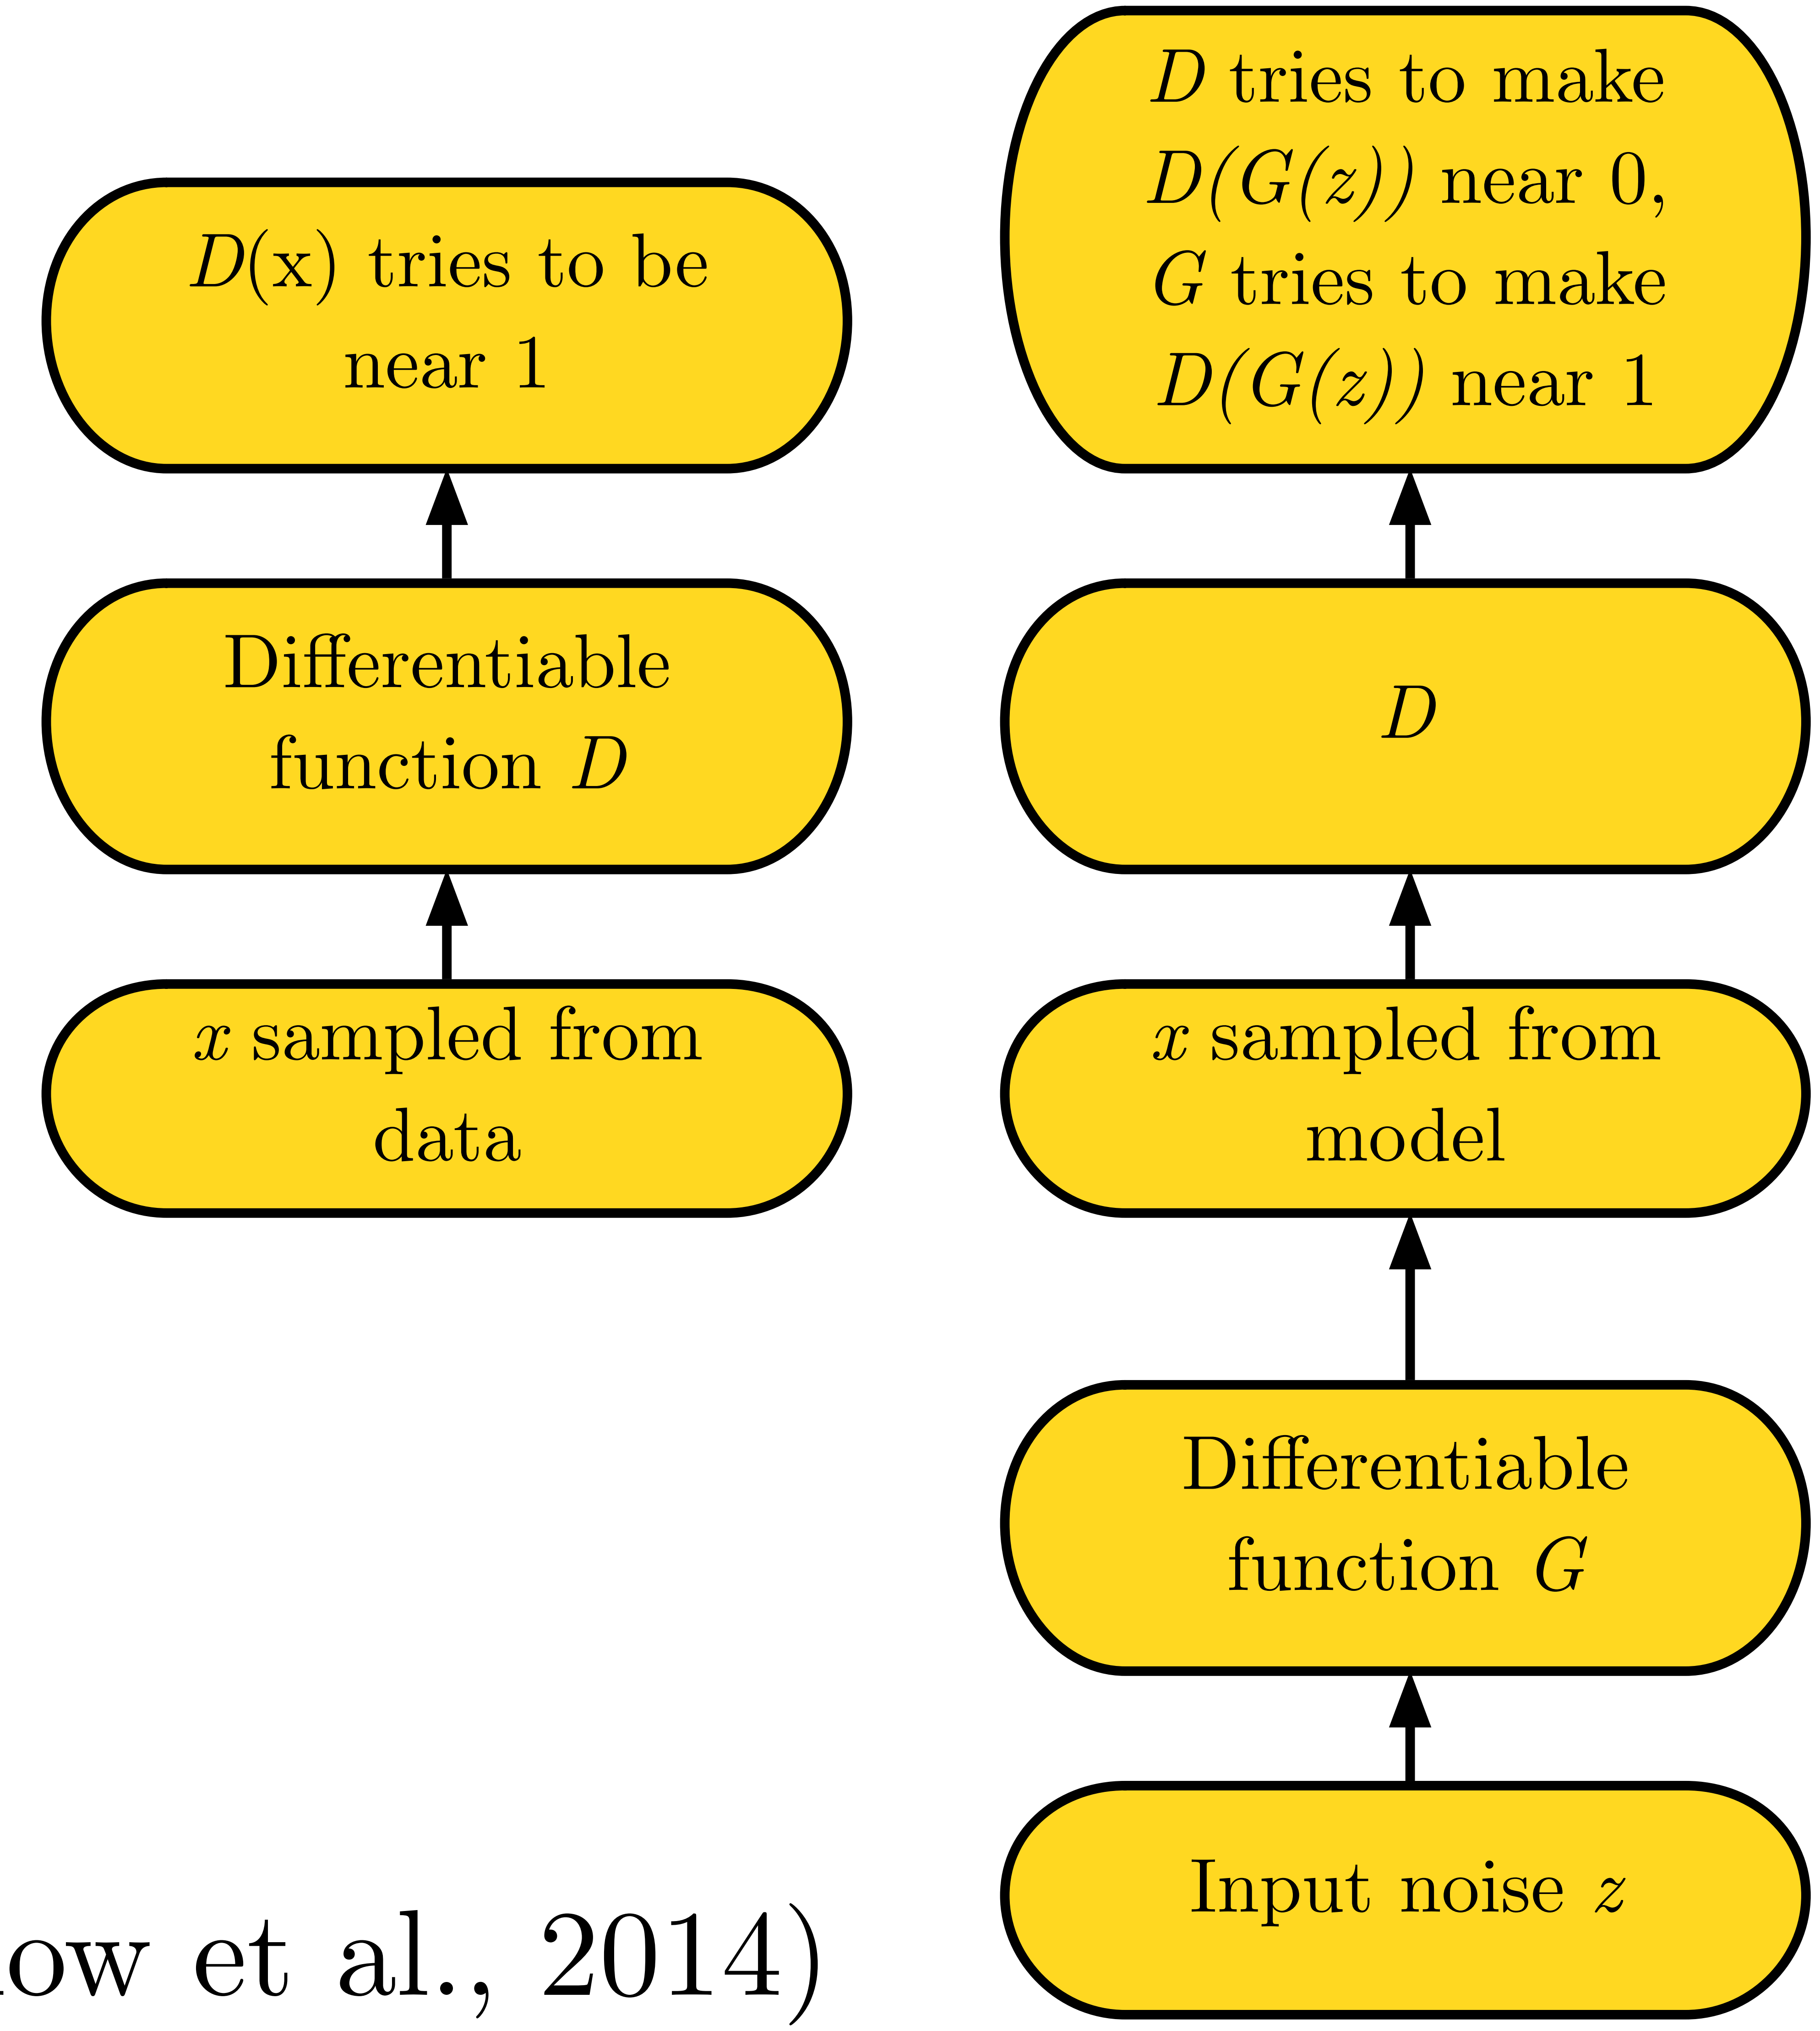

(Goodfellow et al., 2014)

Self-Attention GAN

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

Goldfish

Redshank

Broccoli

Tiger Cat

Geyser

Indigo Bunting

Stone Wall

Saint Bernard

(Zhang et al., 2018)

Self-Play

1959: Arthur Samuel's checkers agent

(Silver et al, 2017)

(OpenAI, 2017)

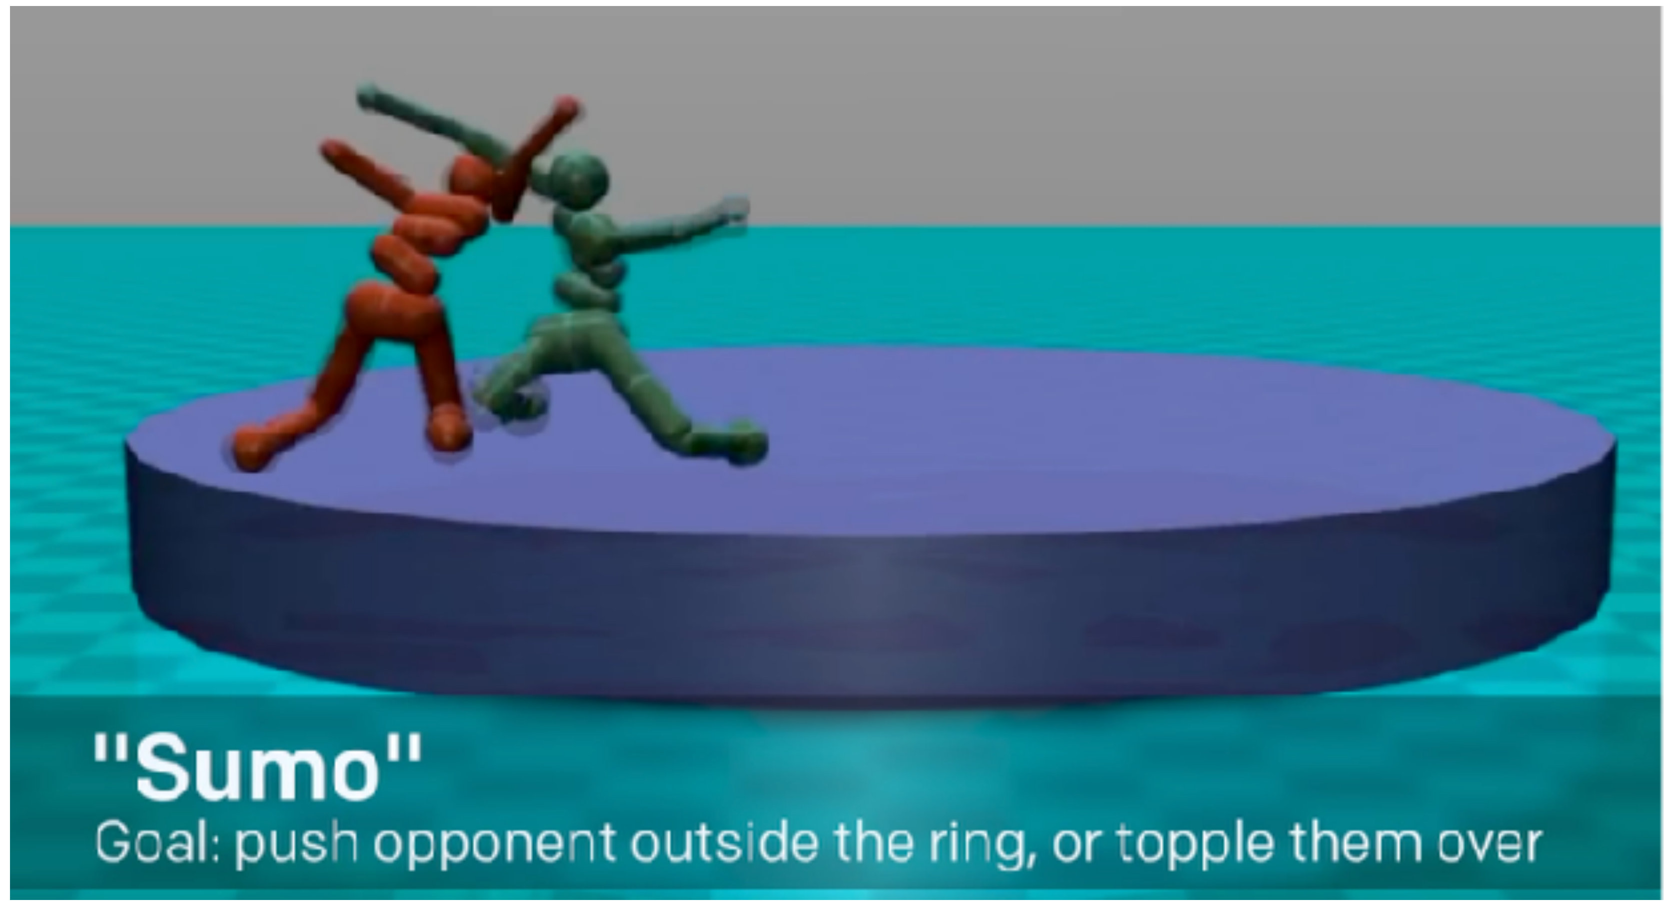

(Bansal et al, 2017)

(Goodfellow 2018)

What can you do with GANs?

- Simulated environments and training data
- Missing data
    - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Model-based optimization
- Automated customization
- Domain adaptation

Autonomous Driving Data

Input labels

Synthesized image

(Wang et al., 2017)

GANs for simulated training data

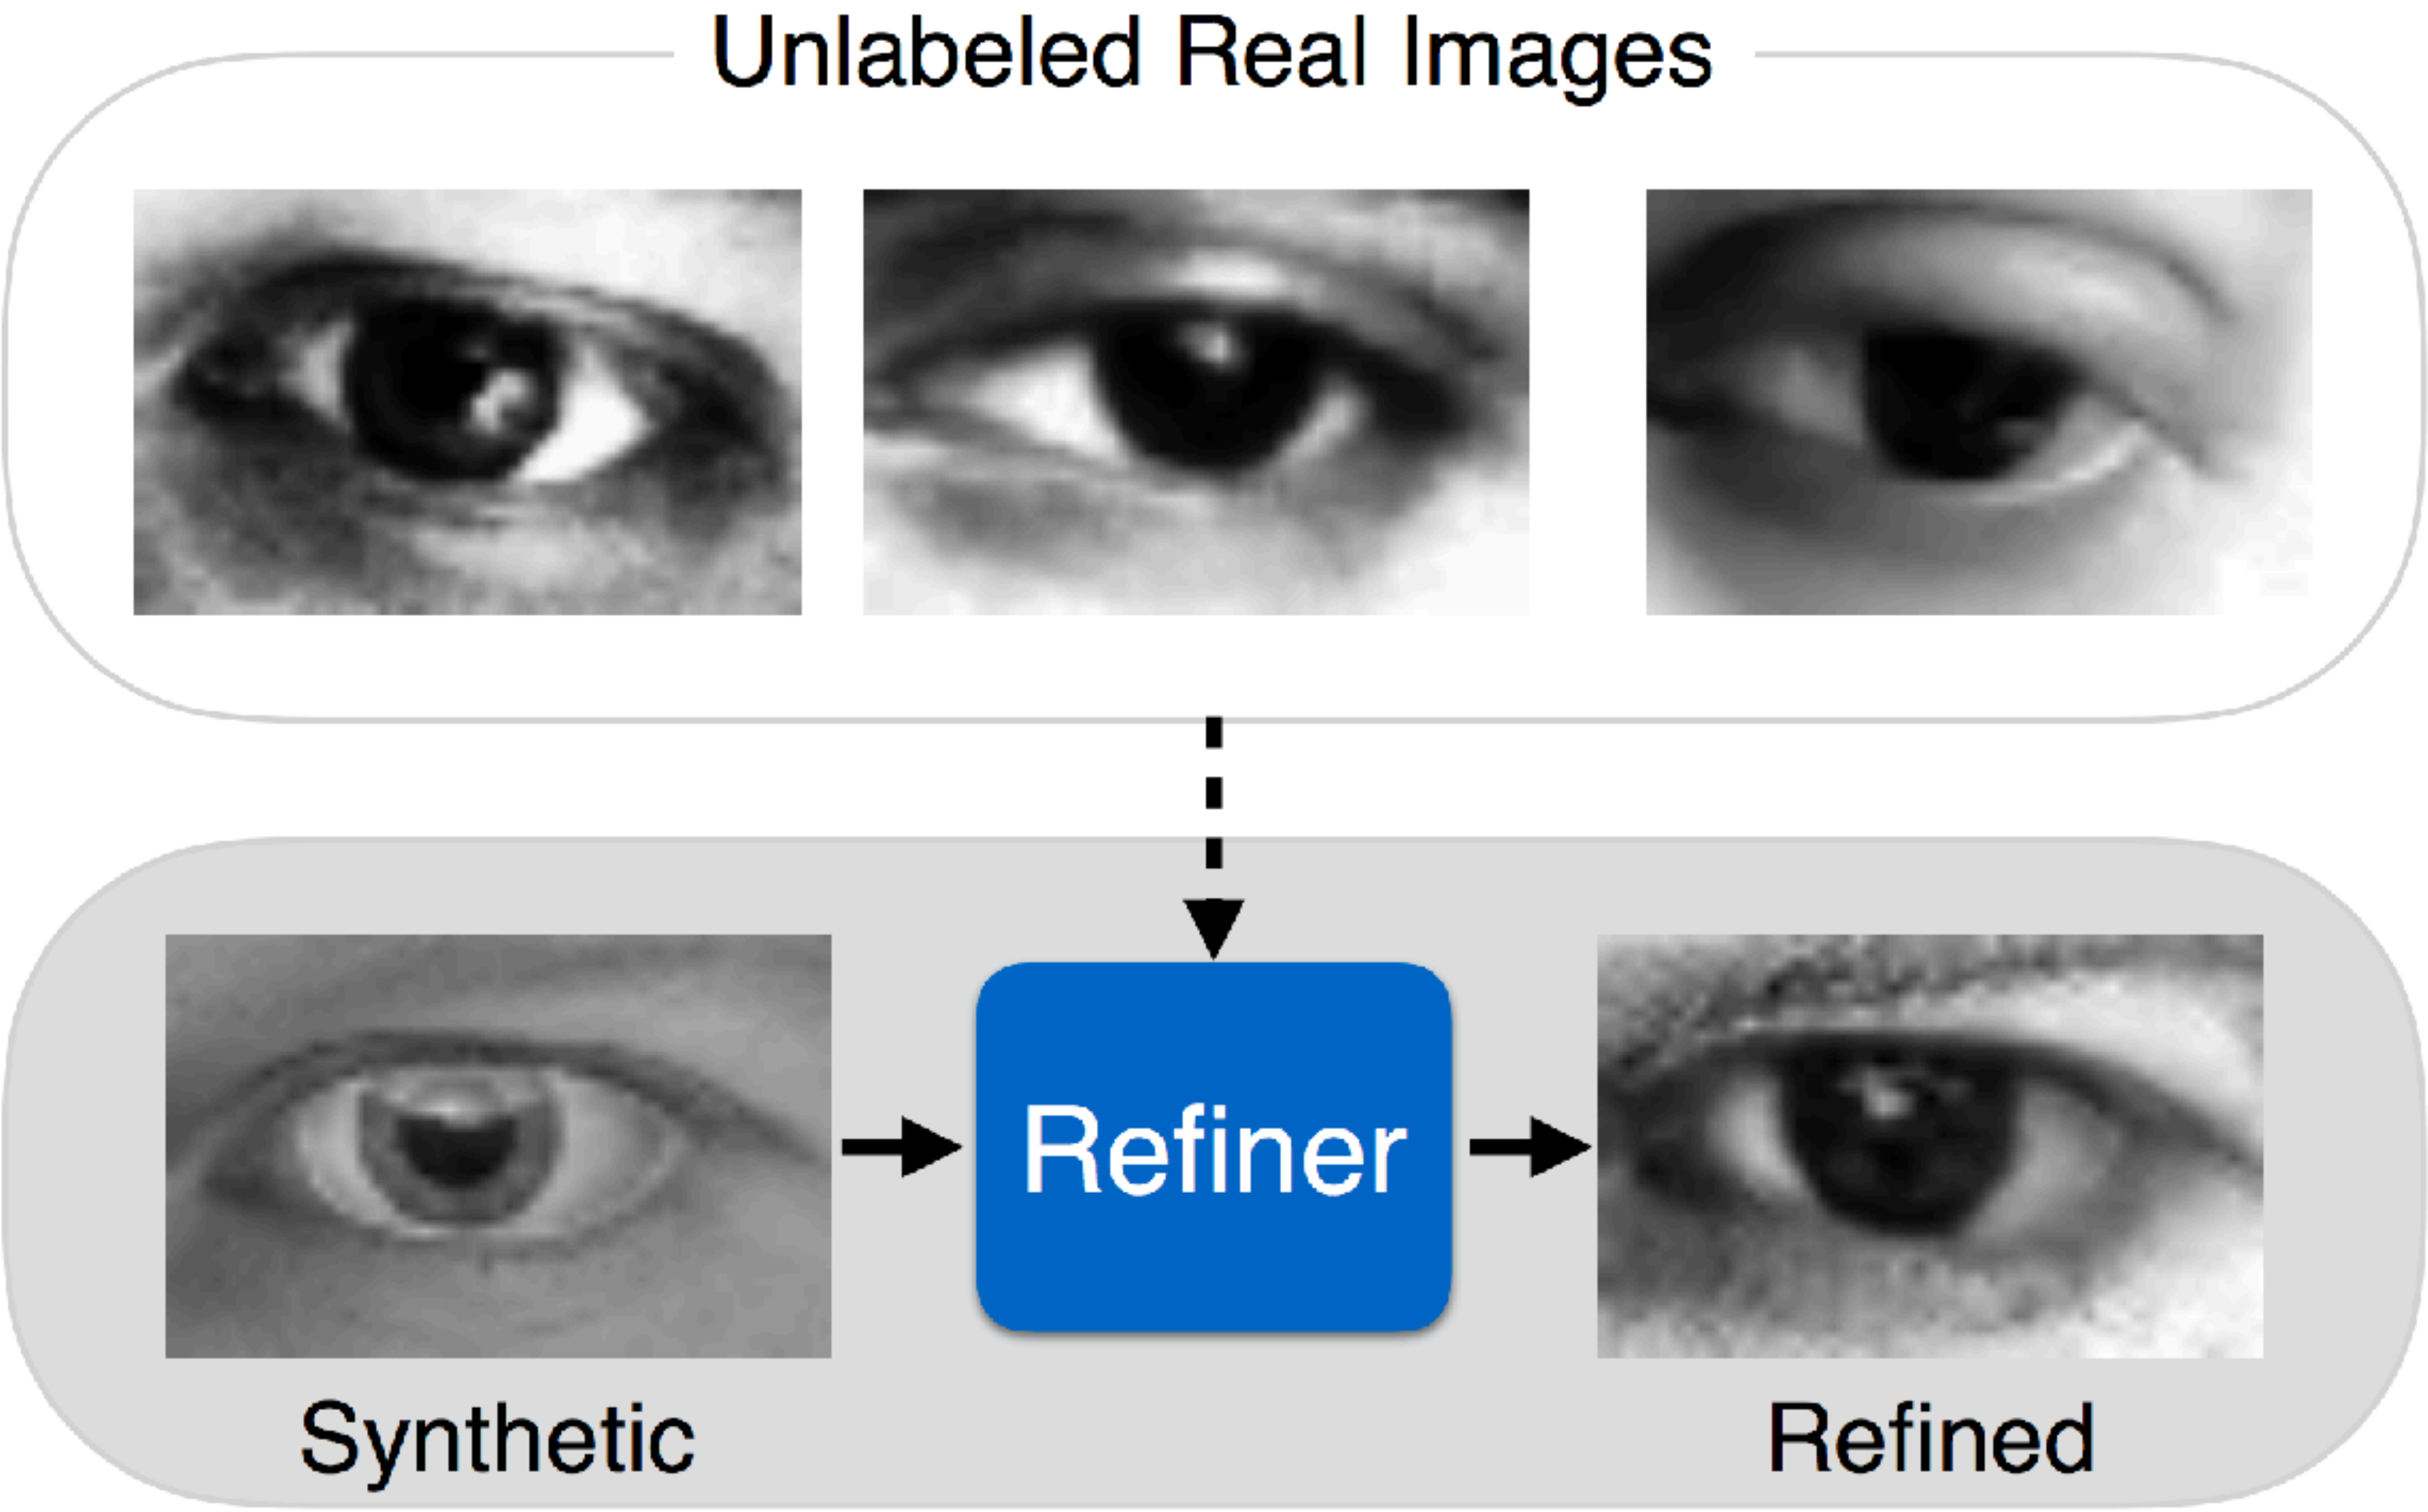

(Shrivastava et al., 2016)

What can you do with GANs?

- Simulated environments and training data
- Missing data
    - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Automated customization
- Domain adaptation

What is in this image?

(Yeh et al., 2016)

Generative modeling reveals a face

(Yeh et al., 2016)

What can you do with GANs?

- Simulated environments and training data
- Missing data
    - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Model-based optimization
- Automated customization
- Domain adaptation

Supervised Discriminator for Semi-Supervised Learning

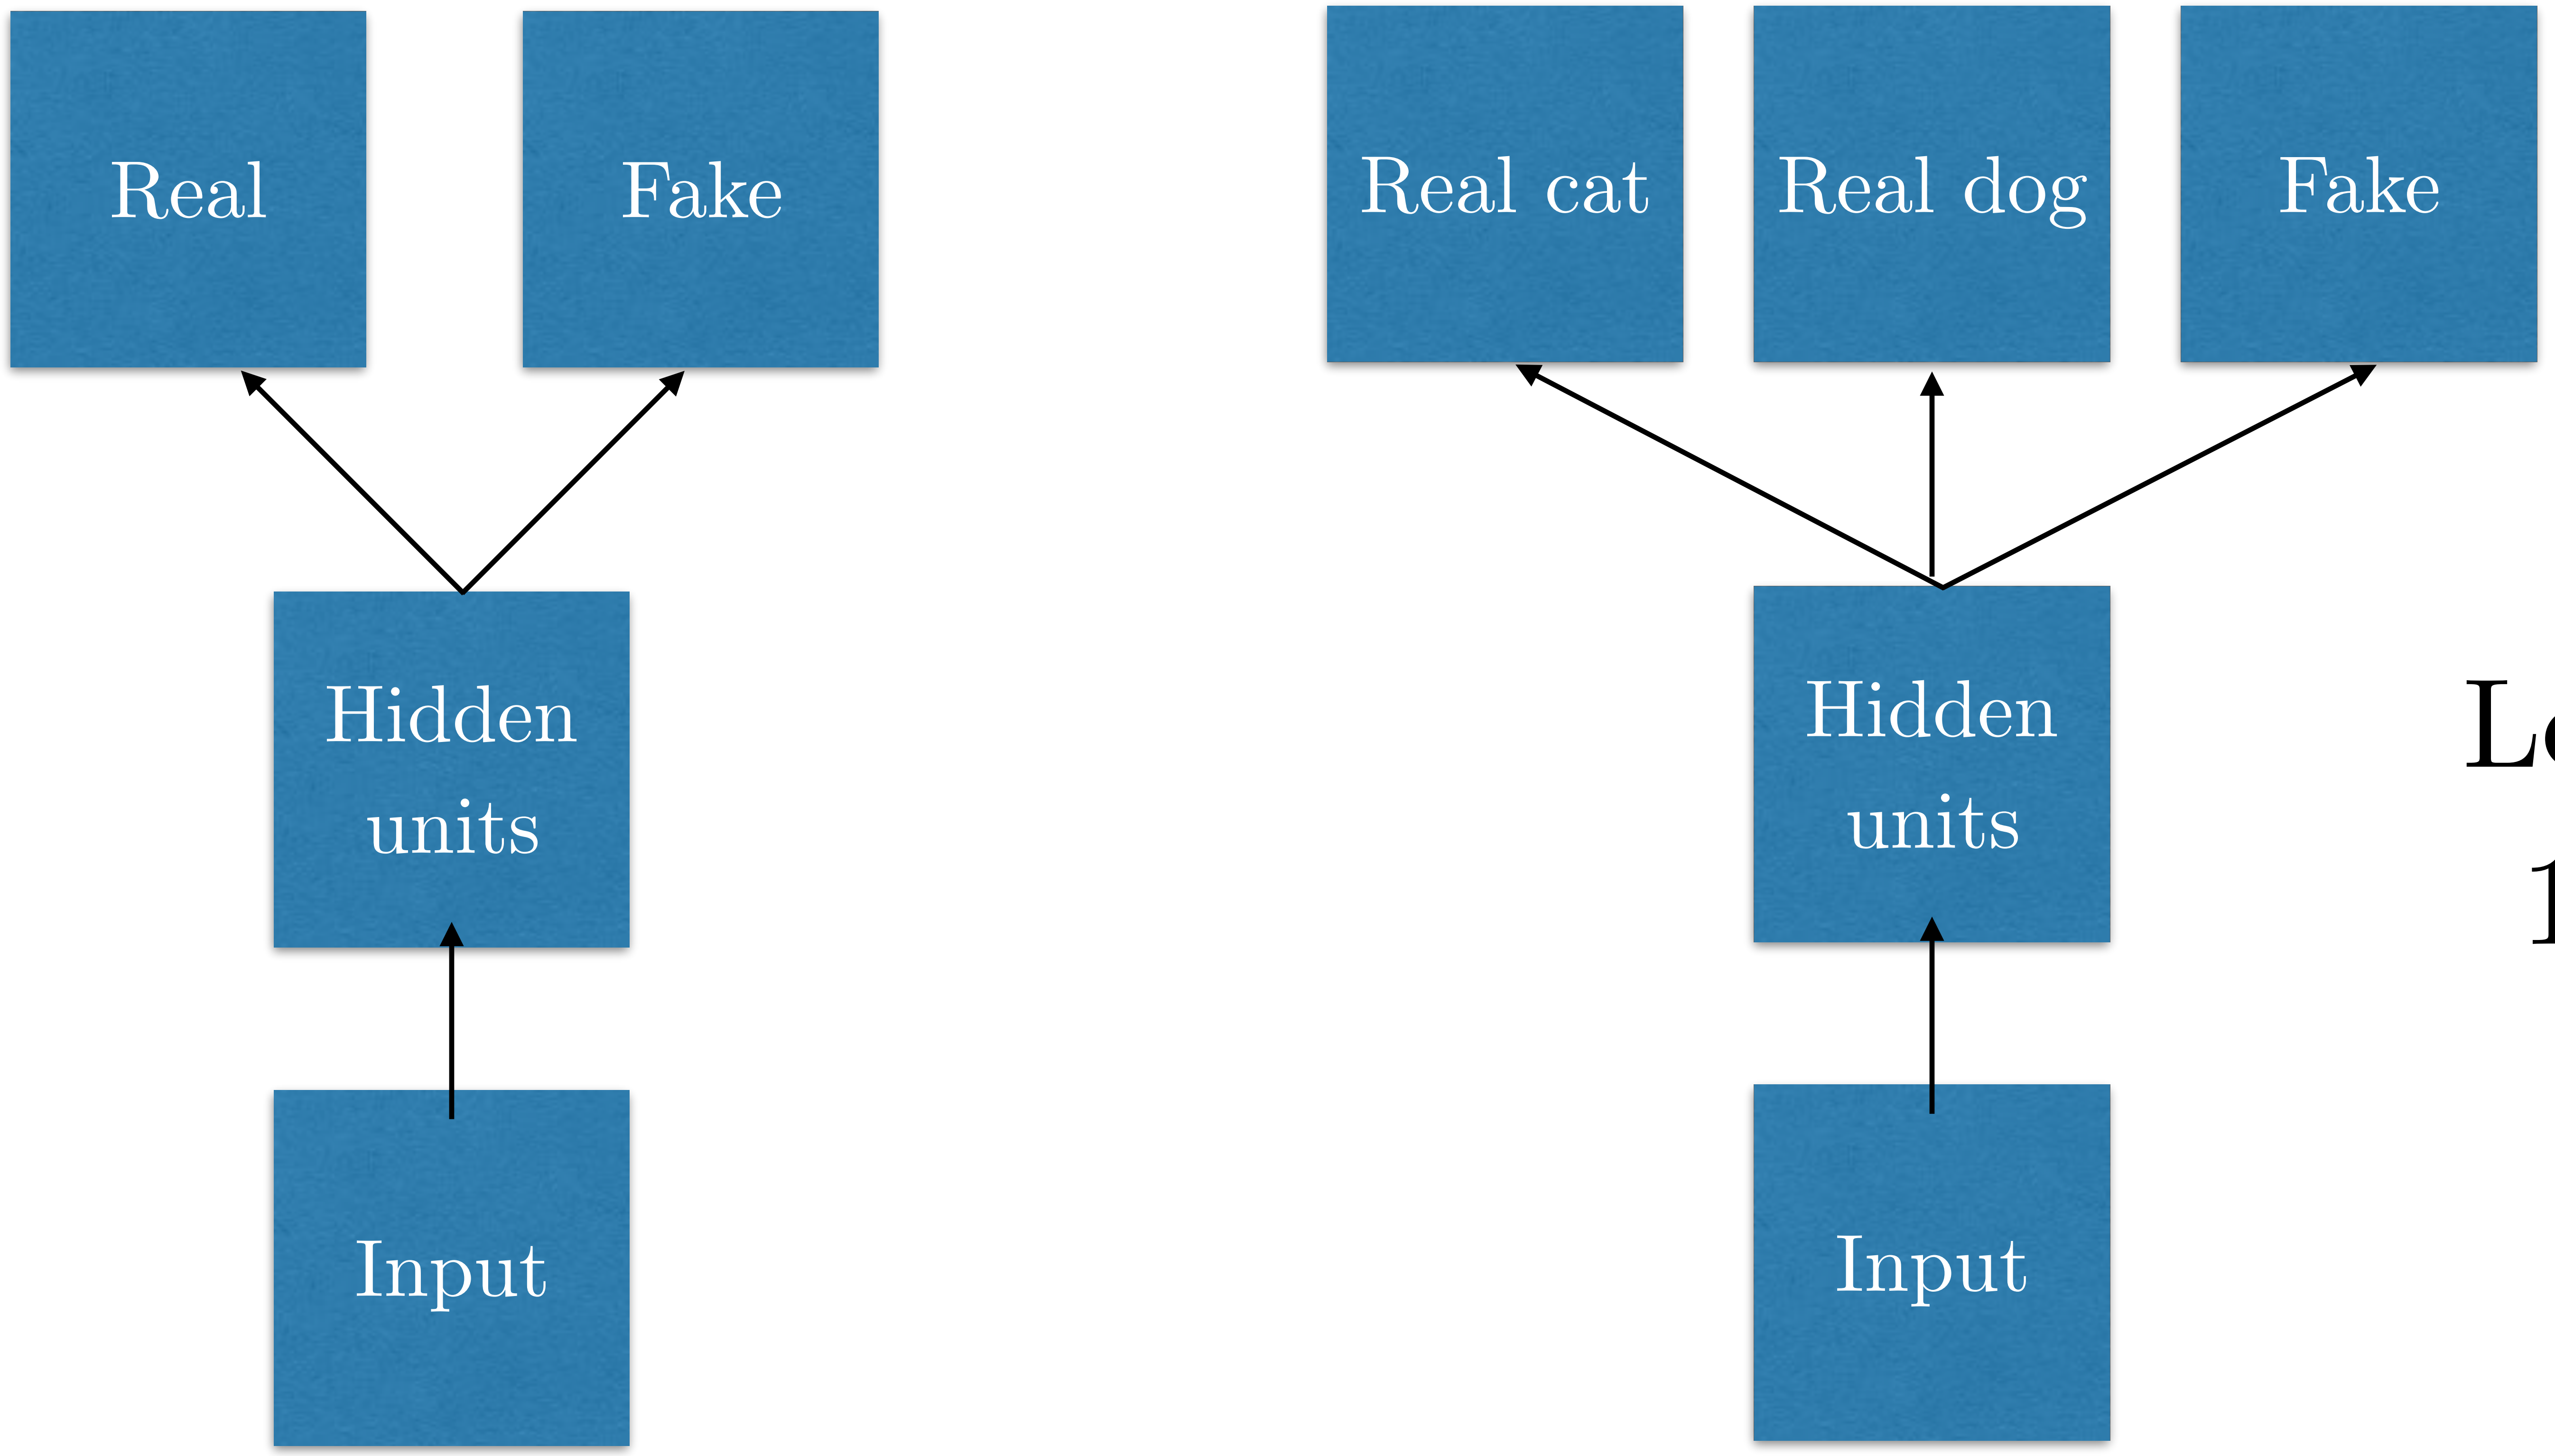

Learn to read with 100 labels rather than 60,000

(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)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Model-based optimization
- Automated customization
- Domain adaptation

Next Video Frame Prediction

Ground Truth

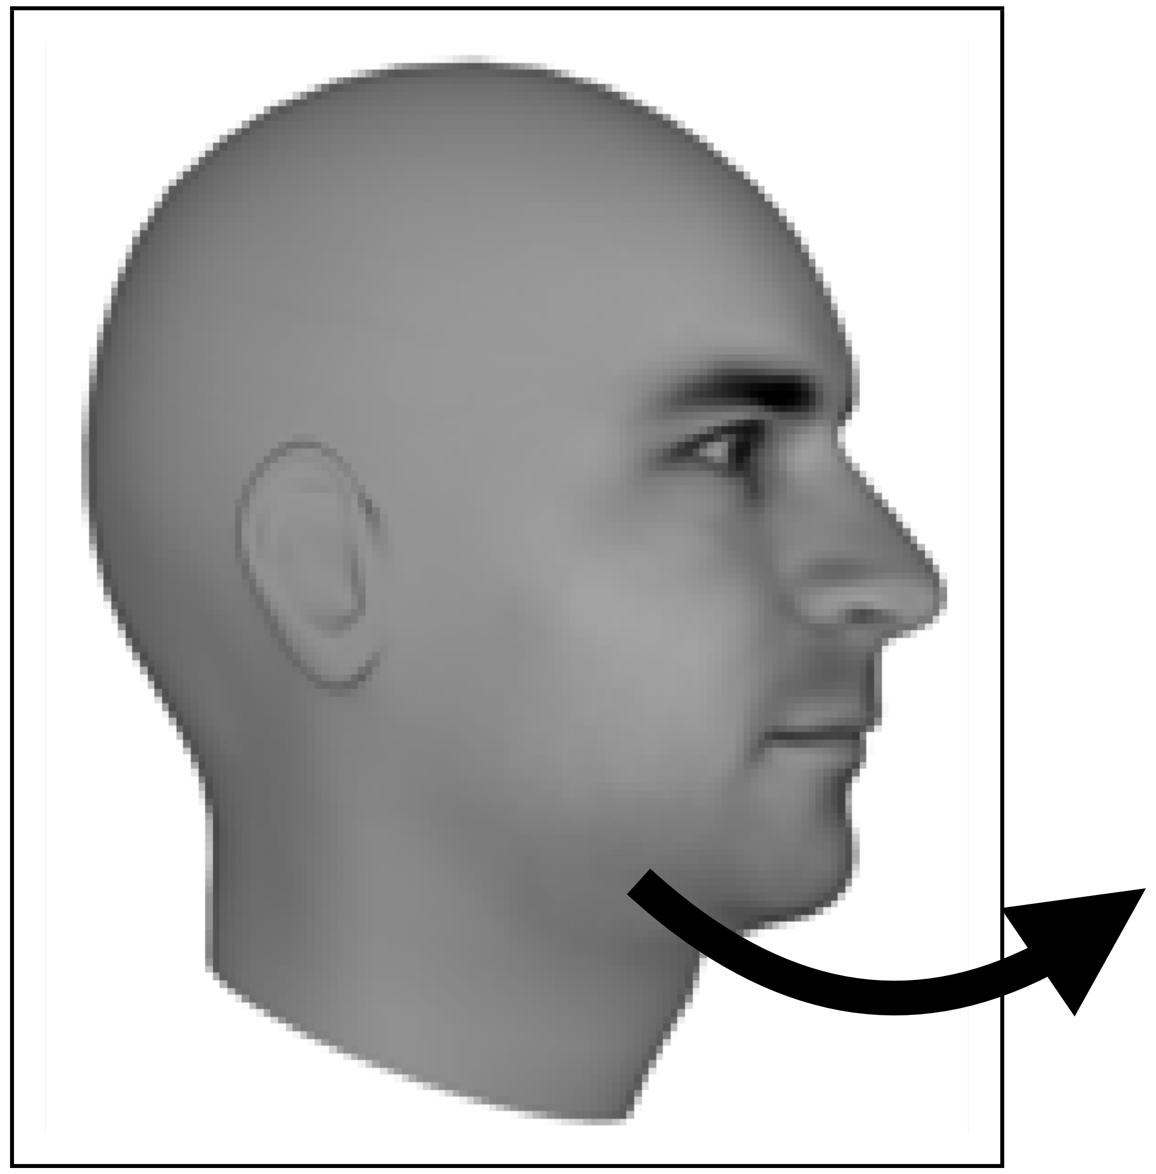

What happens next?

(Lotter et al 2016)

Next Video Frame Prediction

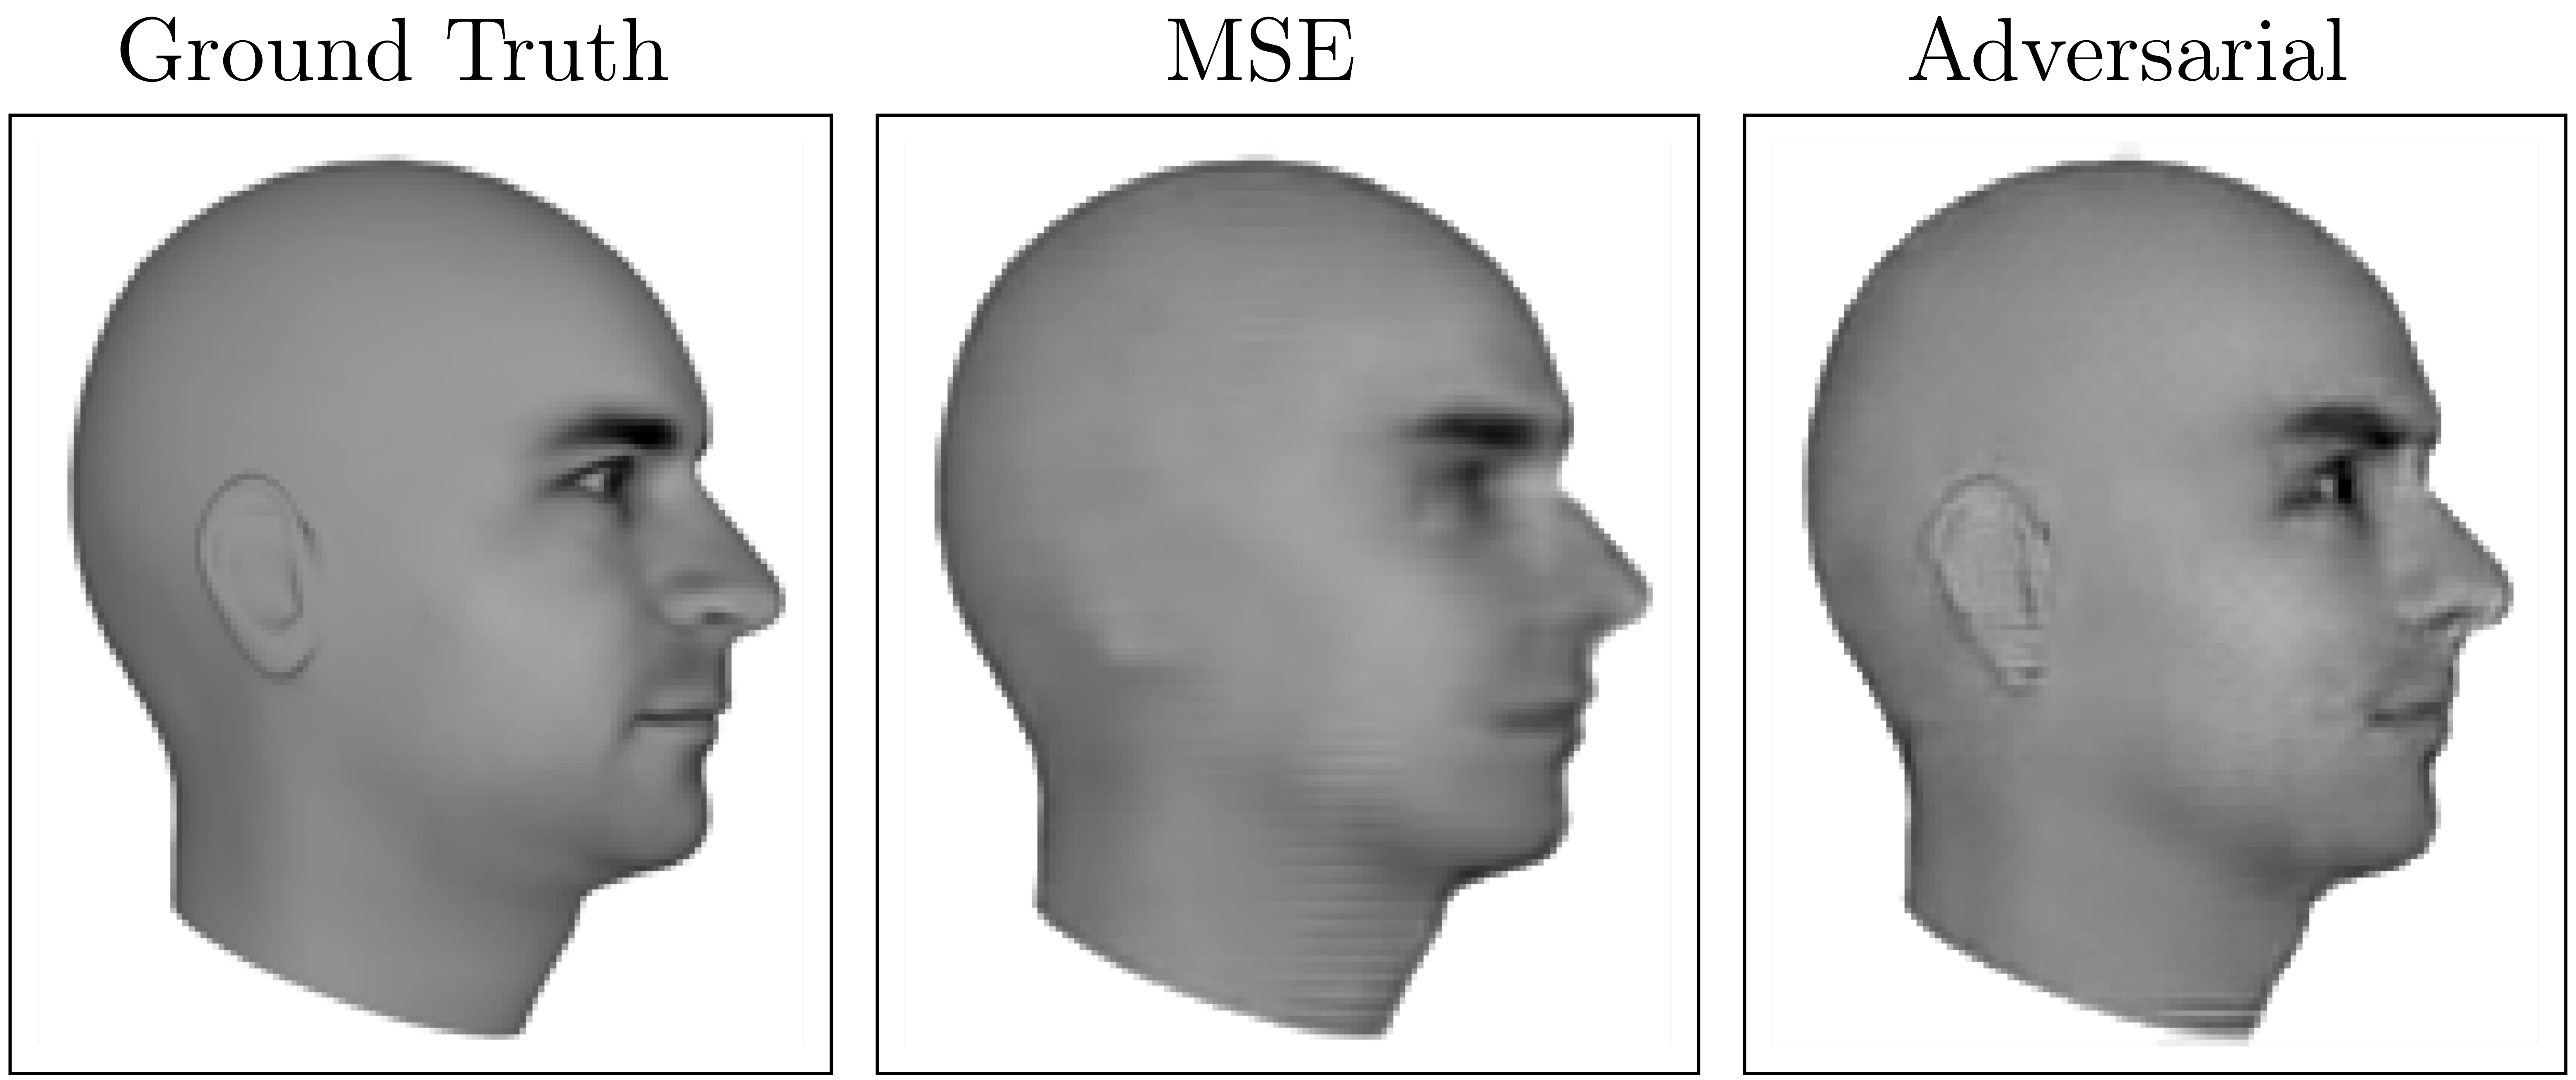

(Lotter et al 2016)

Next Video Frame(s) Prediction

Mean Squared Error

Mean Absolute Error

Adversarial

(Mathieu et al. 2015)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Model-based optimization
- Automated customization
- Domain adaptation

iGAN

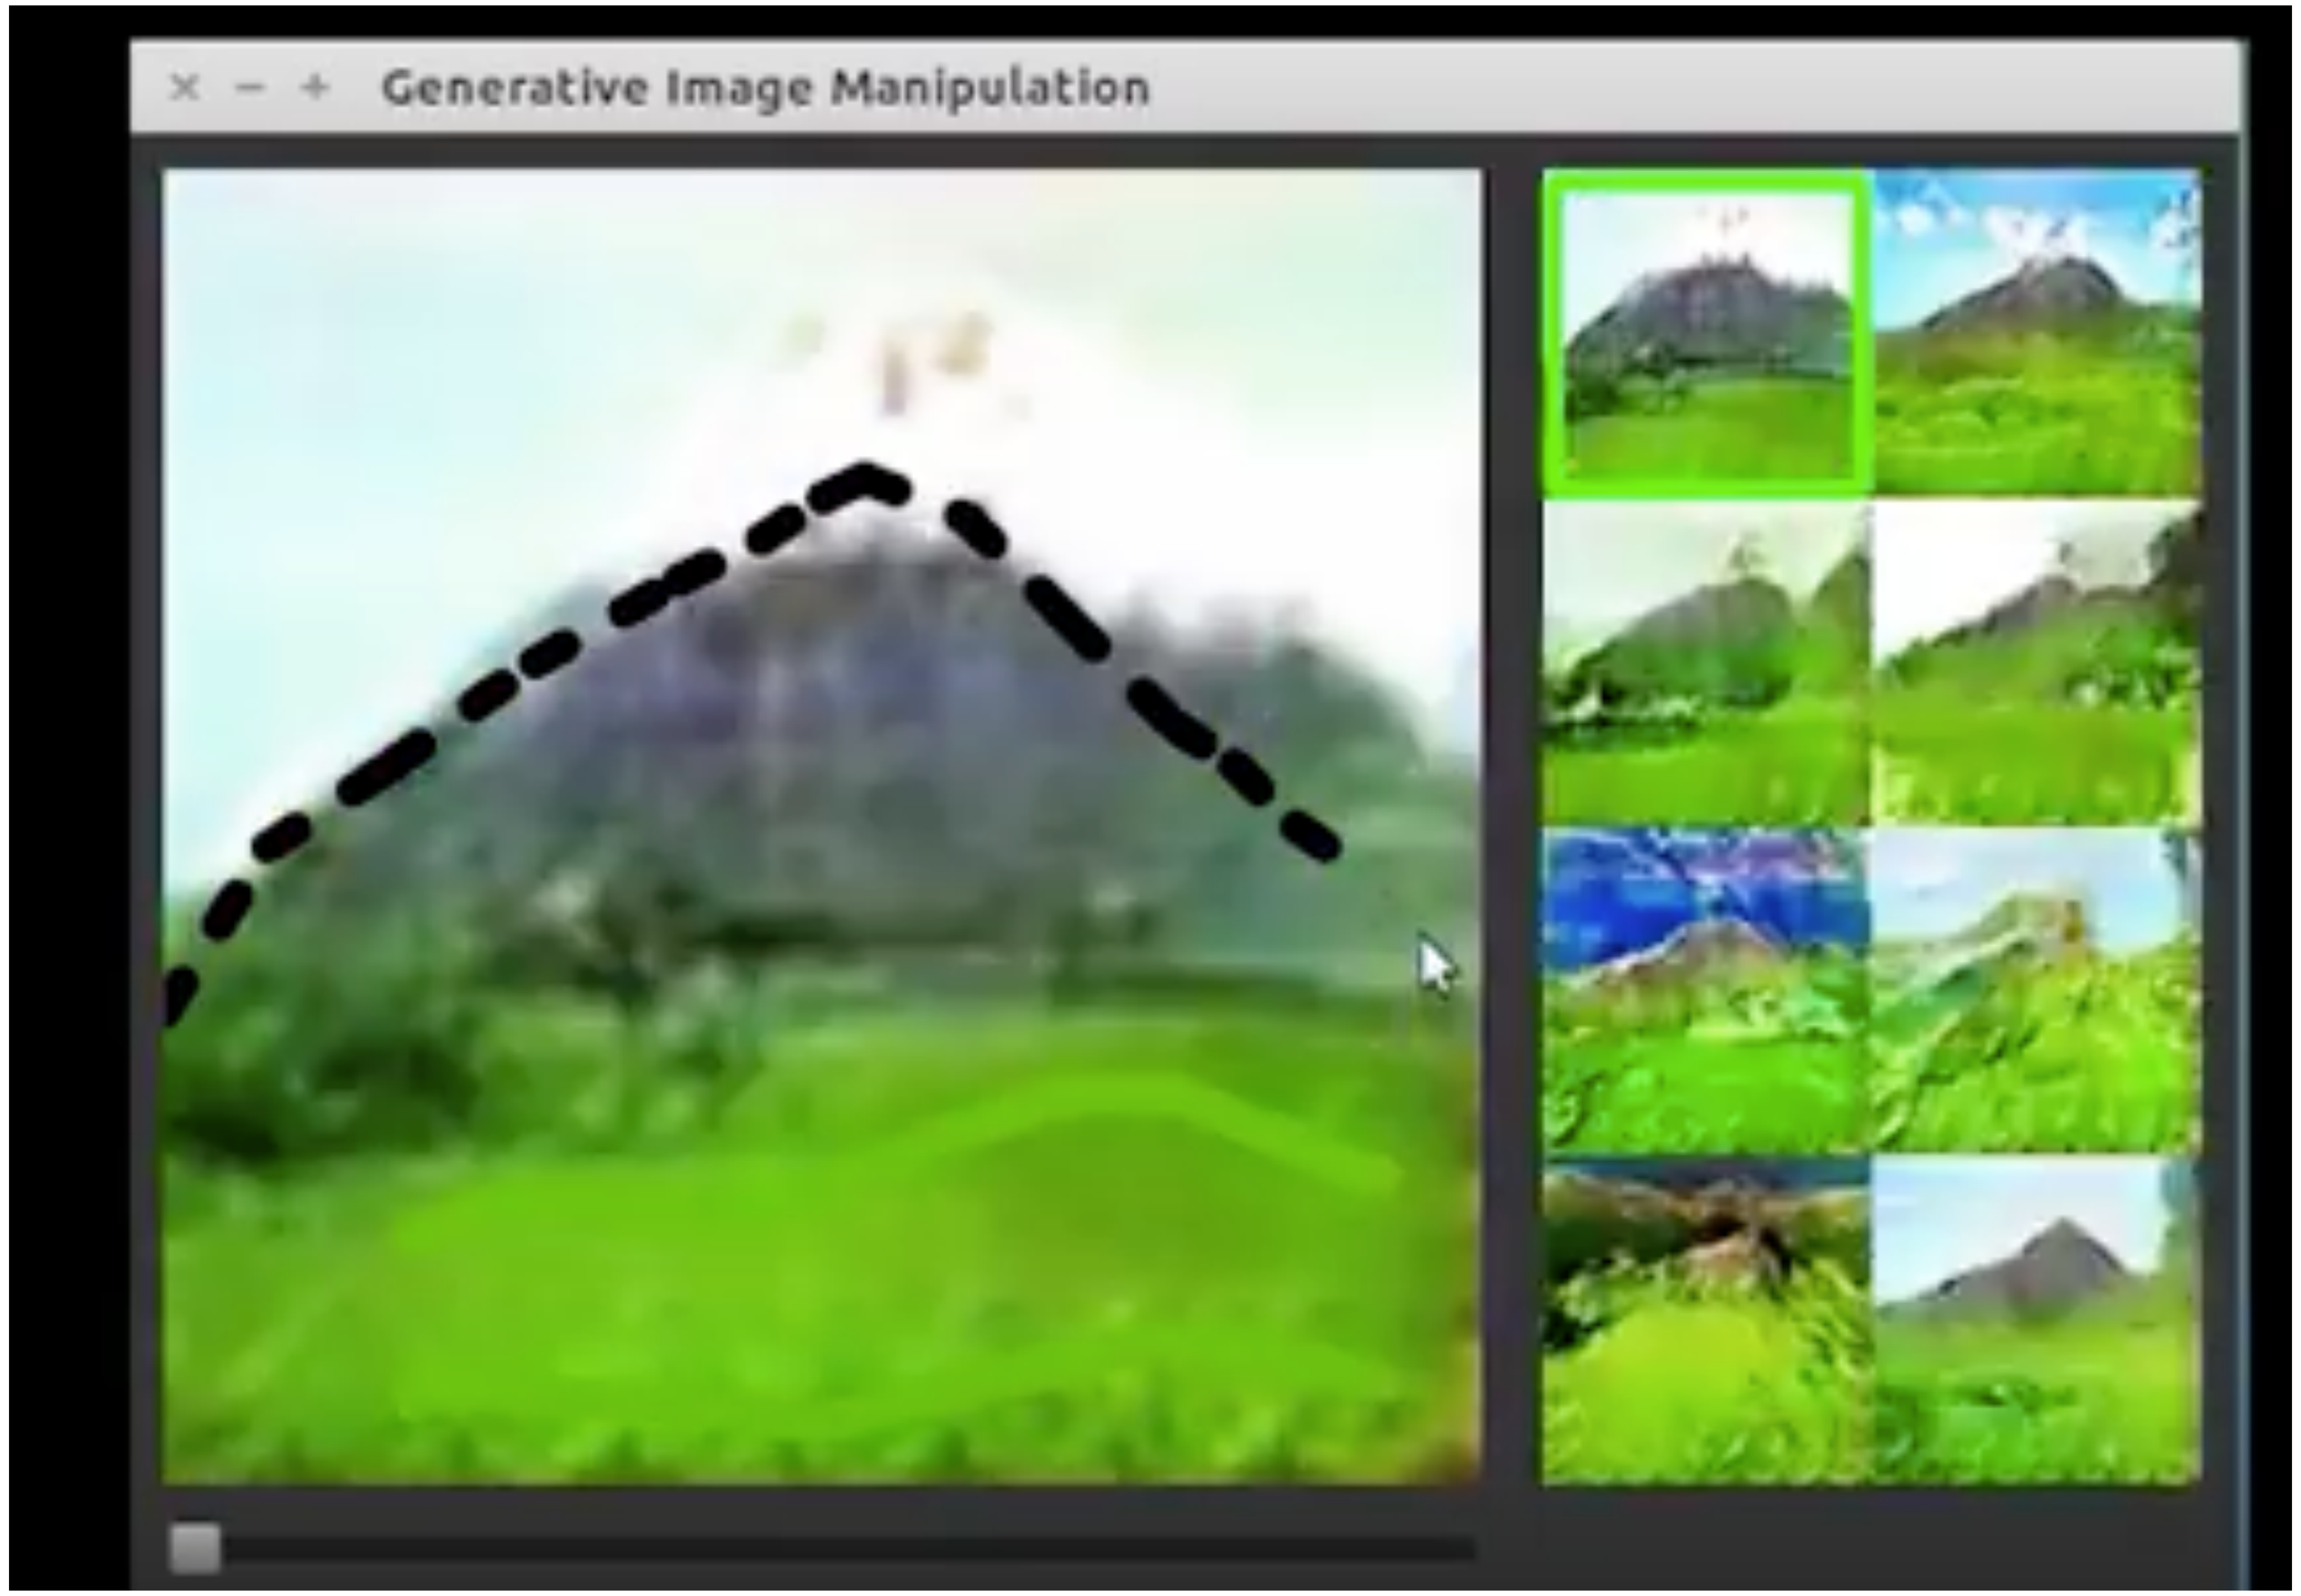

youtube

(Zhu et al., 2016)

Image to Image Translation

(Isola et al., 2016)

Unsupervised Image-to-Image Translation

Day to night

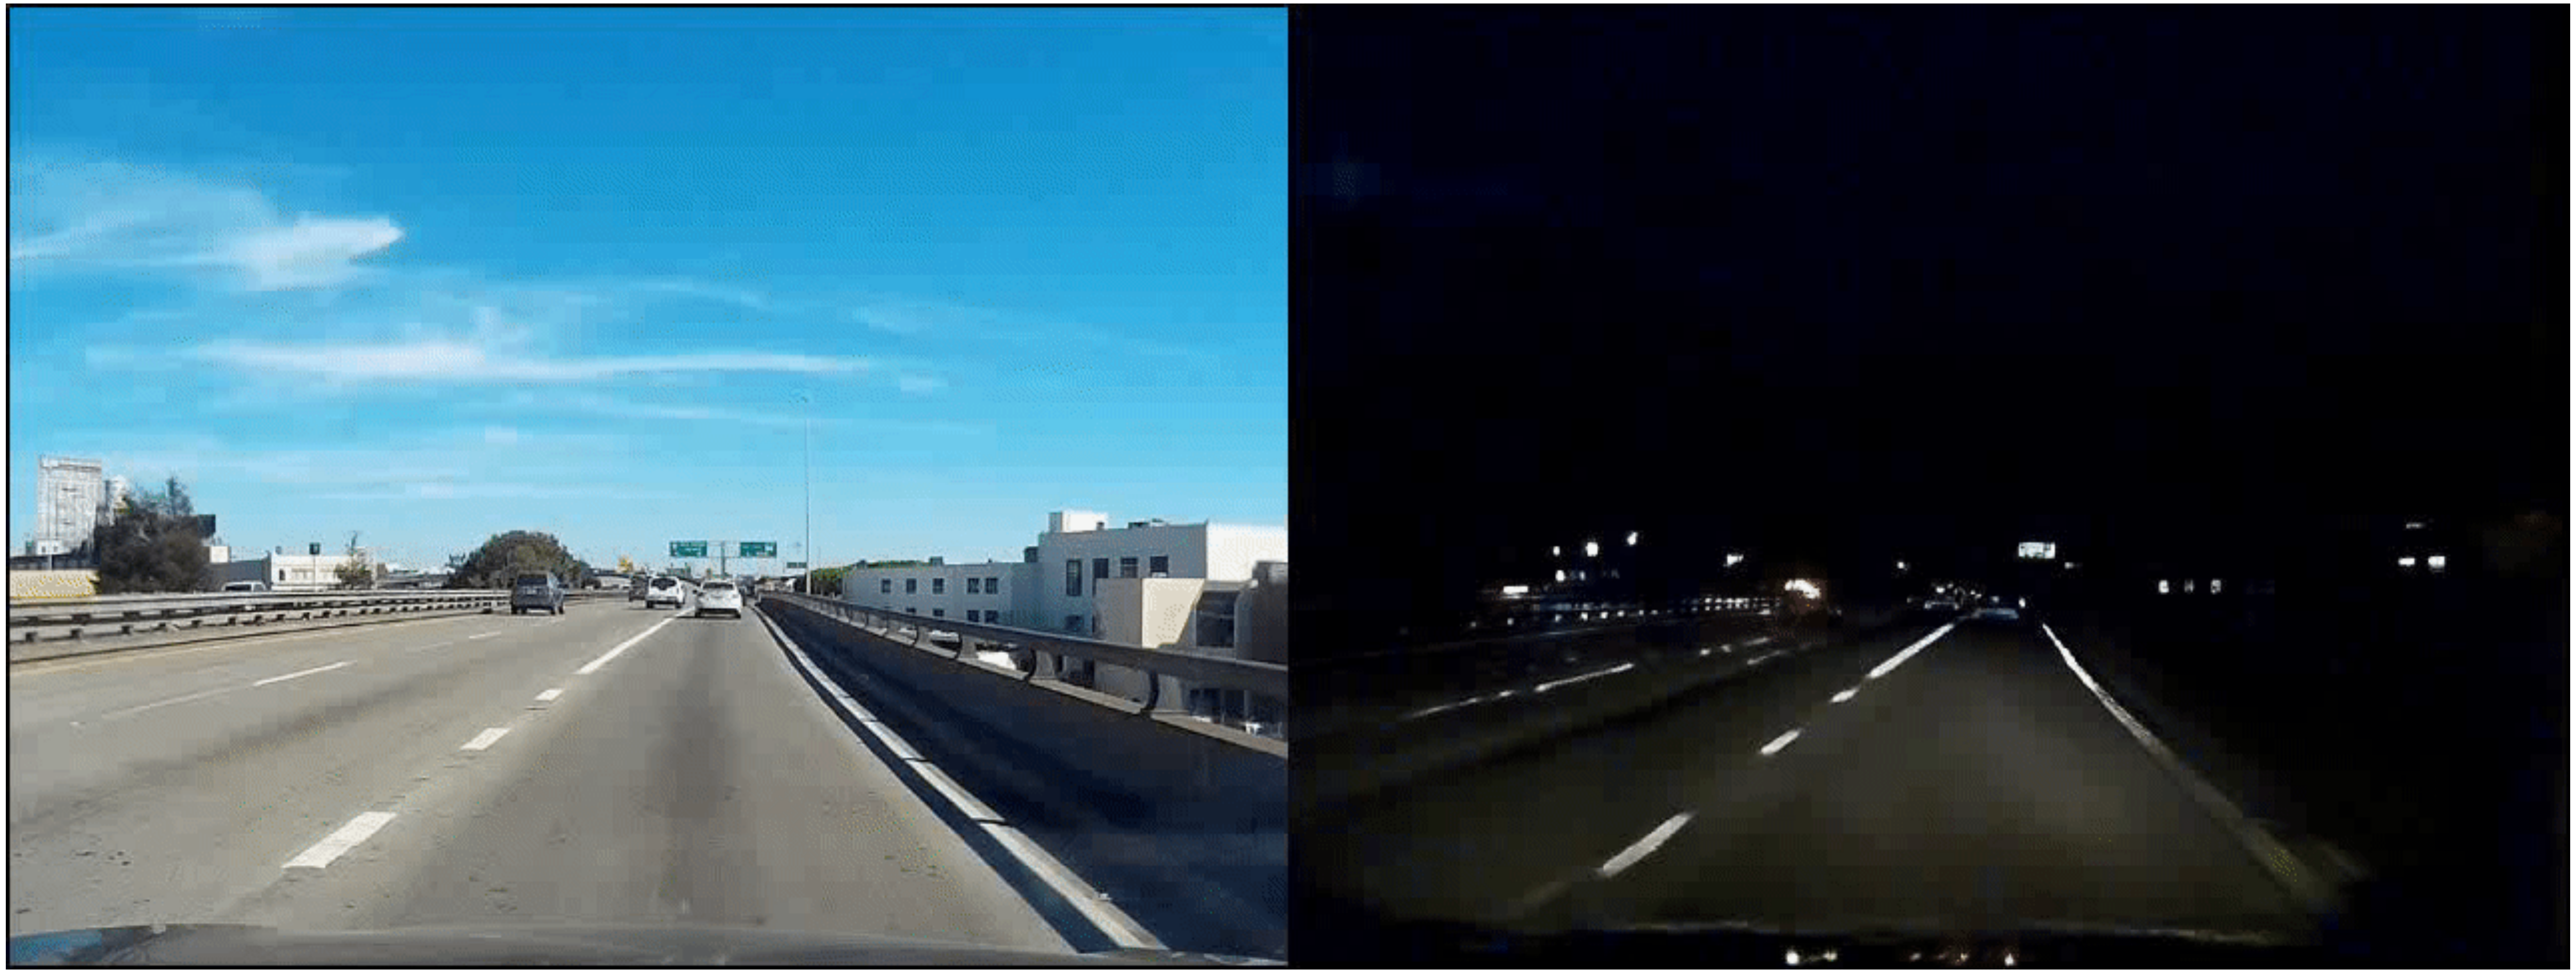

(Liu et al., 2017)

CycleGAN

(Zhu et al., 2017)

Text-to-Image Synthesis

This bird has a yellow belly and tarsus, grey back, wings, and brown throat, nape with a black face

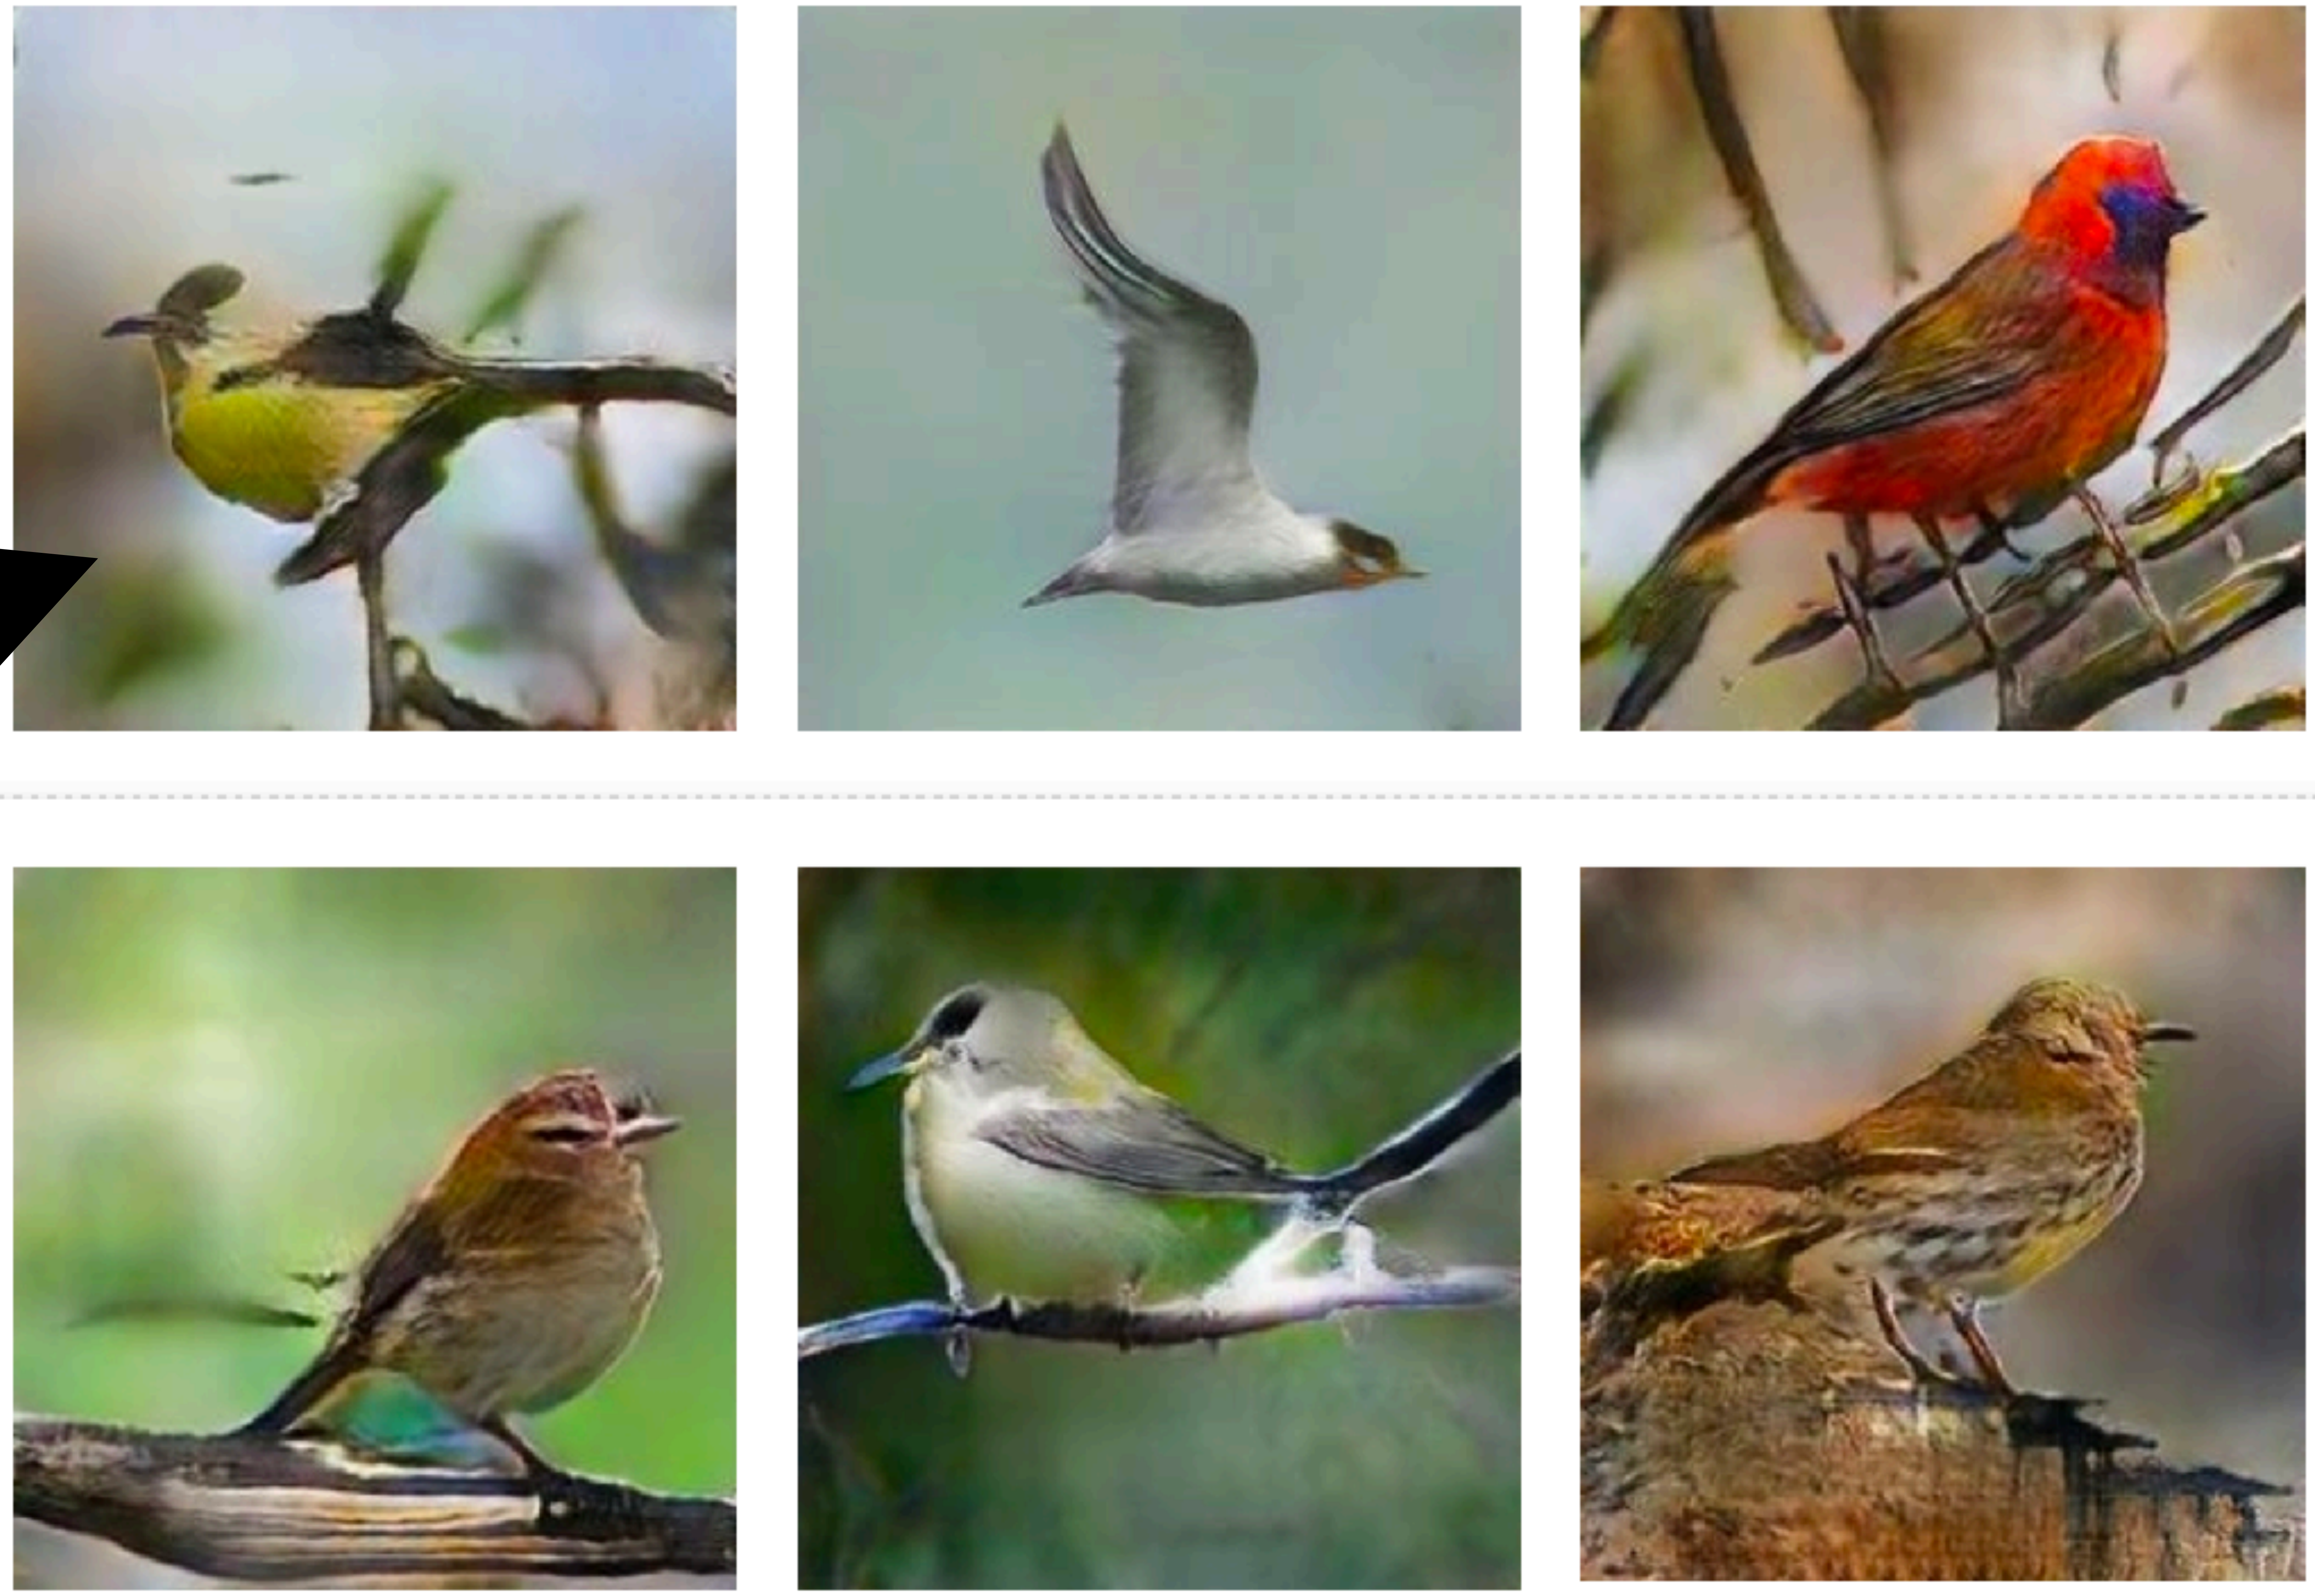

(Zhang et al., 2016)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Model-based optimization
- Automated customization
- Domain adaptation

Designing DNA to optimize protein binding

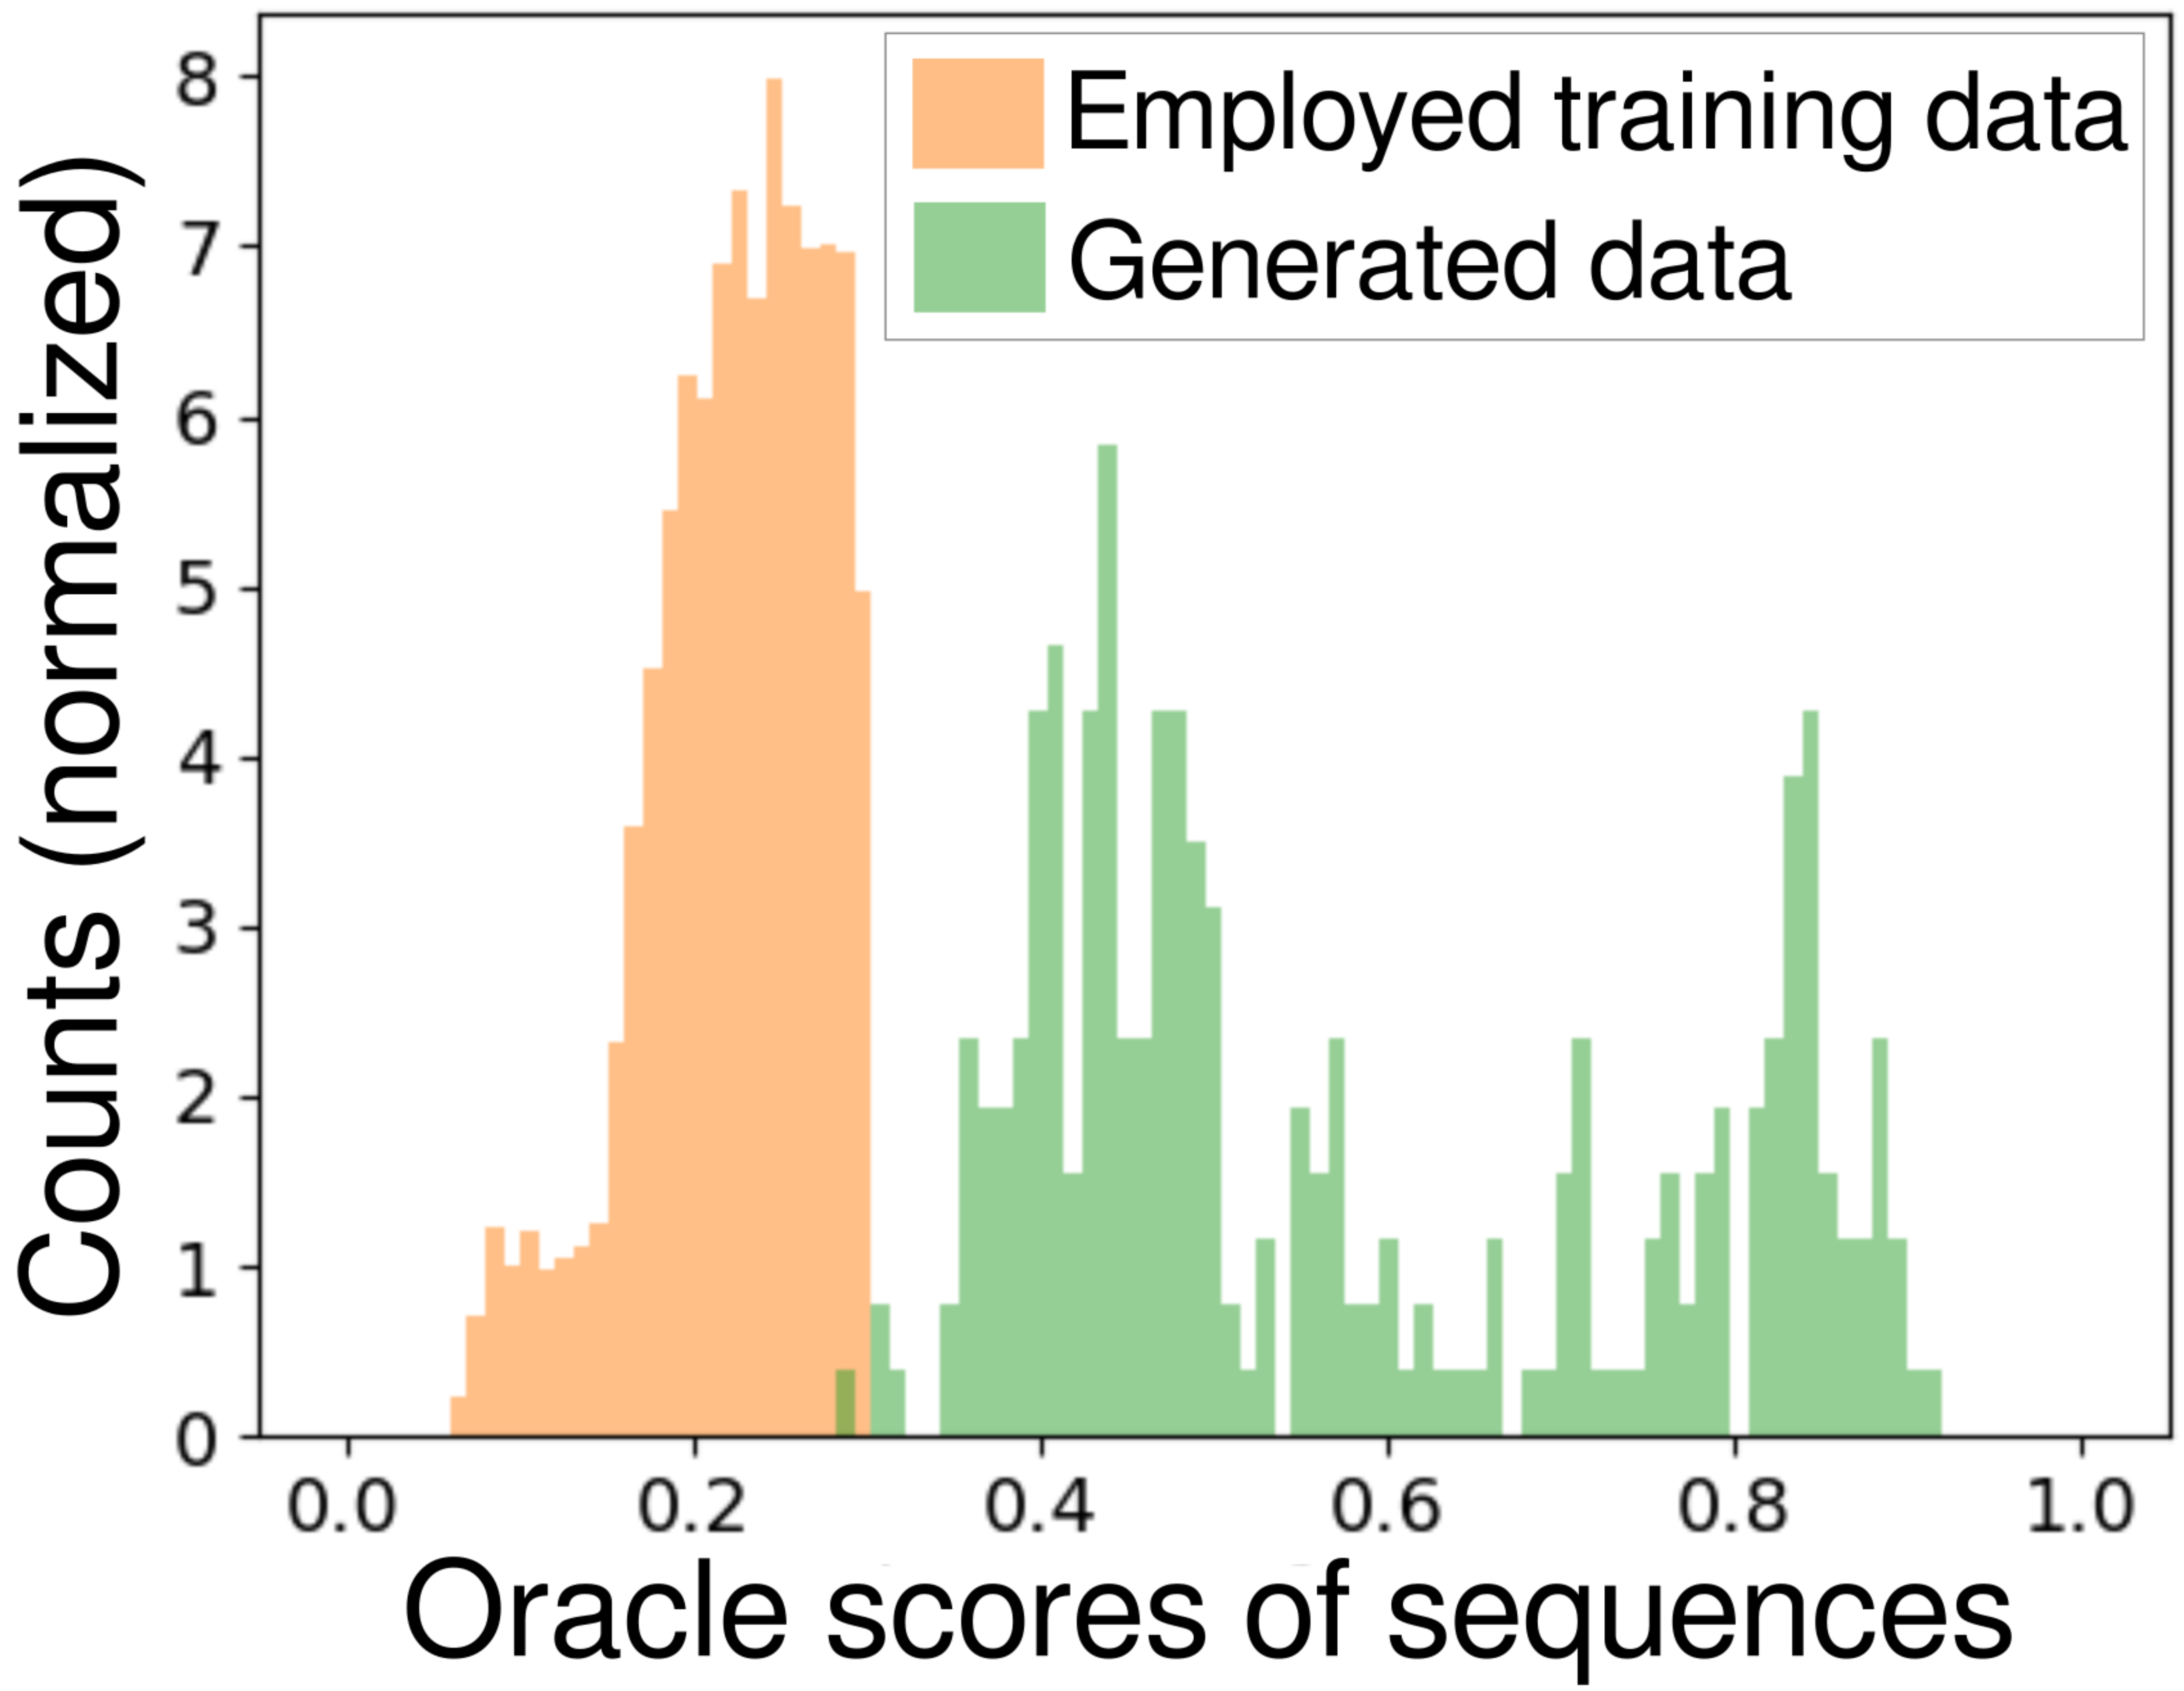

(Killoran et al, 2017)

What can you do with GANs?

- Simulated environments and training data
- Missing data
    - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Automated Customization
- Domain Adaptation

(Goodfellow 2018)

Personalized GANufacturing

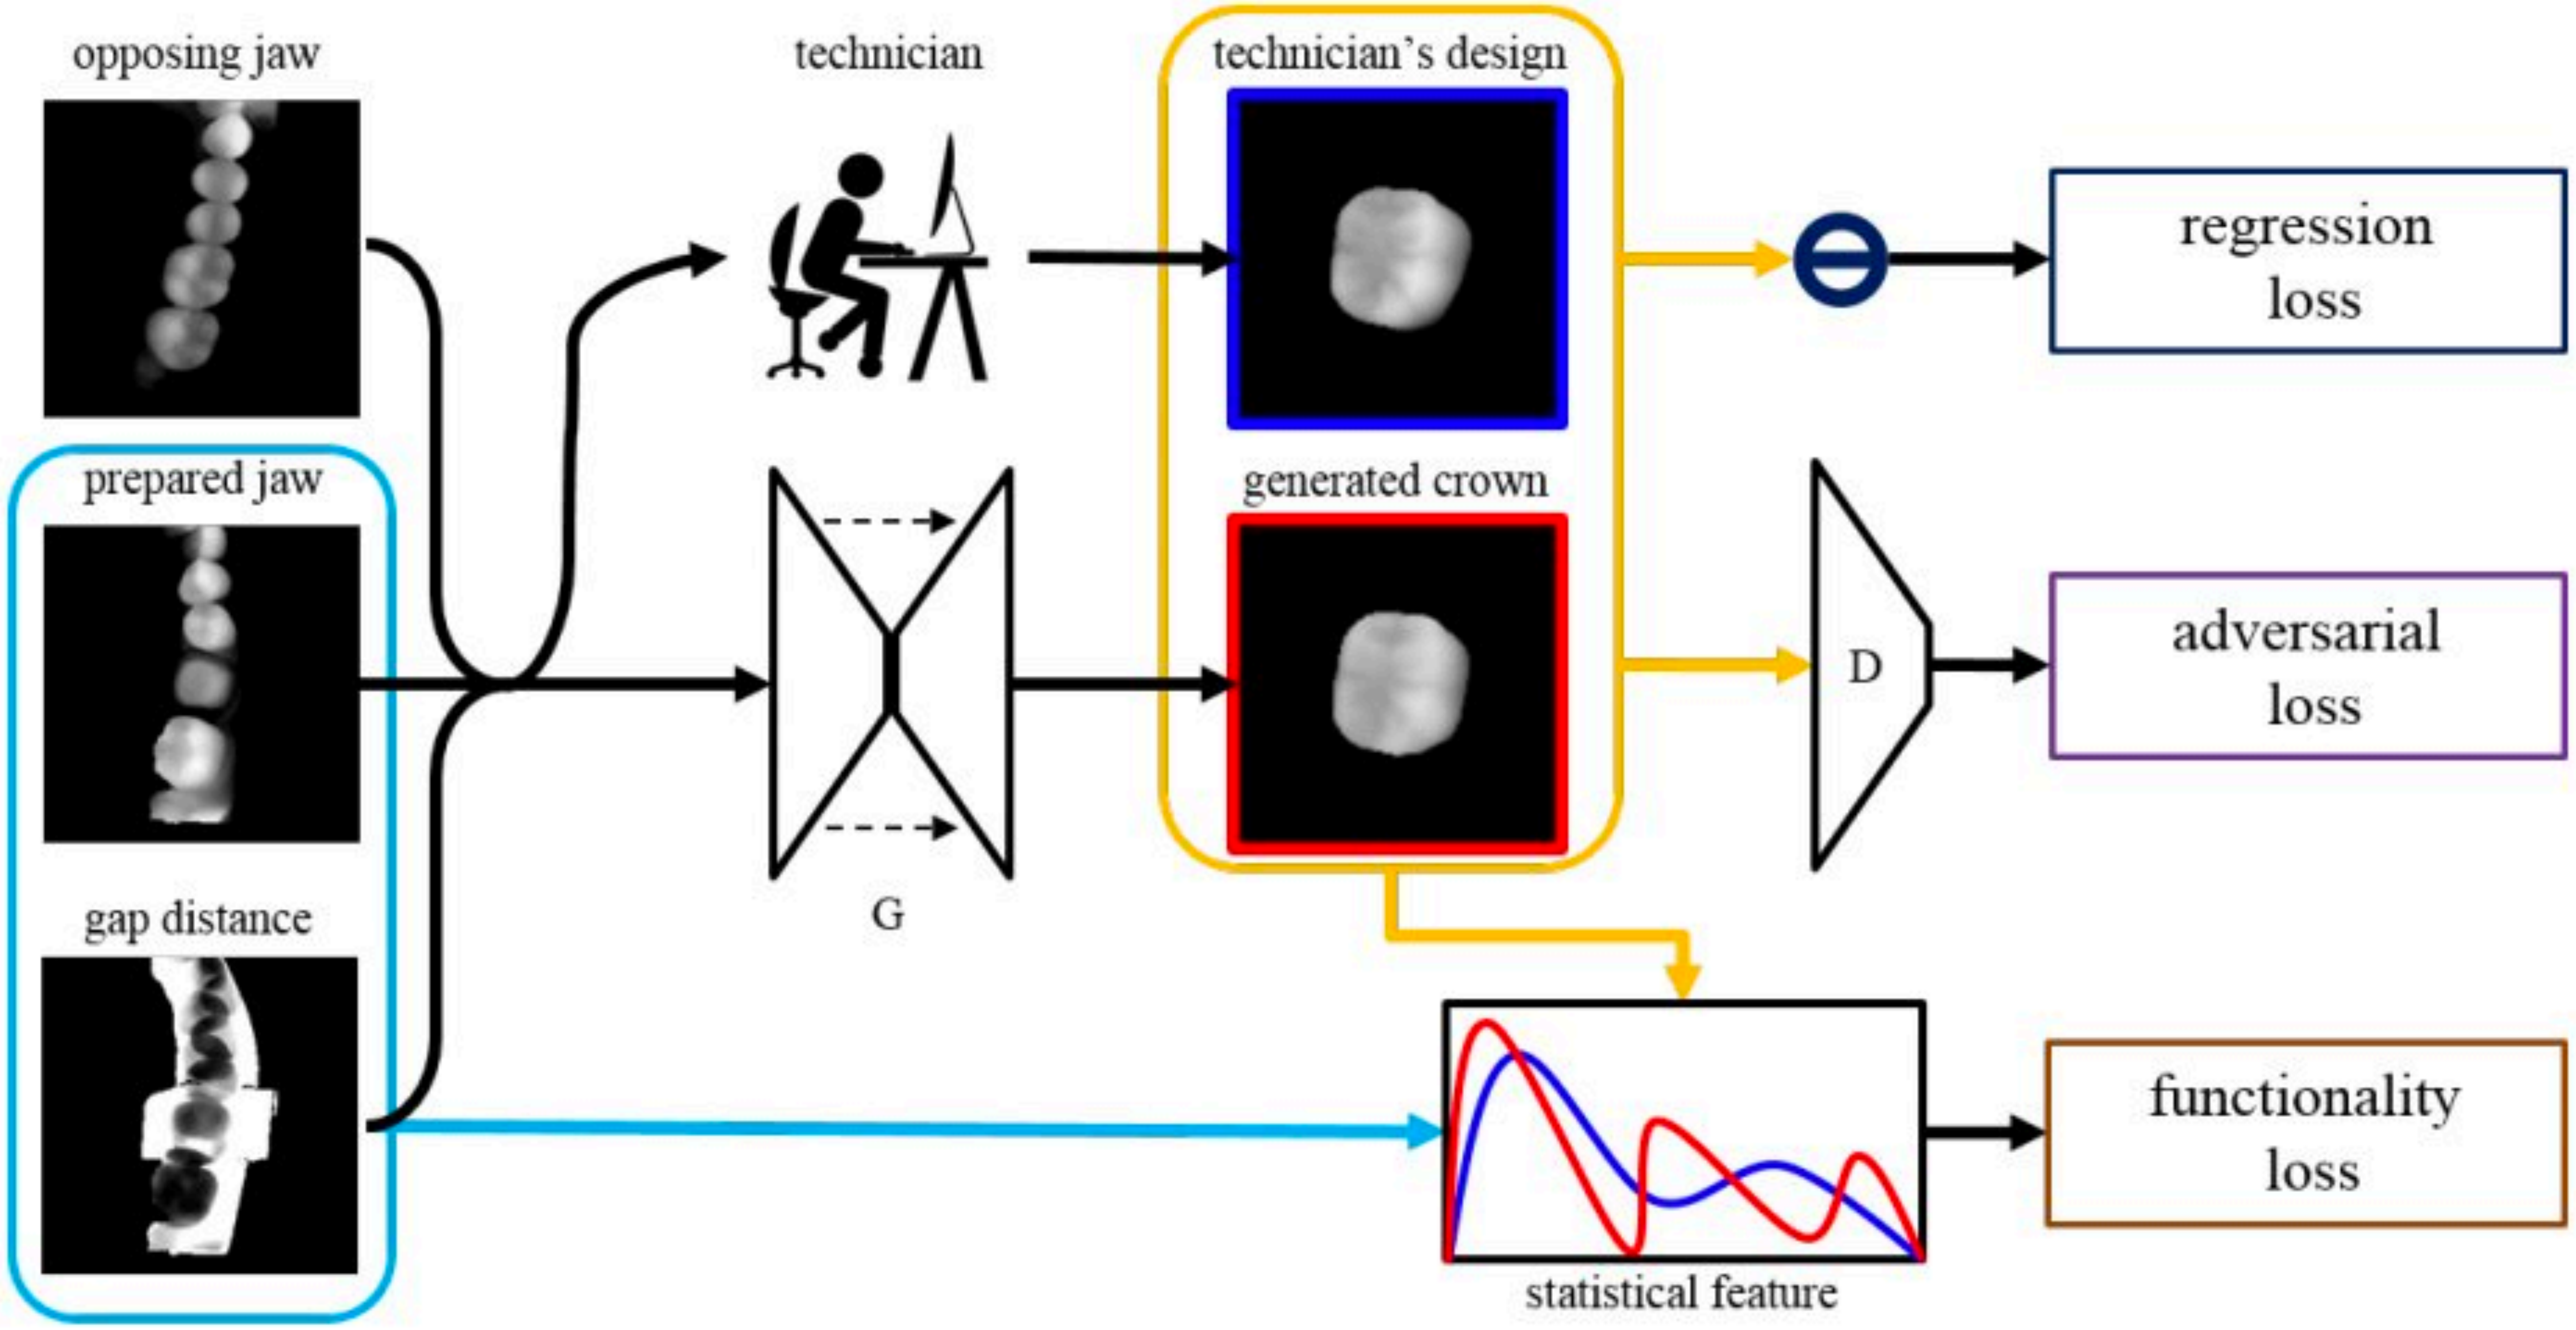

(Hwang et al 2018)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Automated Customization
- Domain Adaptation

Domain Adaptation

- Domain Adversarial Networks (Ganin et al, 2015)

- Professor forcing (Lamb et al, 2016): Domain-Adversarial learning in RNN hidden state

GANs for domain adaptation

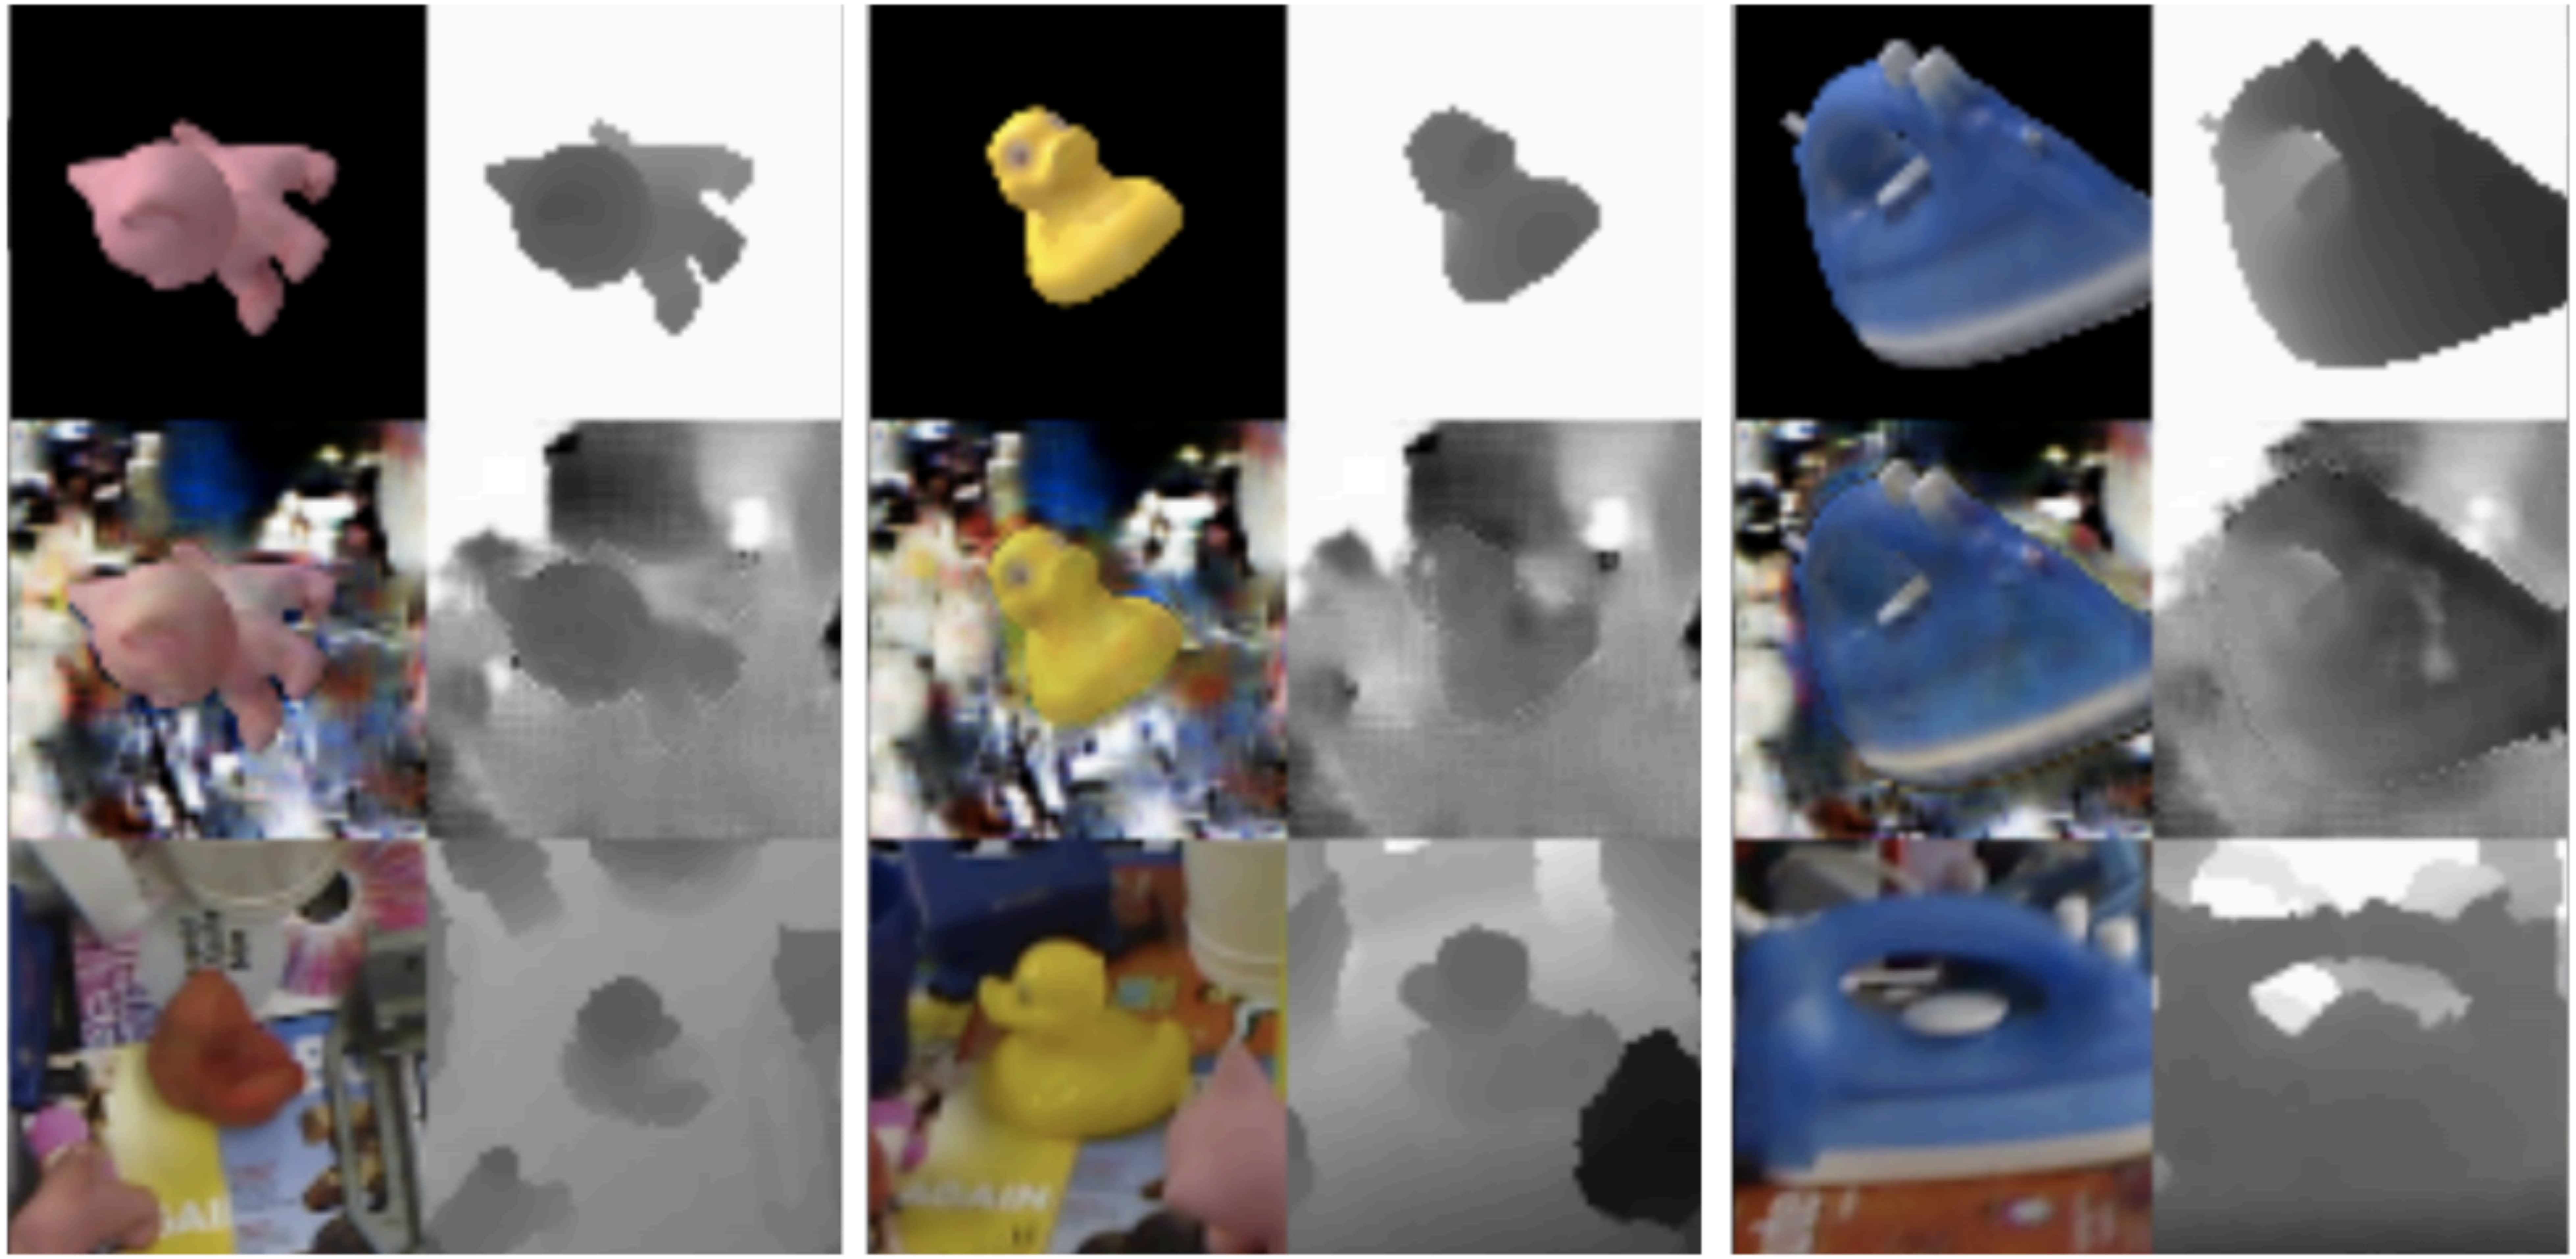

(Bousmalis et al., 2016)

Tips and Tricks

- Spectral normalization (Miyato et al 2017) in both discriminator and generator (Zhang et al 2018)
- Different learning rate for generator and discriminator (Heusel et al 2017)
  - No need to run discriminator more often than generator (Zhang et al 2018)
- Many different loss functions all work well (Lucic et al 2017); spend more time tuning hyperparameters than trying different losses

Track updates at the GAN Zoo

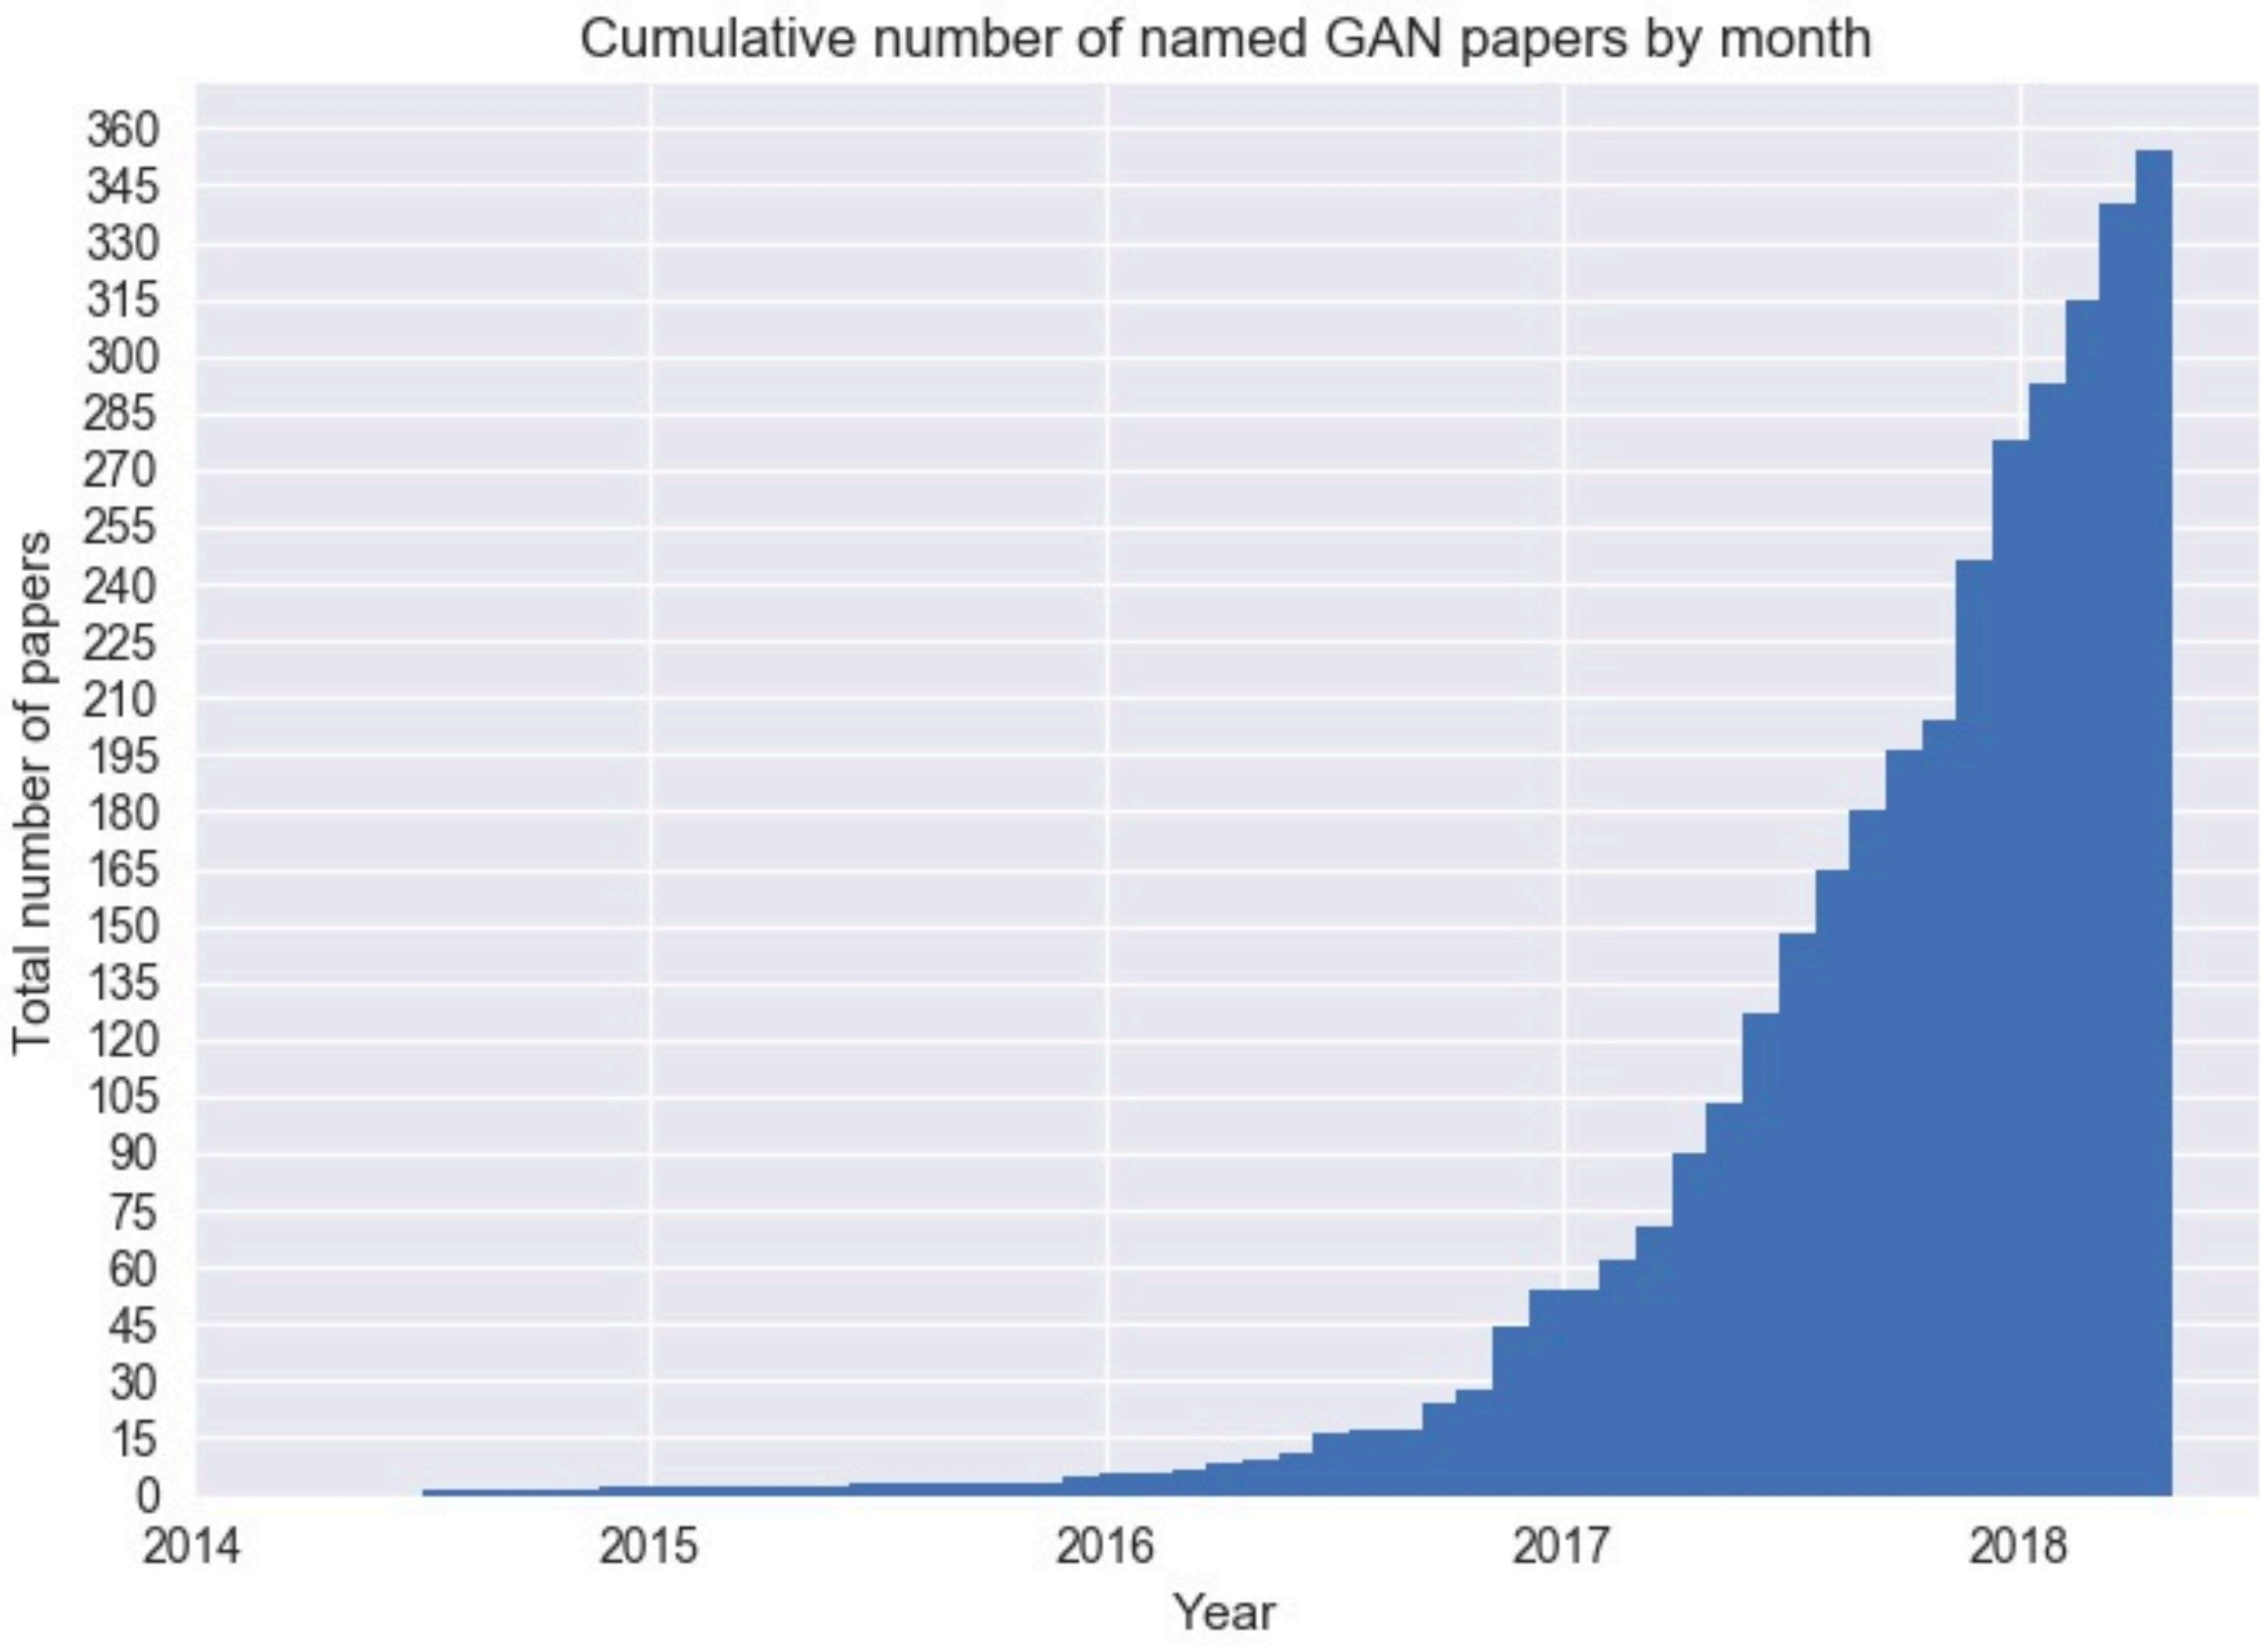

https://github.com/hindupuravinash/the-gan-zoo

(Goodfellow 2018)

Questions