MedGAN ID-CGAN CoGAN LR-GAN CGAN IcGAN
b-GAN LS-GAN AffGAN LAPGAN DiscoGAN MPM-GAN AdaGAN
LSGAN InfoGAN CatGAN AMGAN iGAN IAN

Generative Adversarial Networks

Ian Goodfellow, Staff Research Scientist, Google Brain

NVIDIA Distinguished Lecture Series in Machine Learning

USC, Los Angeles 2017-09-05

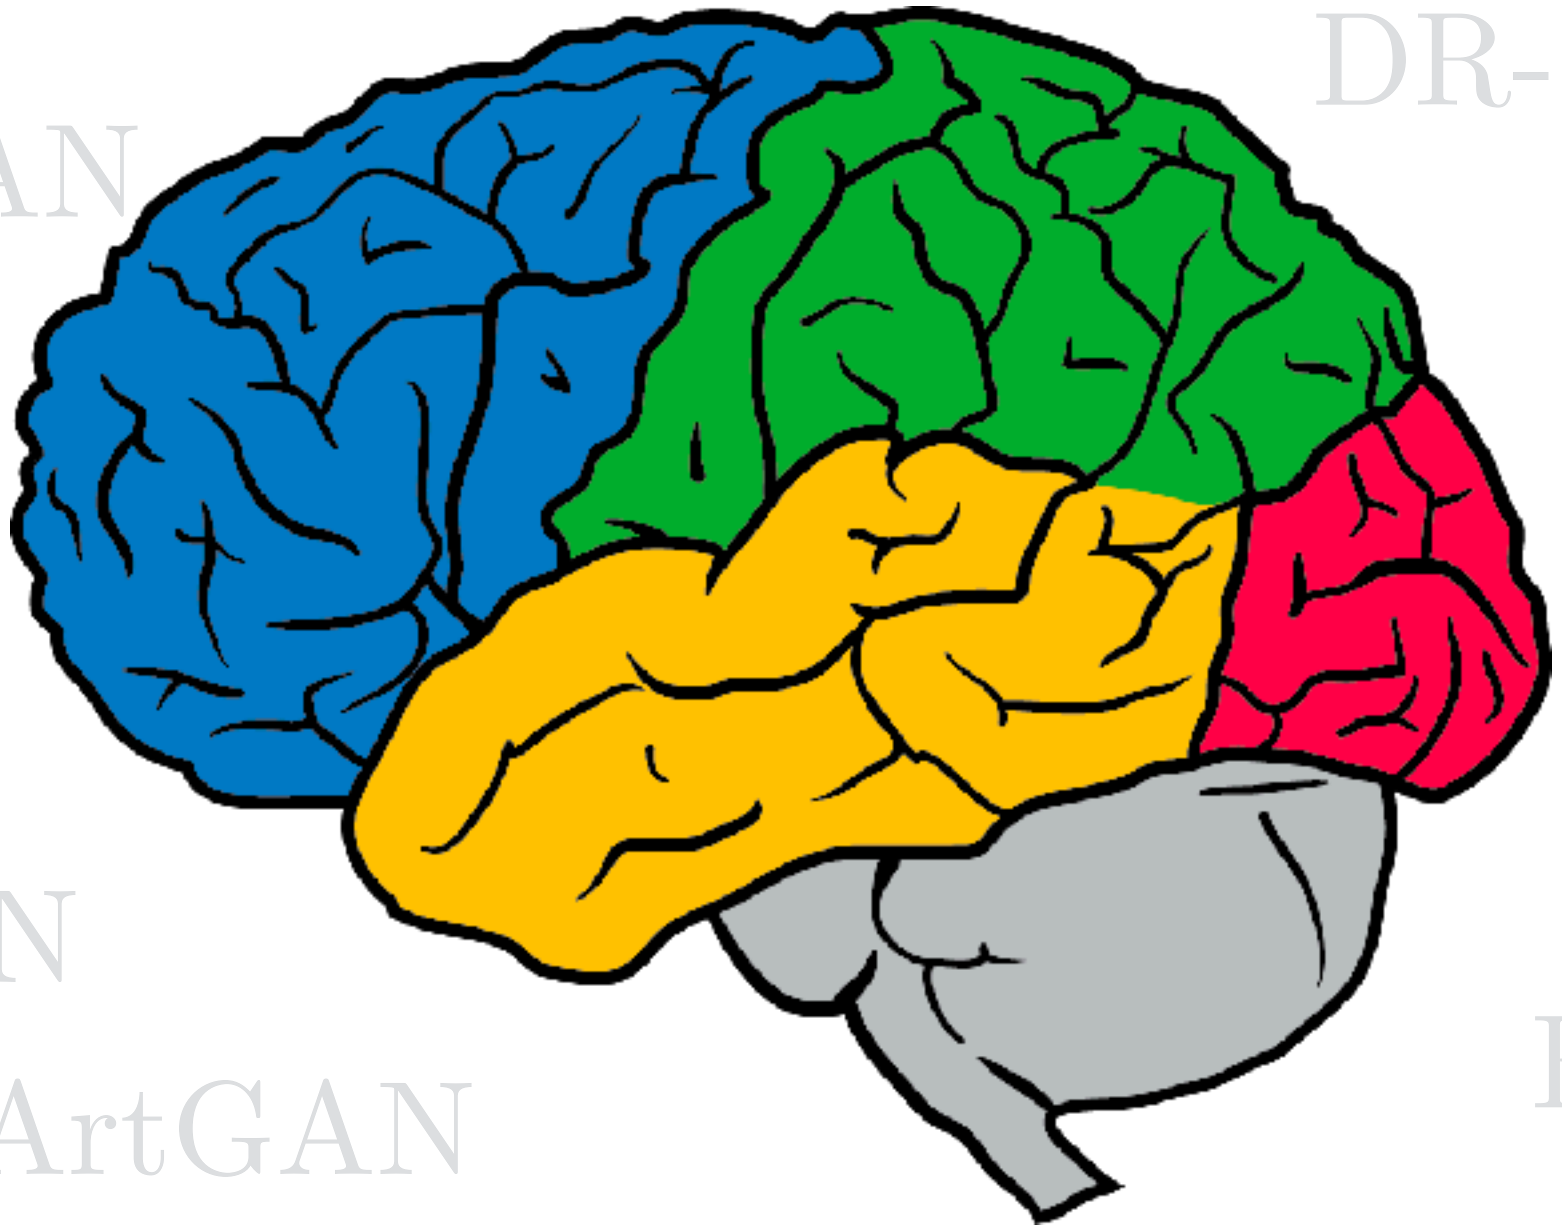

McGAN MIX+GAN MGAN FF-GAN BS-GAN GoGAN C-VAE-GAN C-RNN-GAN DR-GAN DCGAN CCGAN AC-GAN MAGAN 3D-GAN BiGAN GAWWN DualGAN CycleGAN Bayesian GAN GP-GAN AnoGAN EBGAN DTN Context-RNN-GAN MAD-GAN ALI BEGAN AL-CGAN MARTA-GAN f-GAN ArtGAN MalGAN

Generative Modeling

- Density estimation

- Sample generation

Training examples　　Model samples

(Goodfellow 2017)

Maximum Likelihood

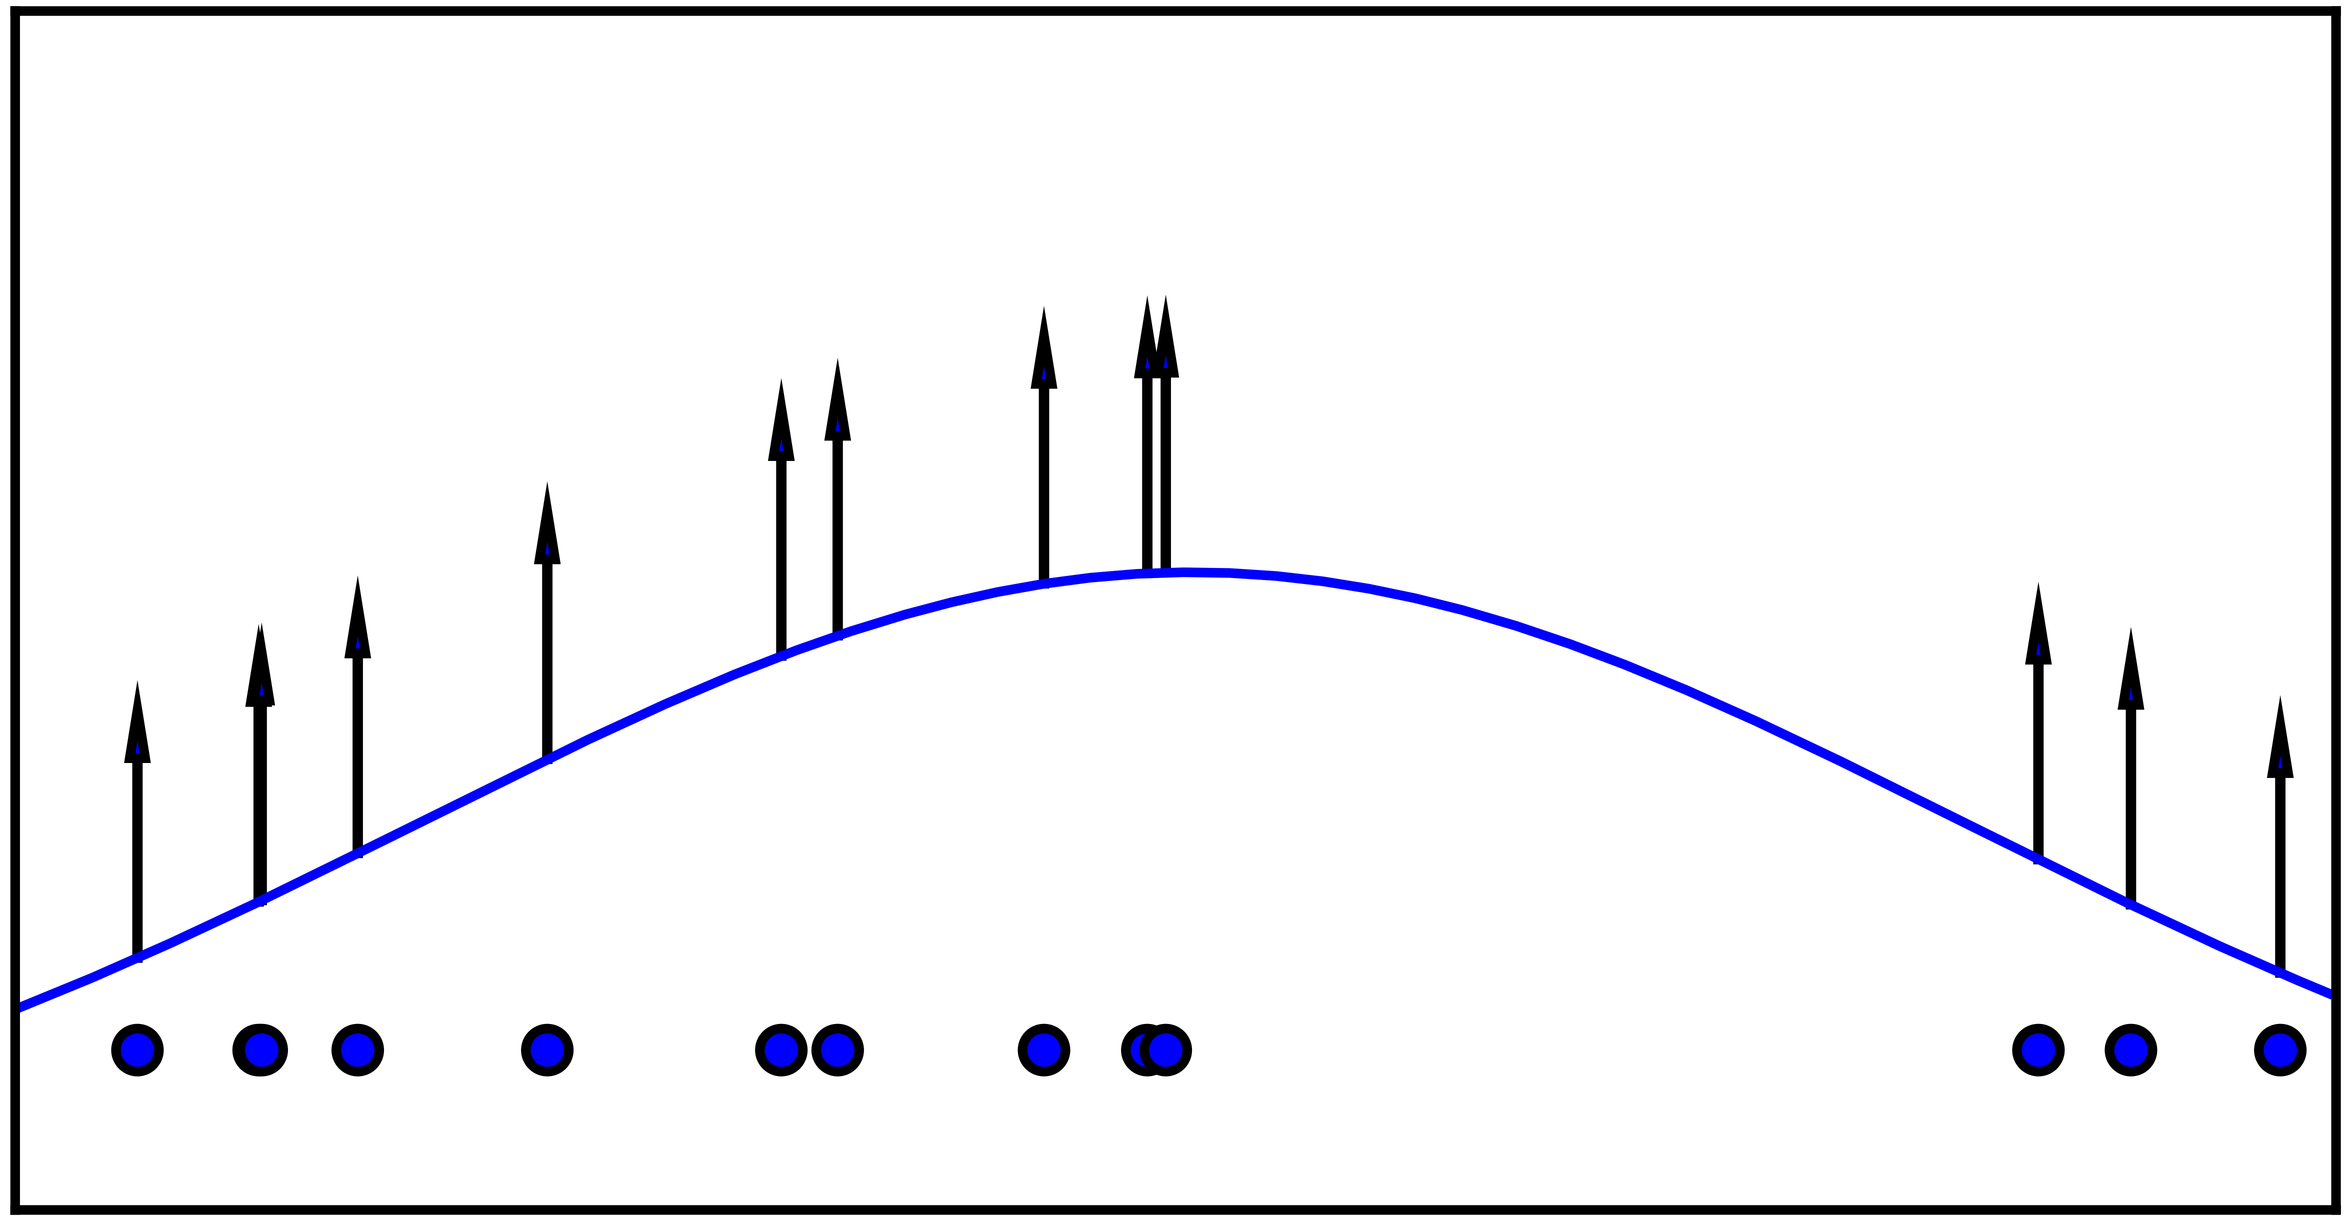

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} \mathbb{E}_{x \sim p_{\text{data}}} \log p_{\text{model}}(\boldsymbol{x} \mid \boldsymbol{\theta})$$

Adversarial Nets Framework

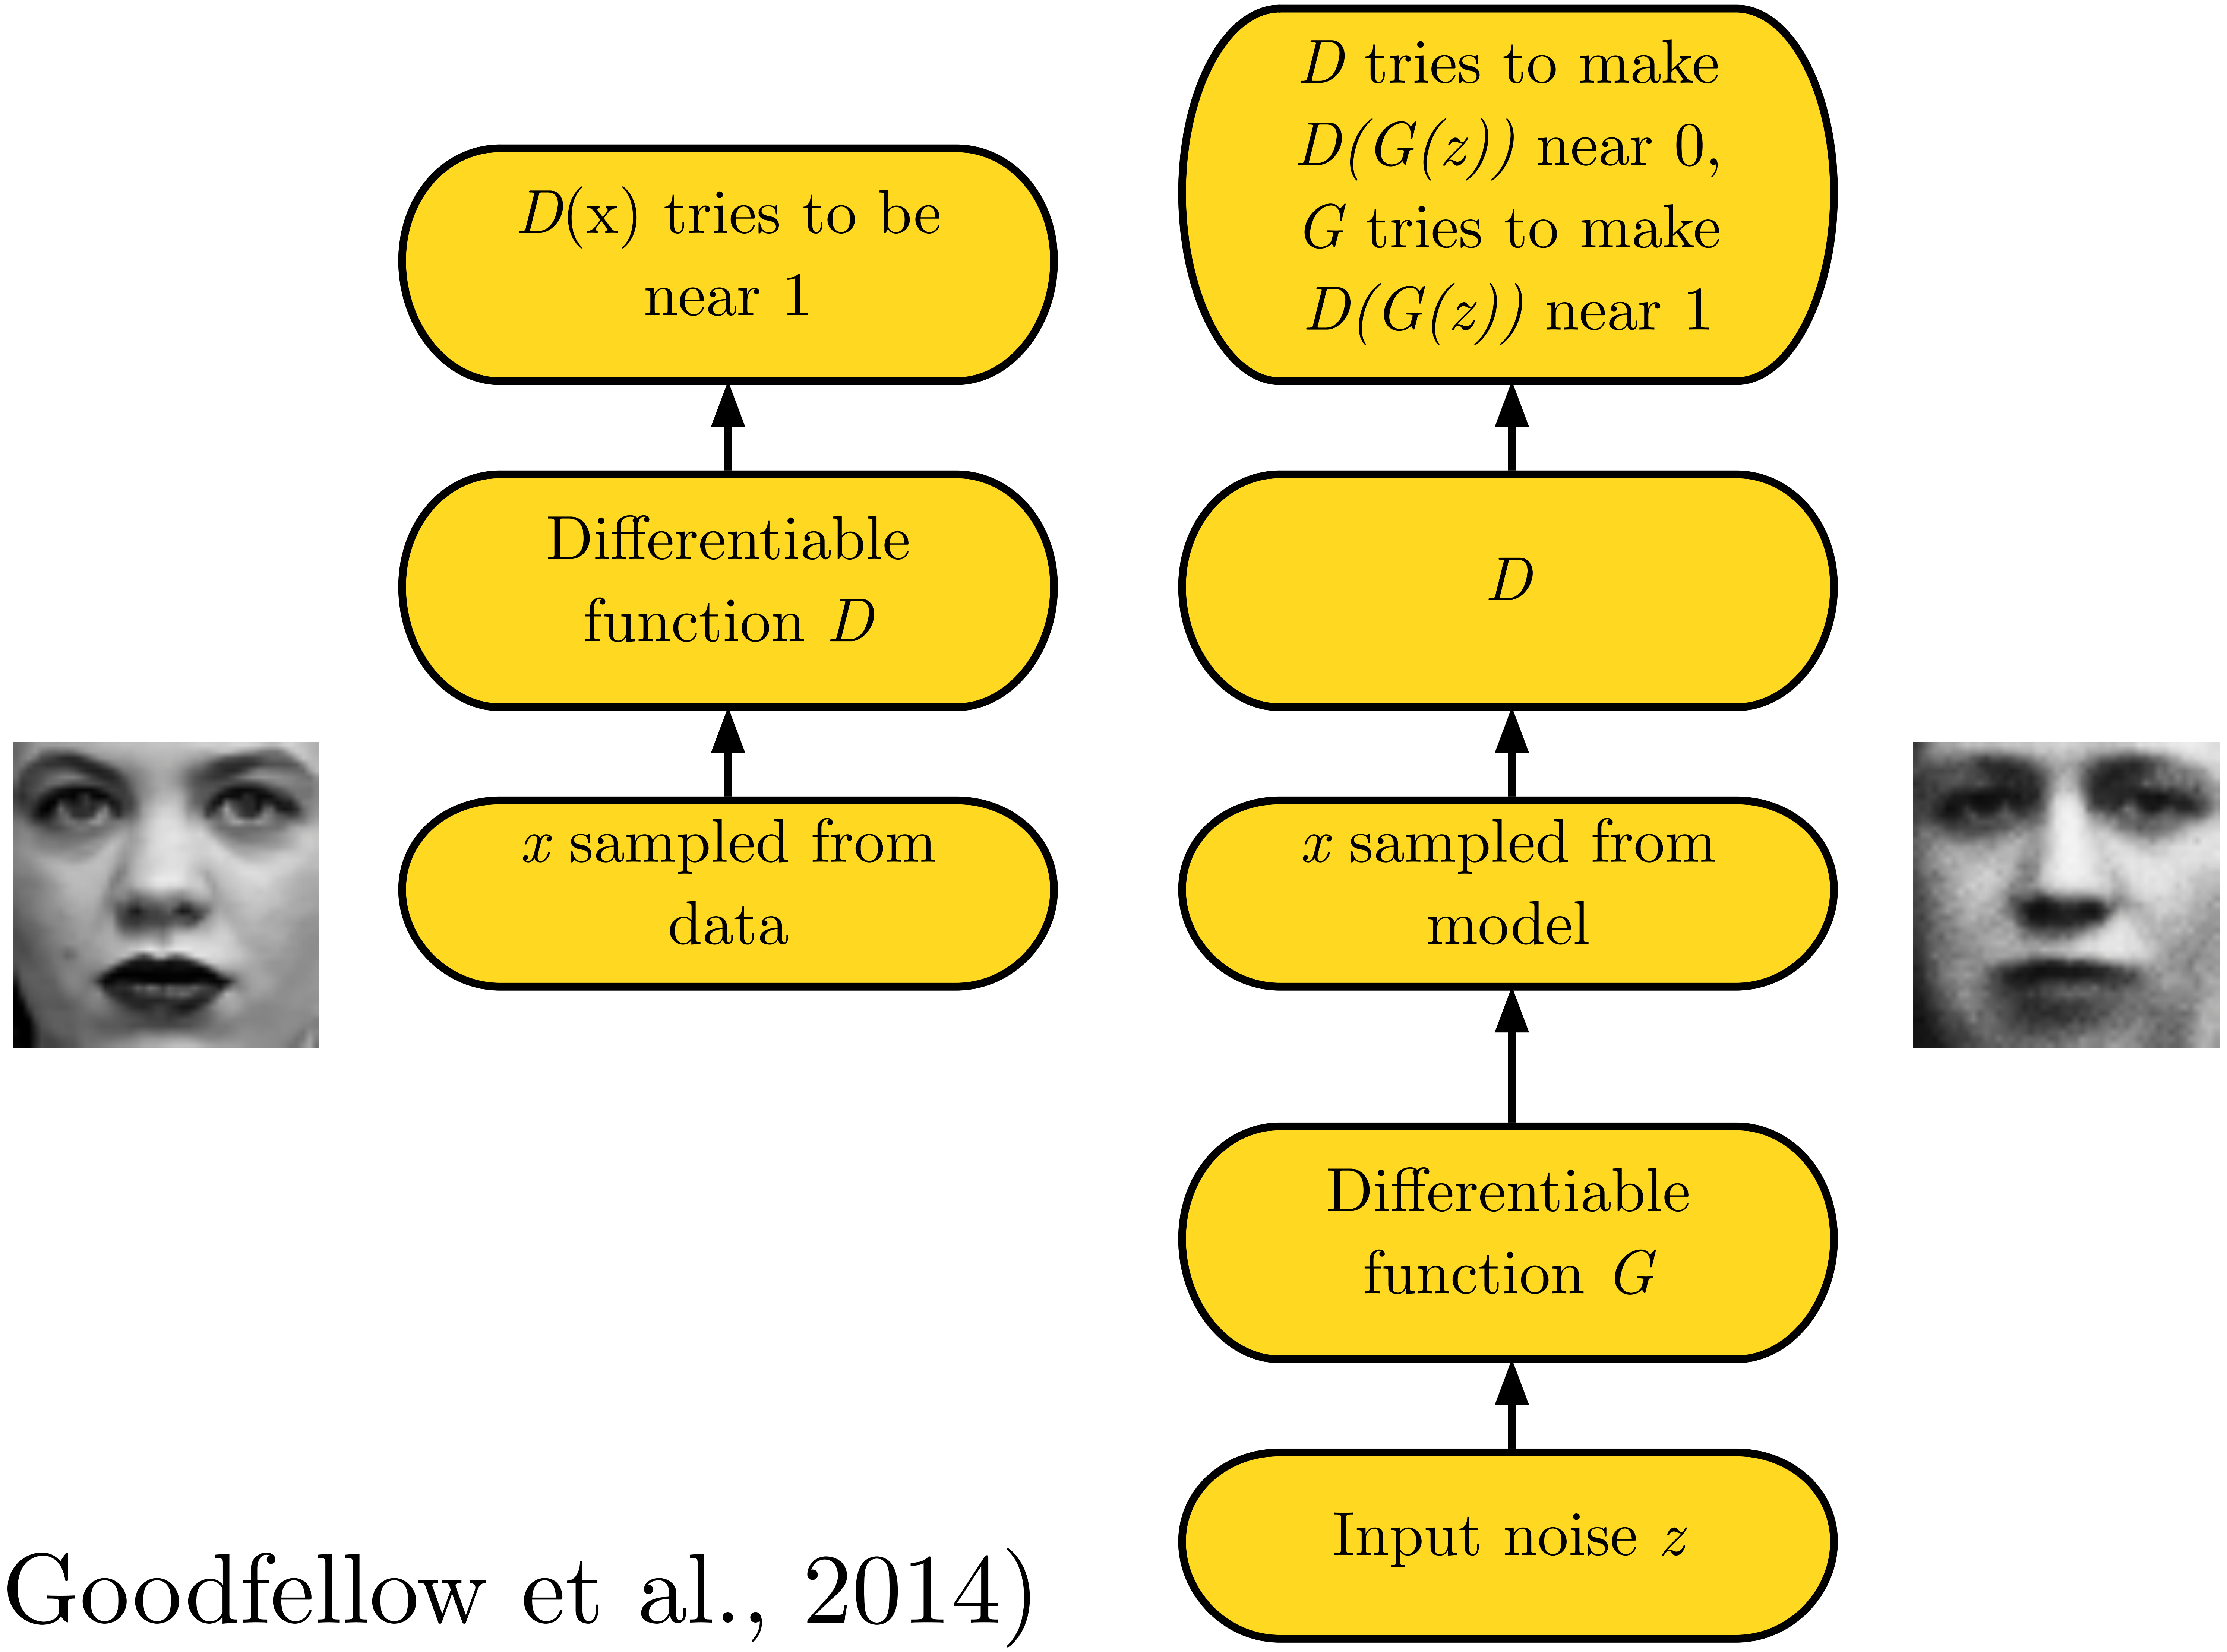

(Goodfellow et al., 2014)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

(Goodfellow 2017)

TEACHING AID

Apple's first research paper tries to solve a problem facing every company working on AI

GANs for simulated training data

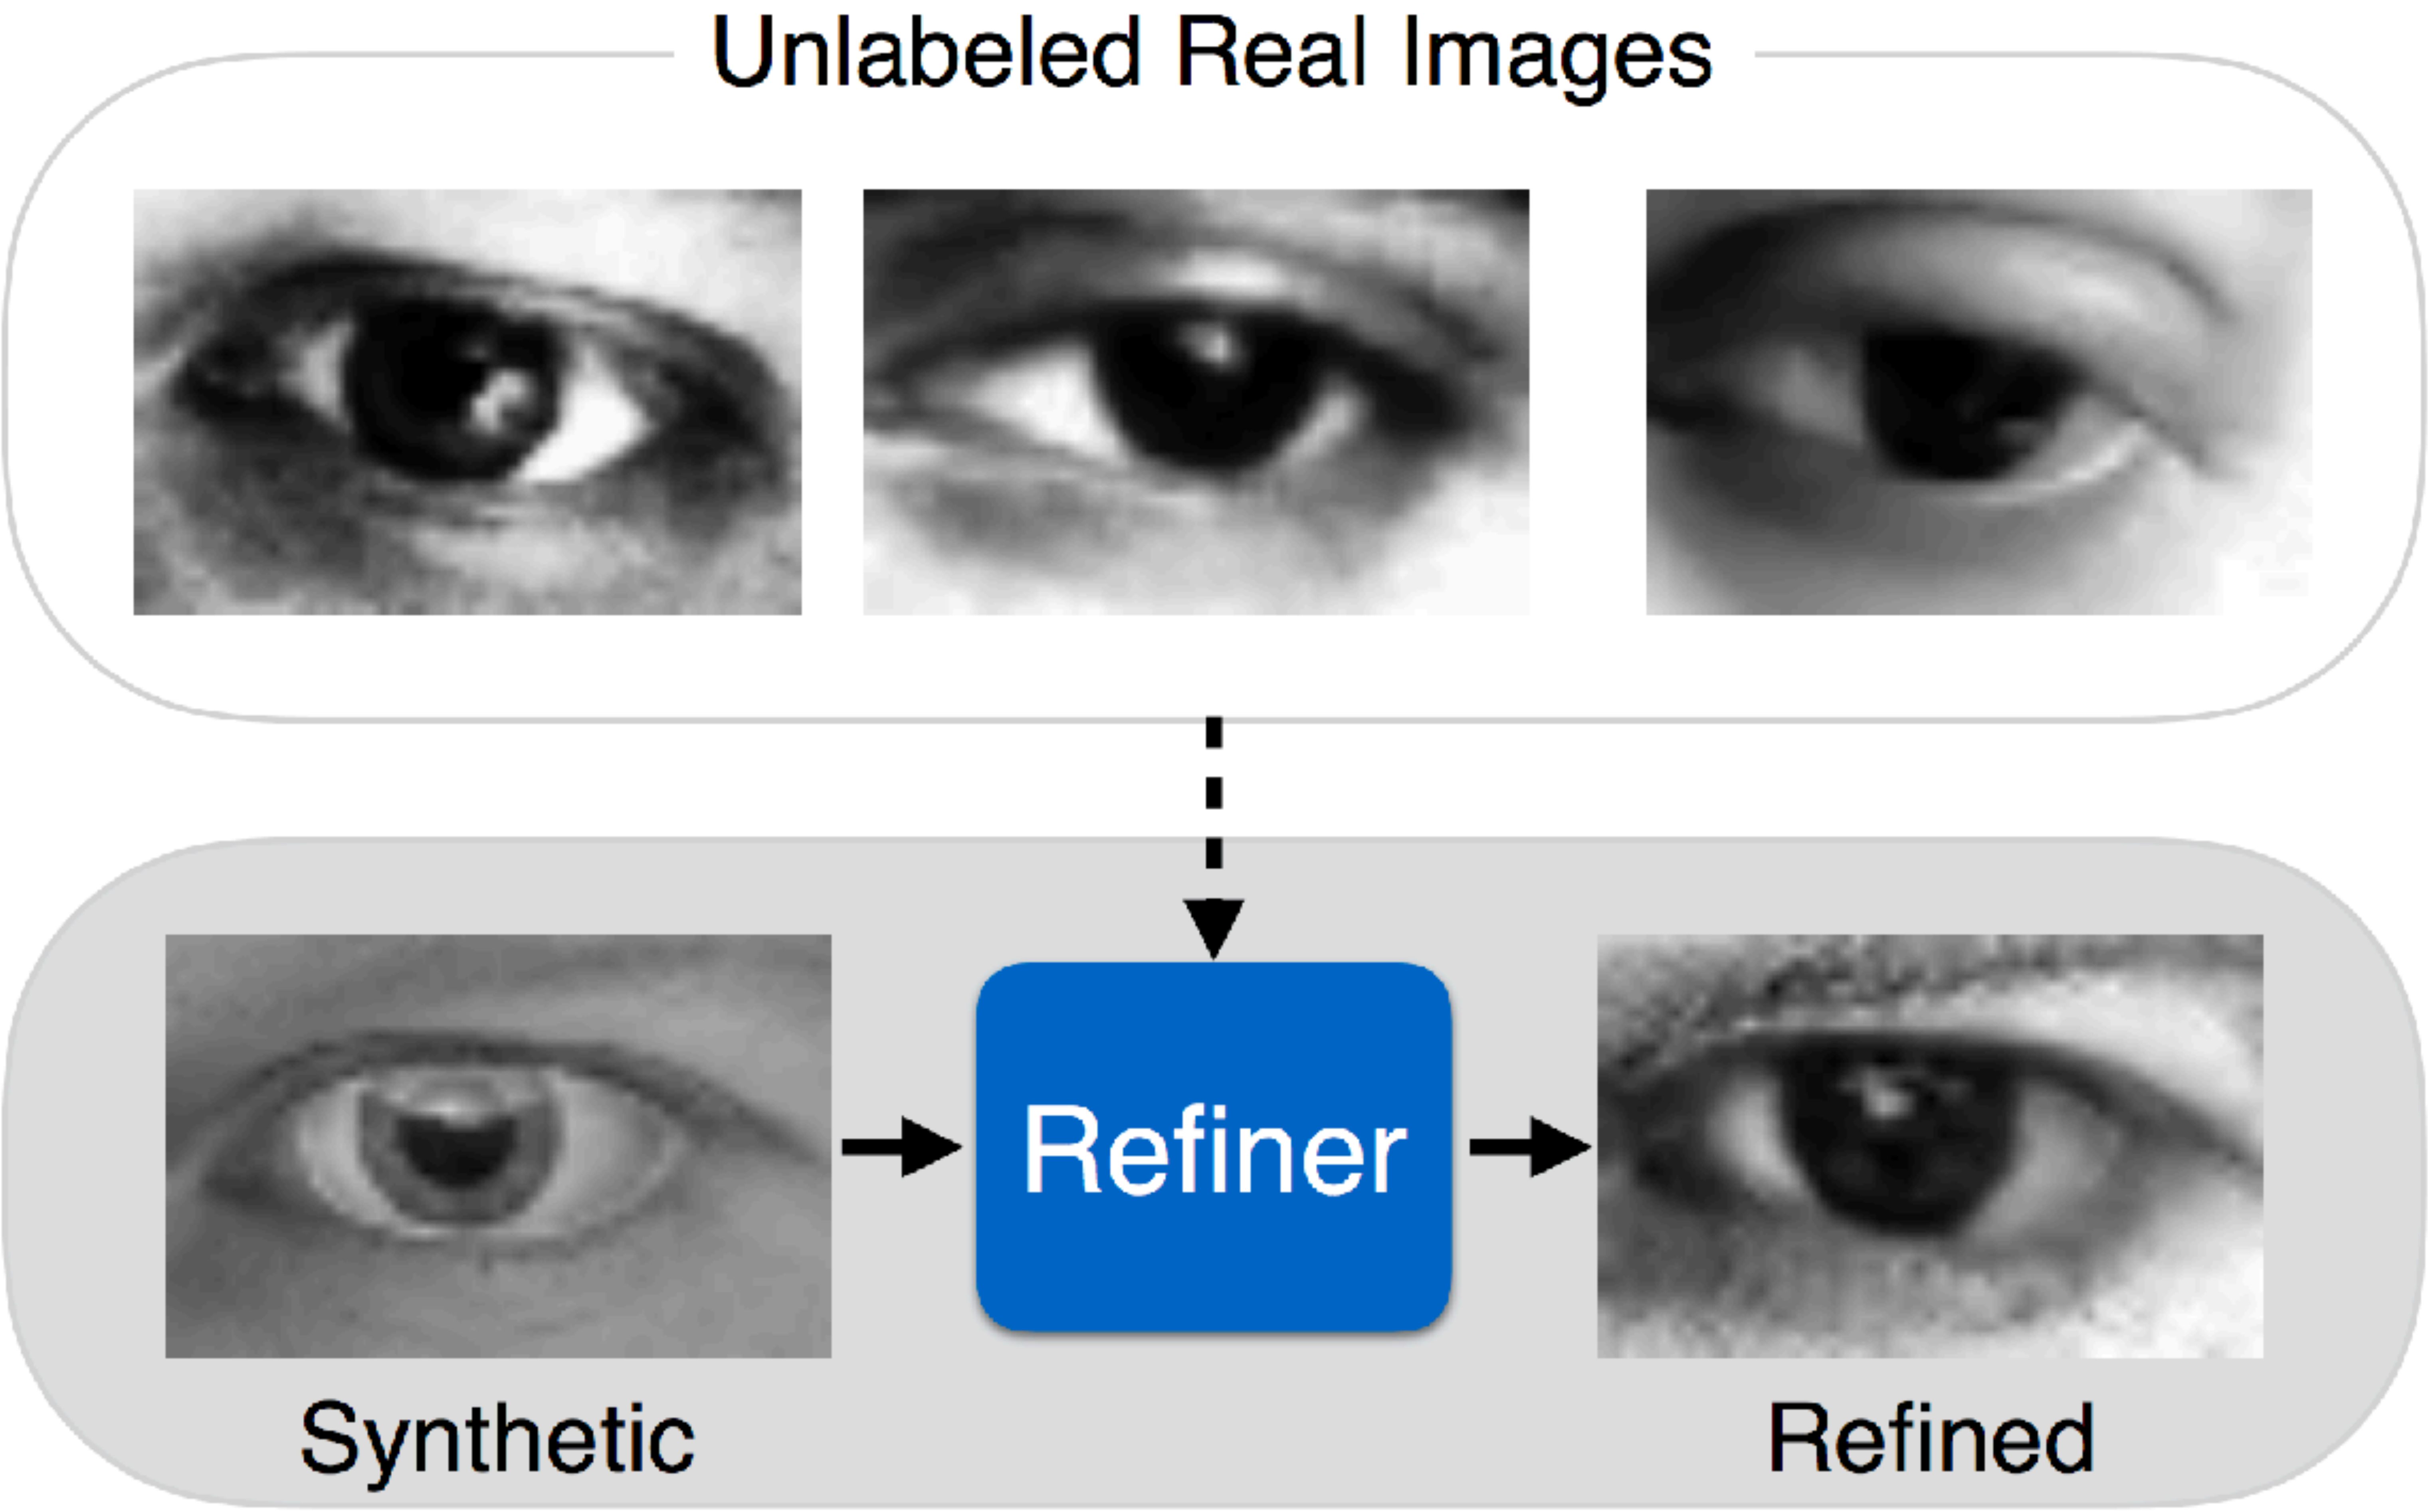

(Shrivastava et al., 2016)

GANs for domain adaptation

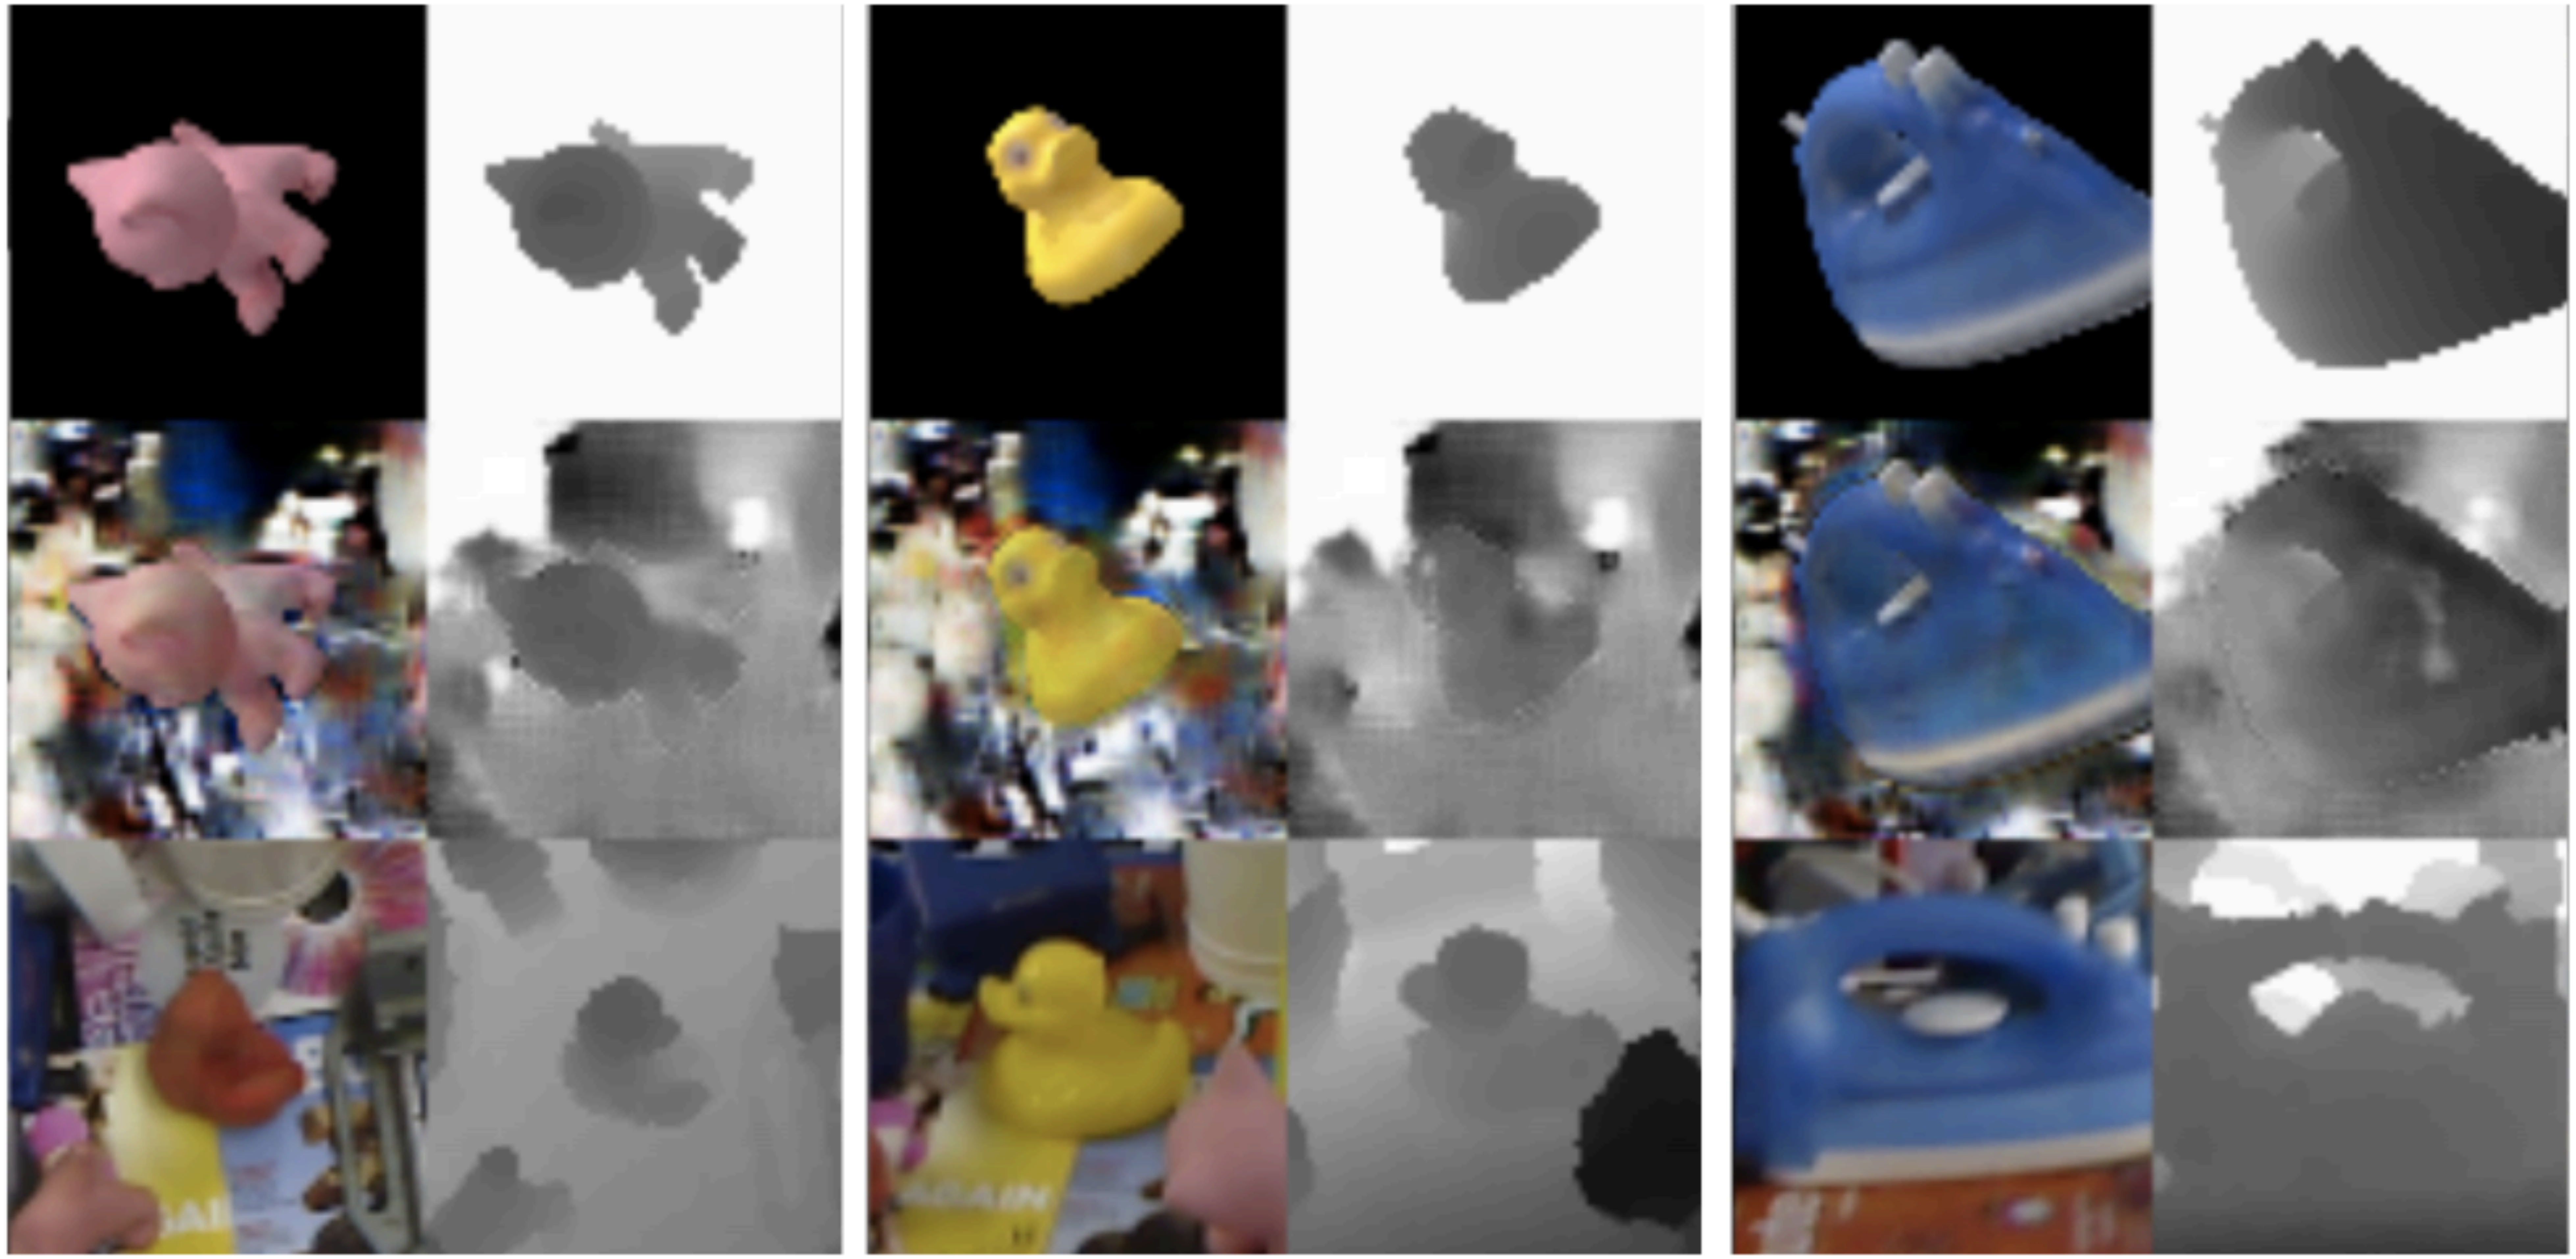

(Bousmalis et al., 2016)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

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
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

Supervised Discriminator

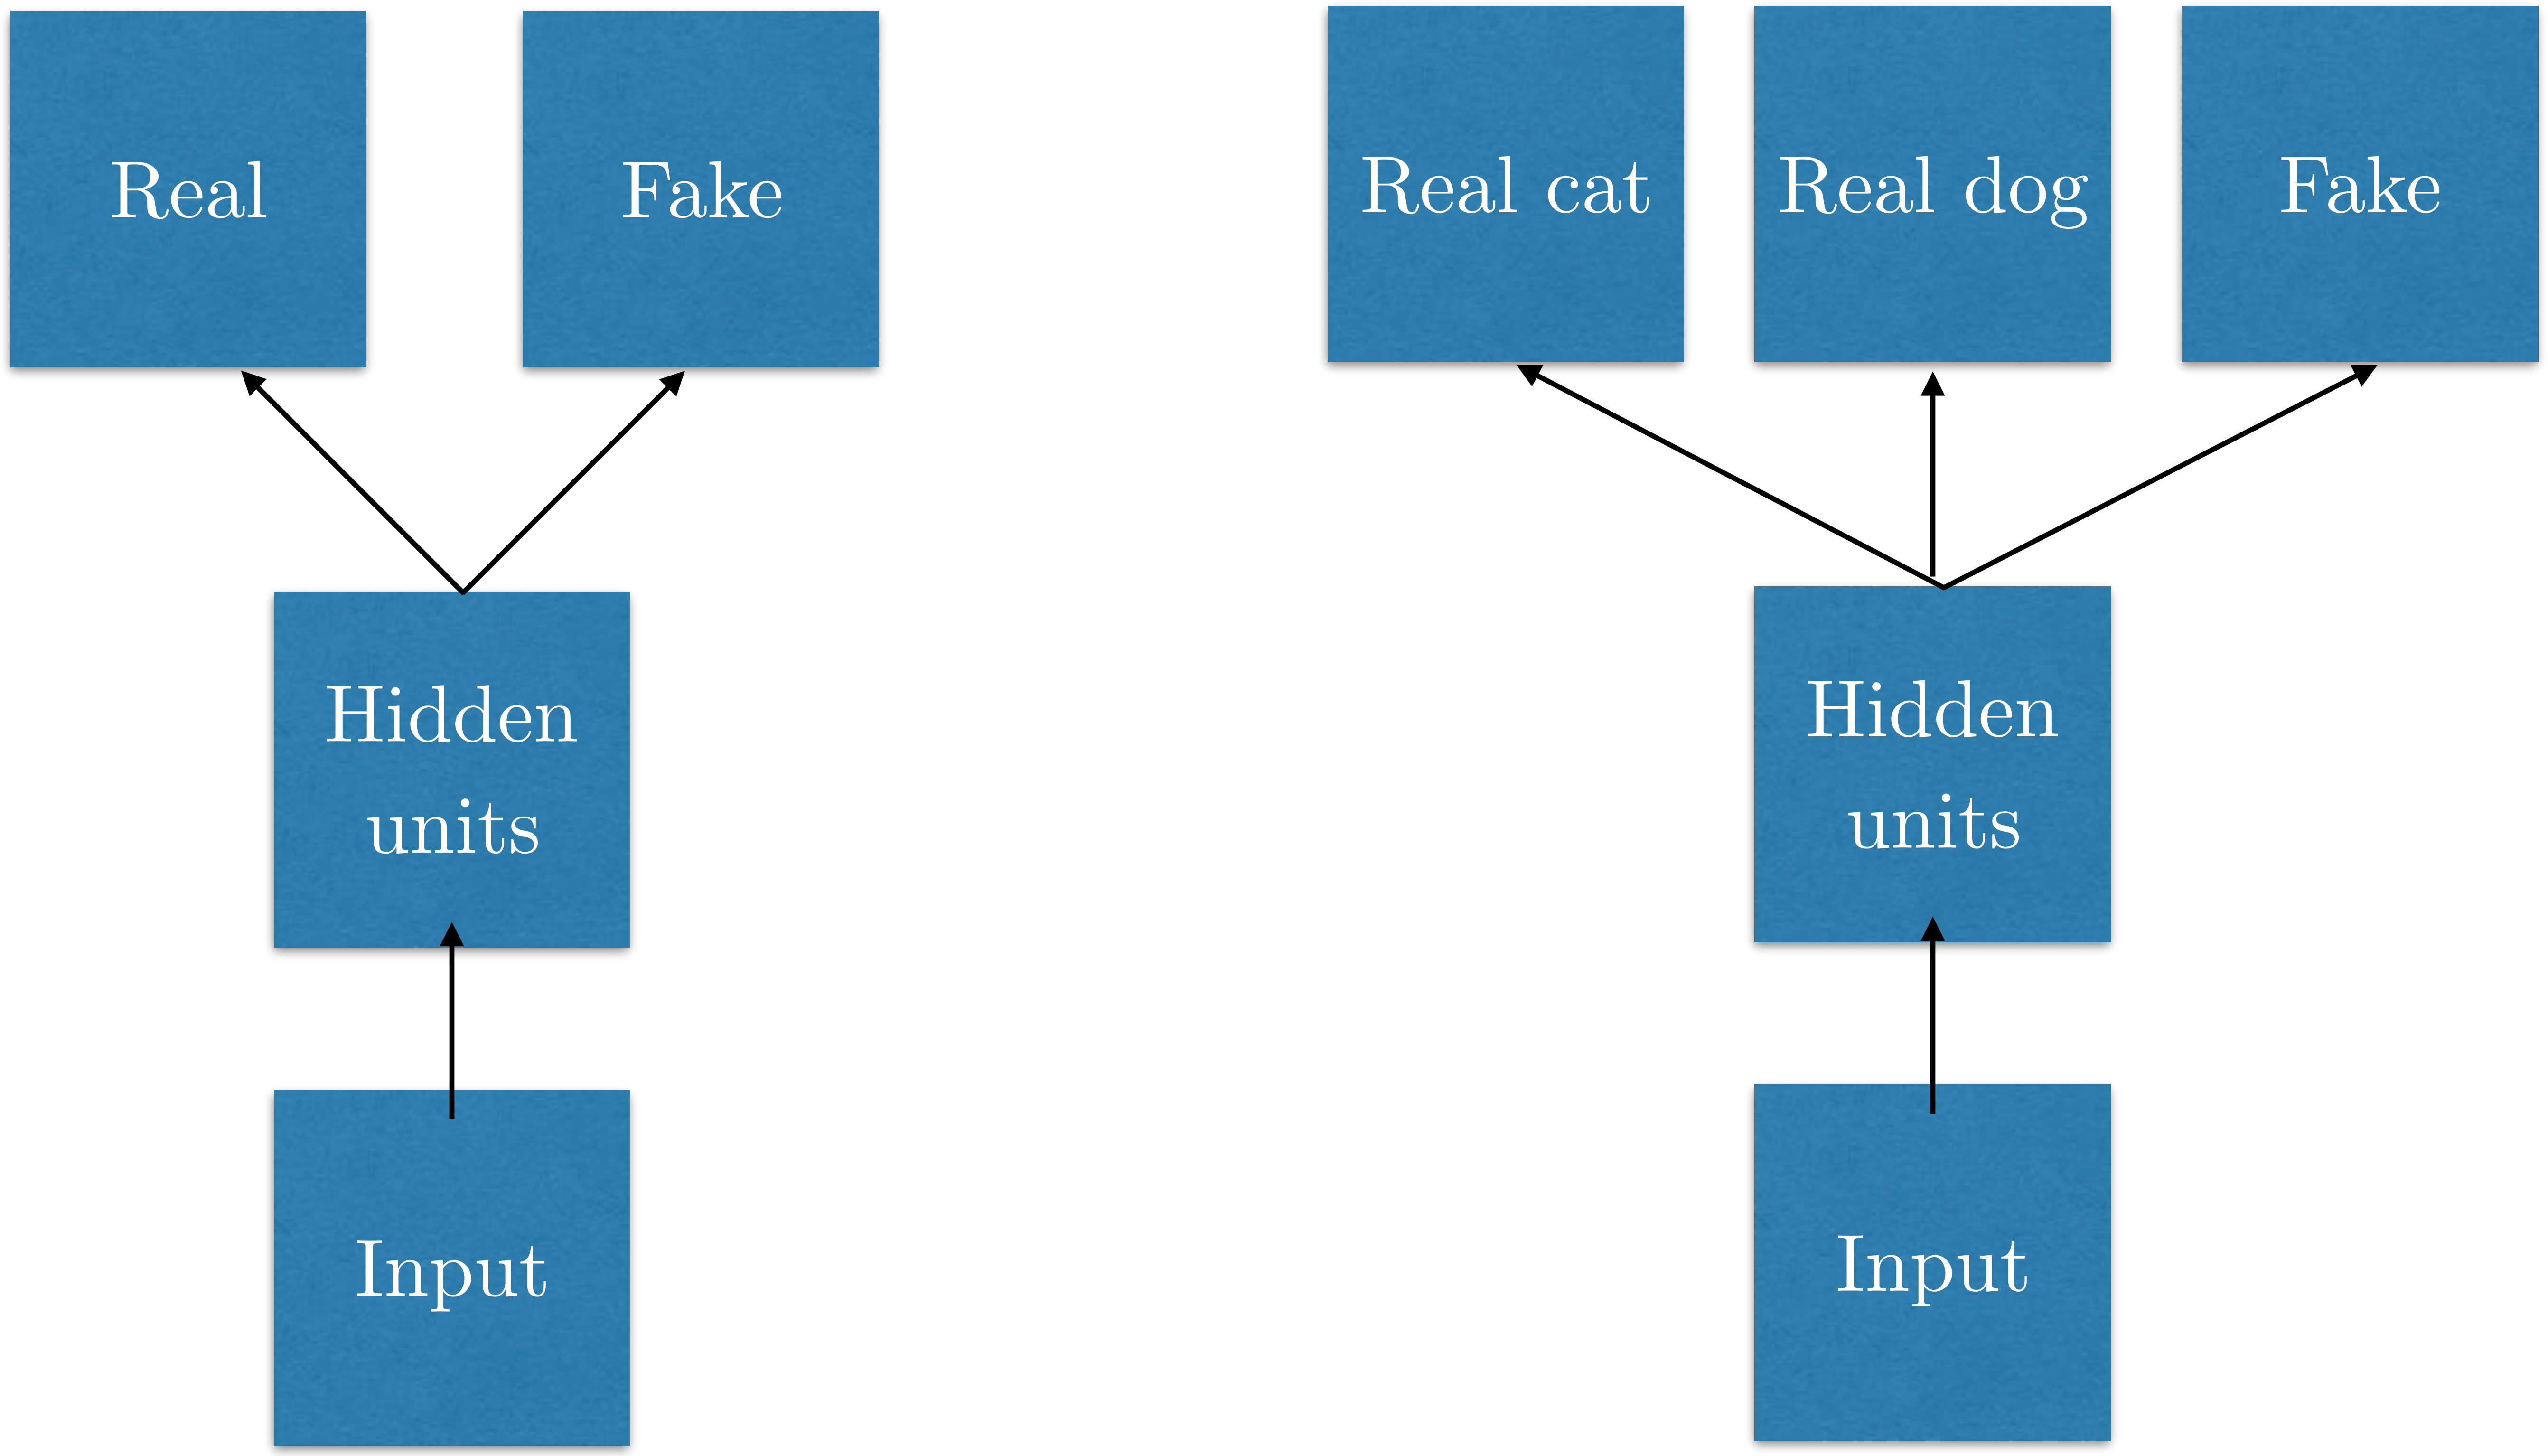

(Odena 2016, Salimans et al 2016)

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
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

Next Video Frame Prediction

Ground Truth

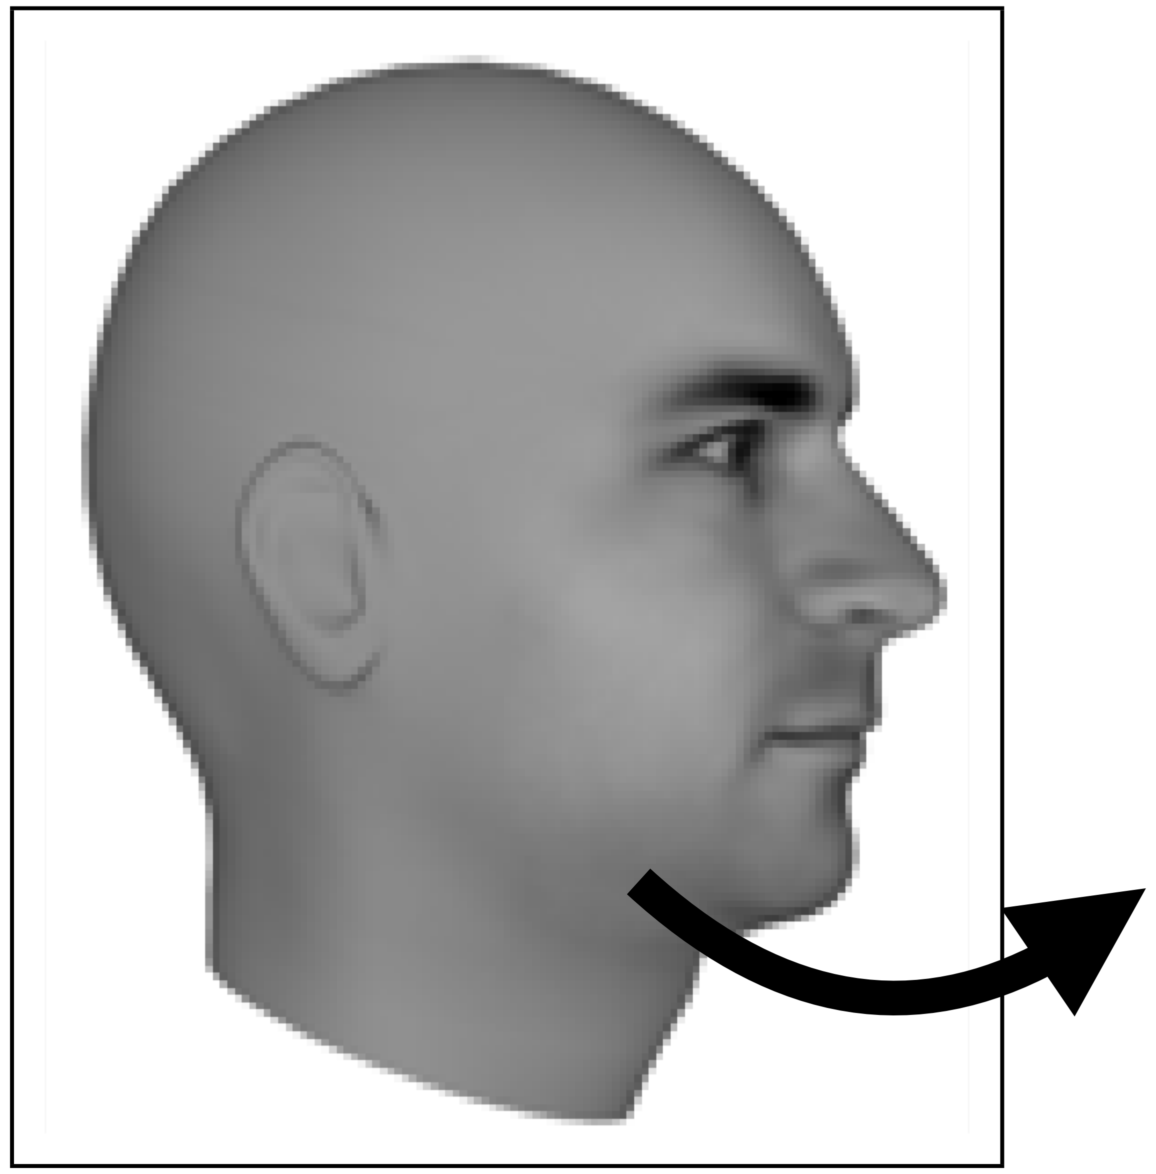

What happens next?

(Lotter et al 2016)

Next Video Frame Prediction

Ground Truth MSE Adversarial

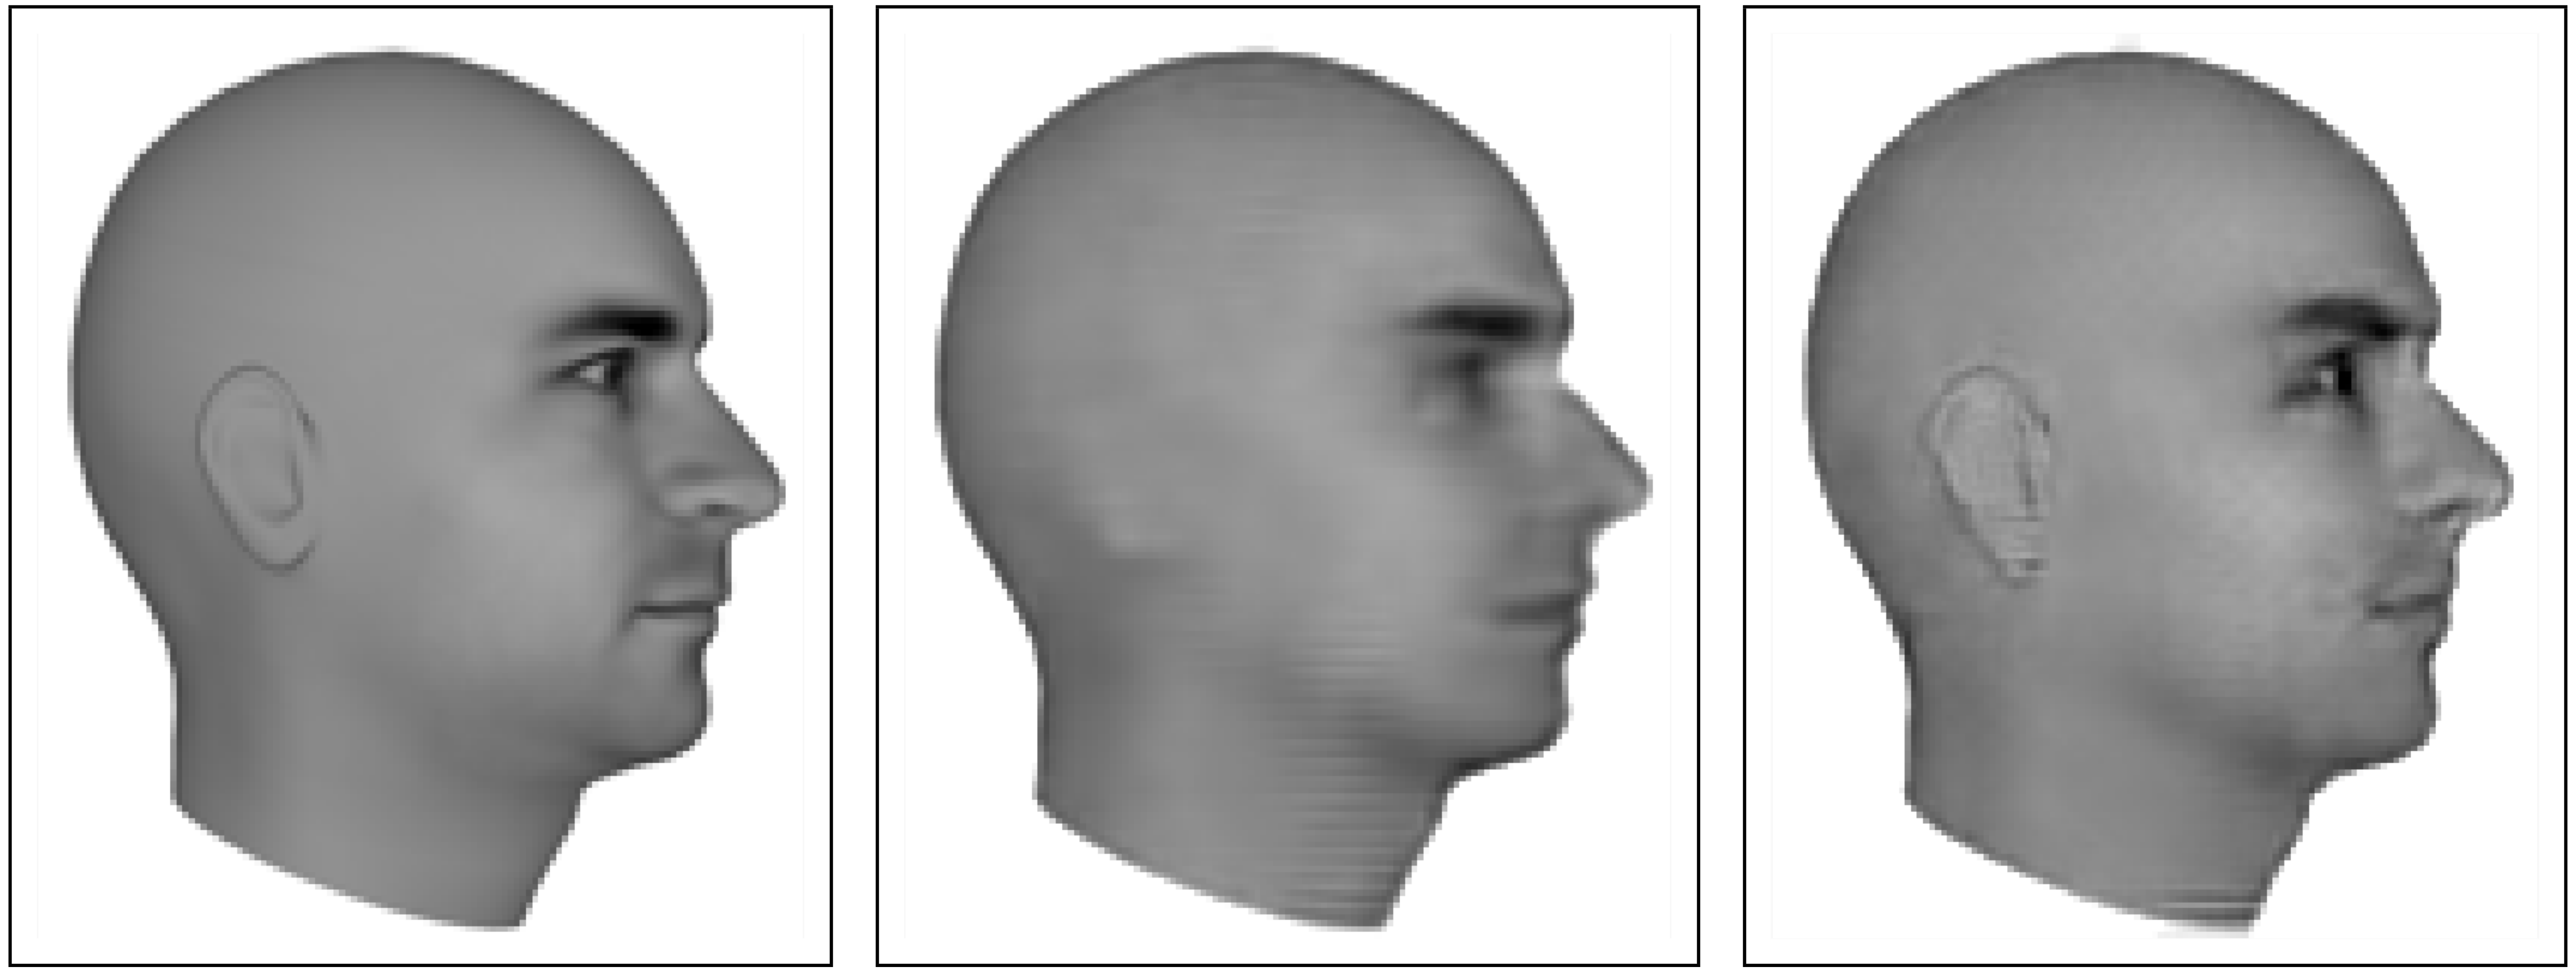

(Lotter et al 2016)

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
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

iGAN

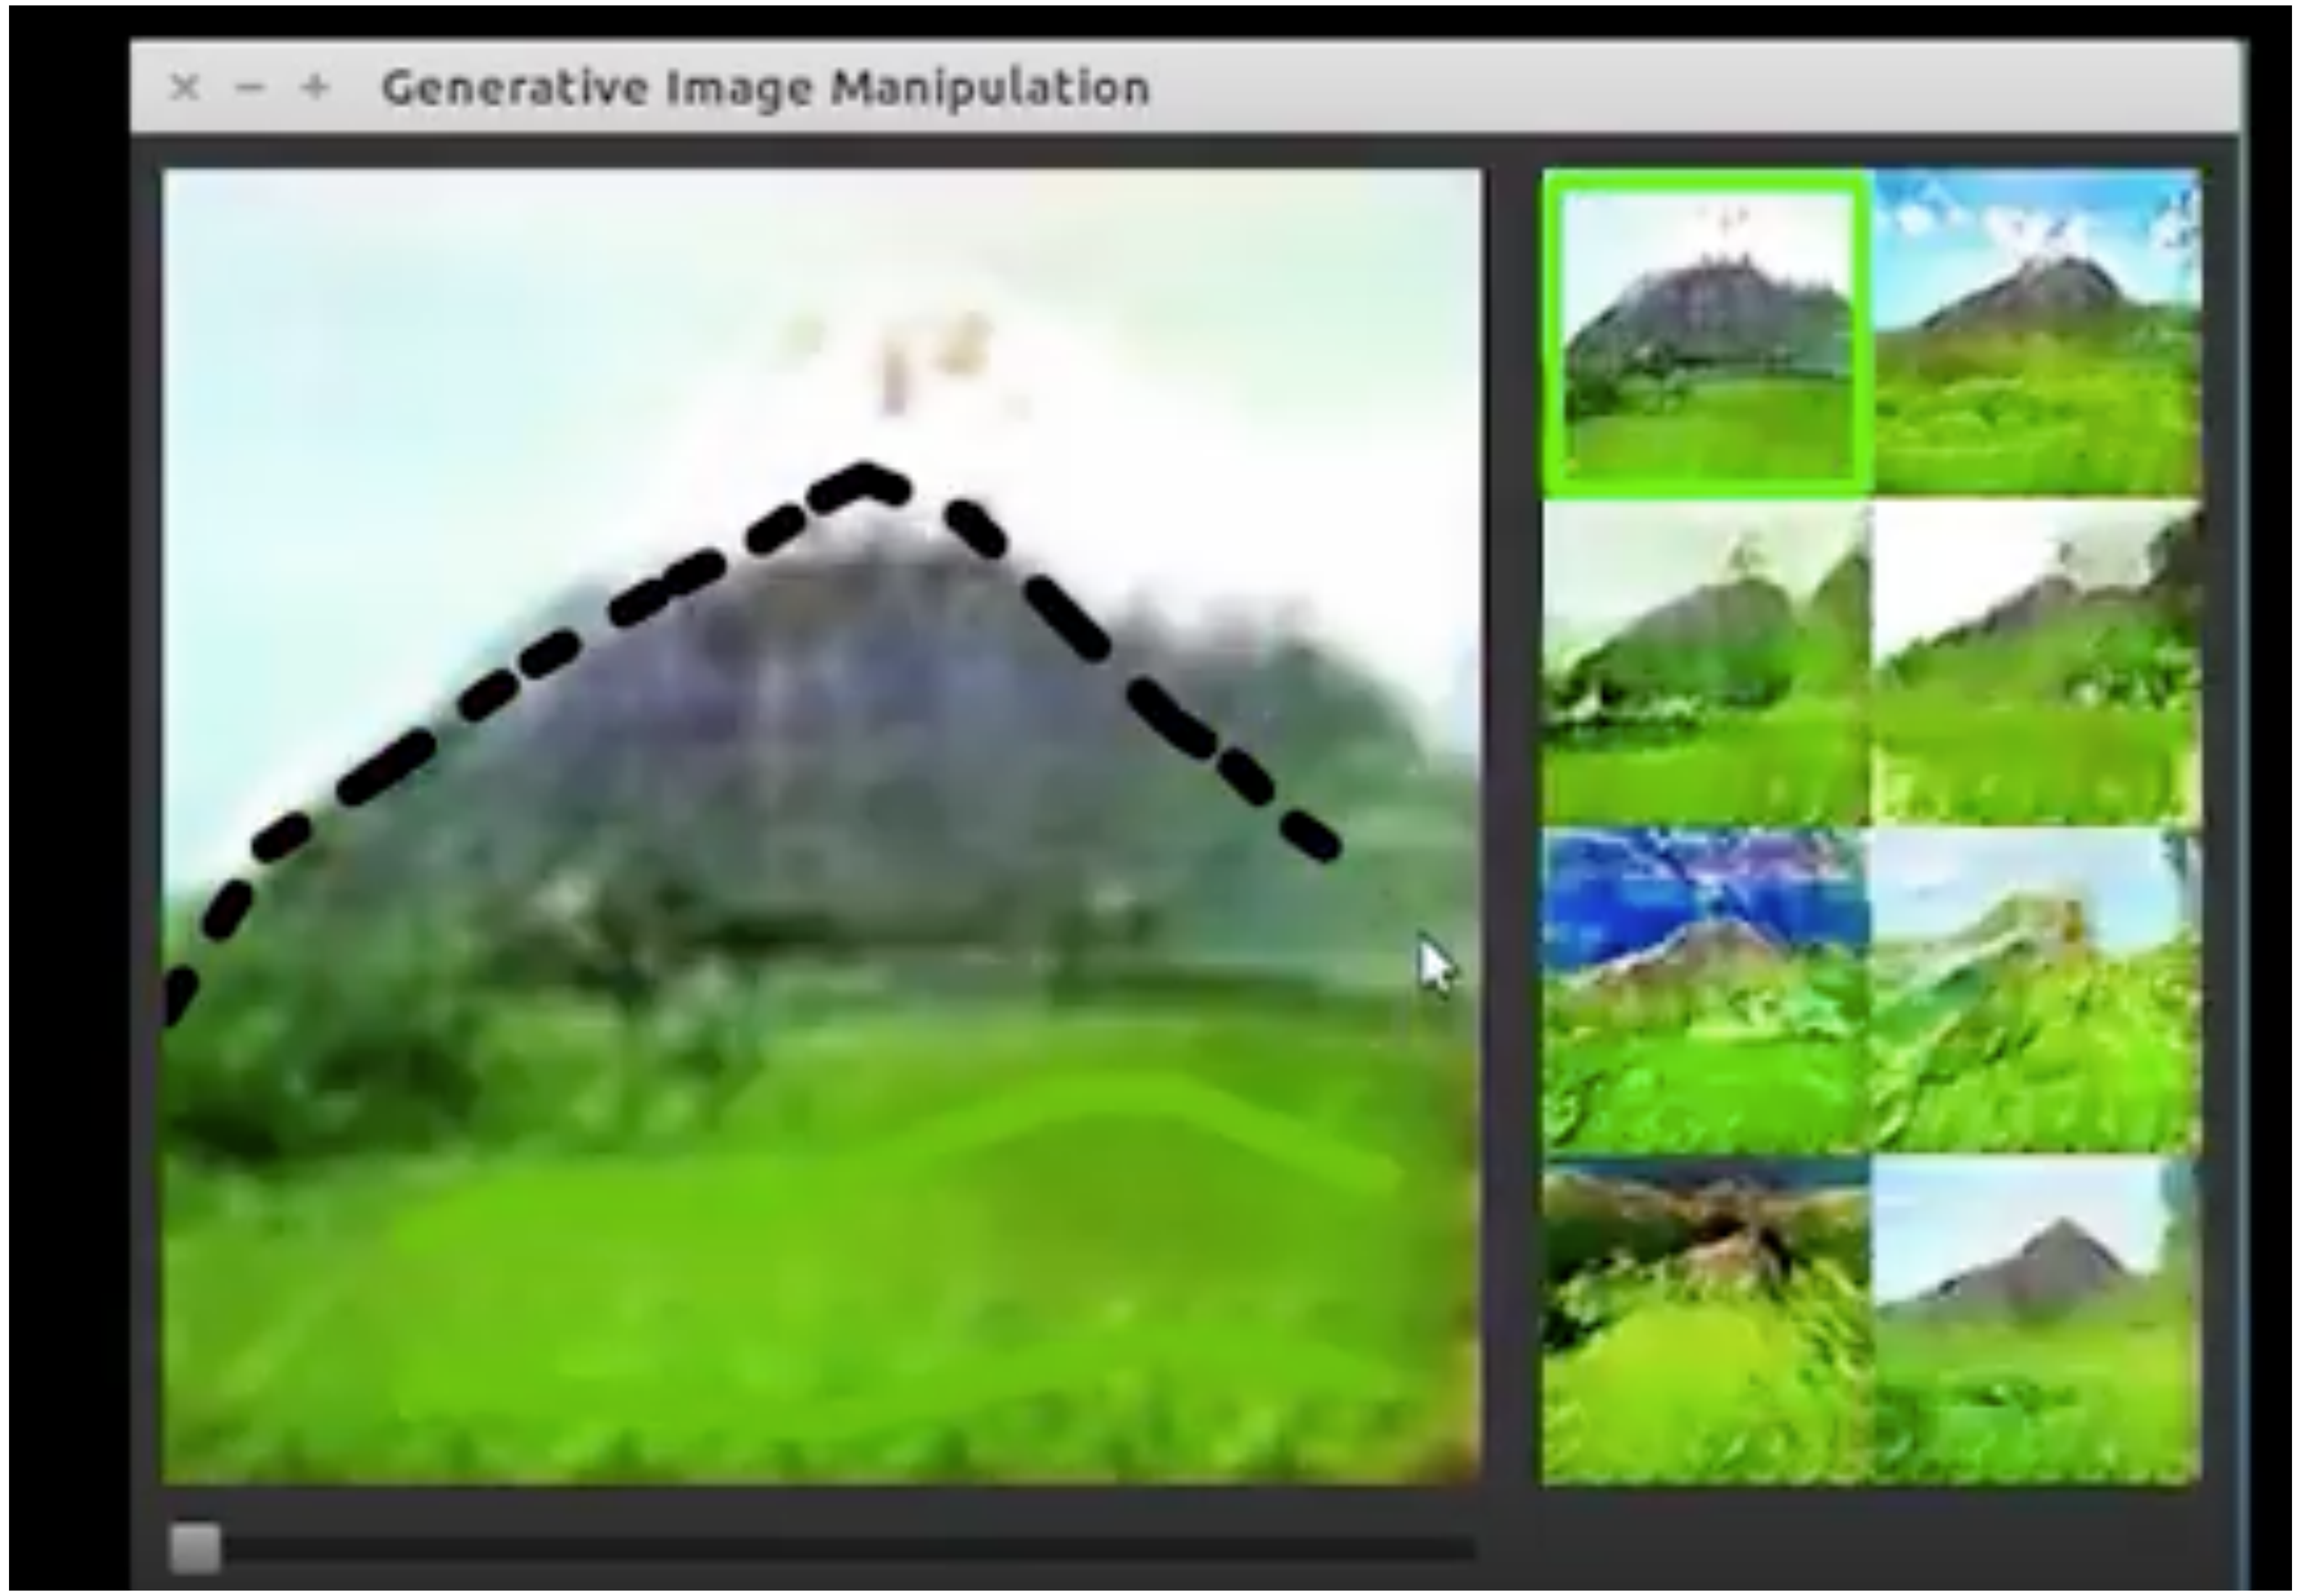

youtube

(Zhu et al., 2016)

Introspective Adversarial Networks

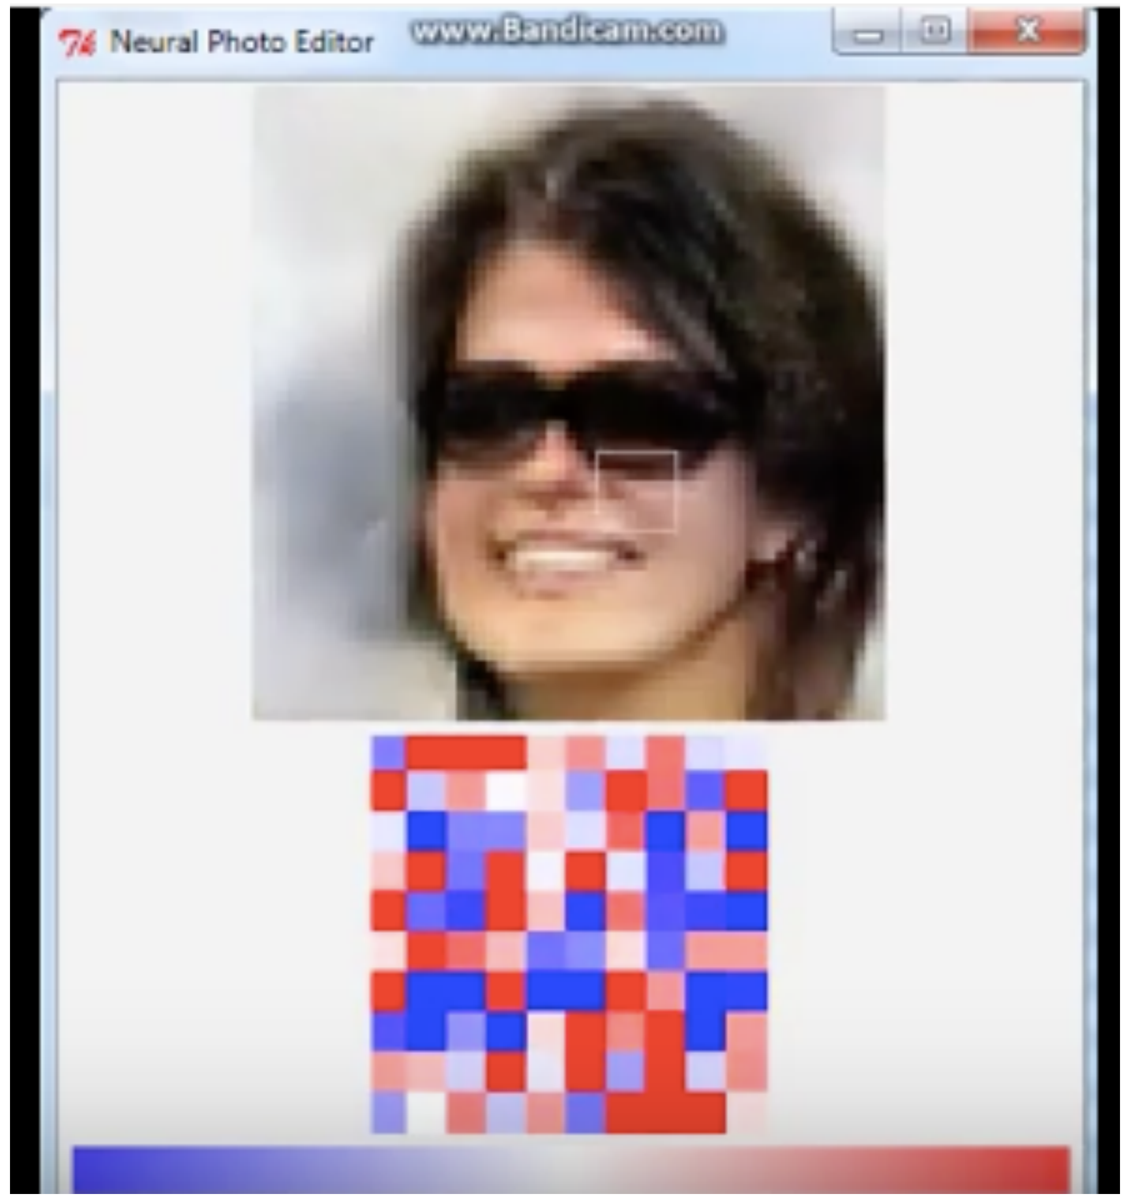

youtube

(Brock 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

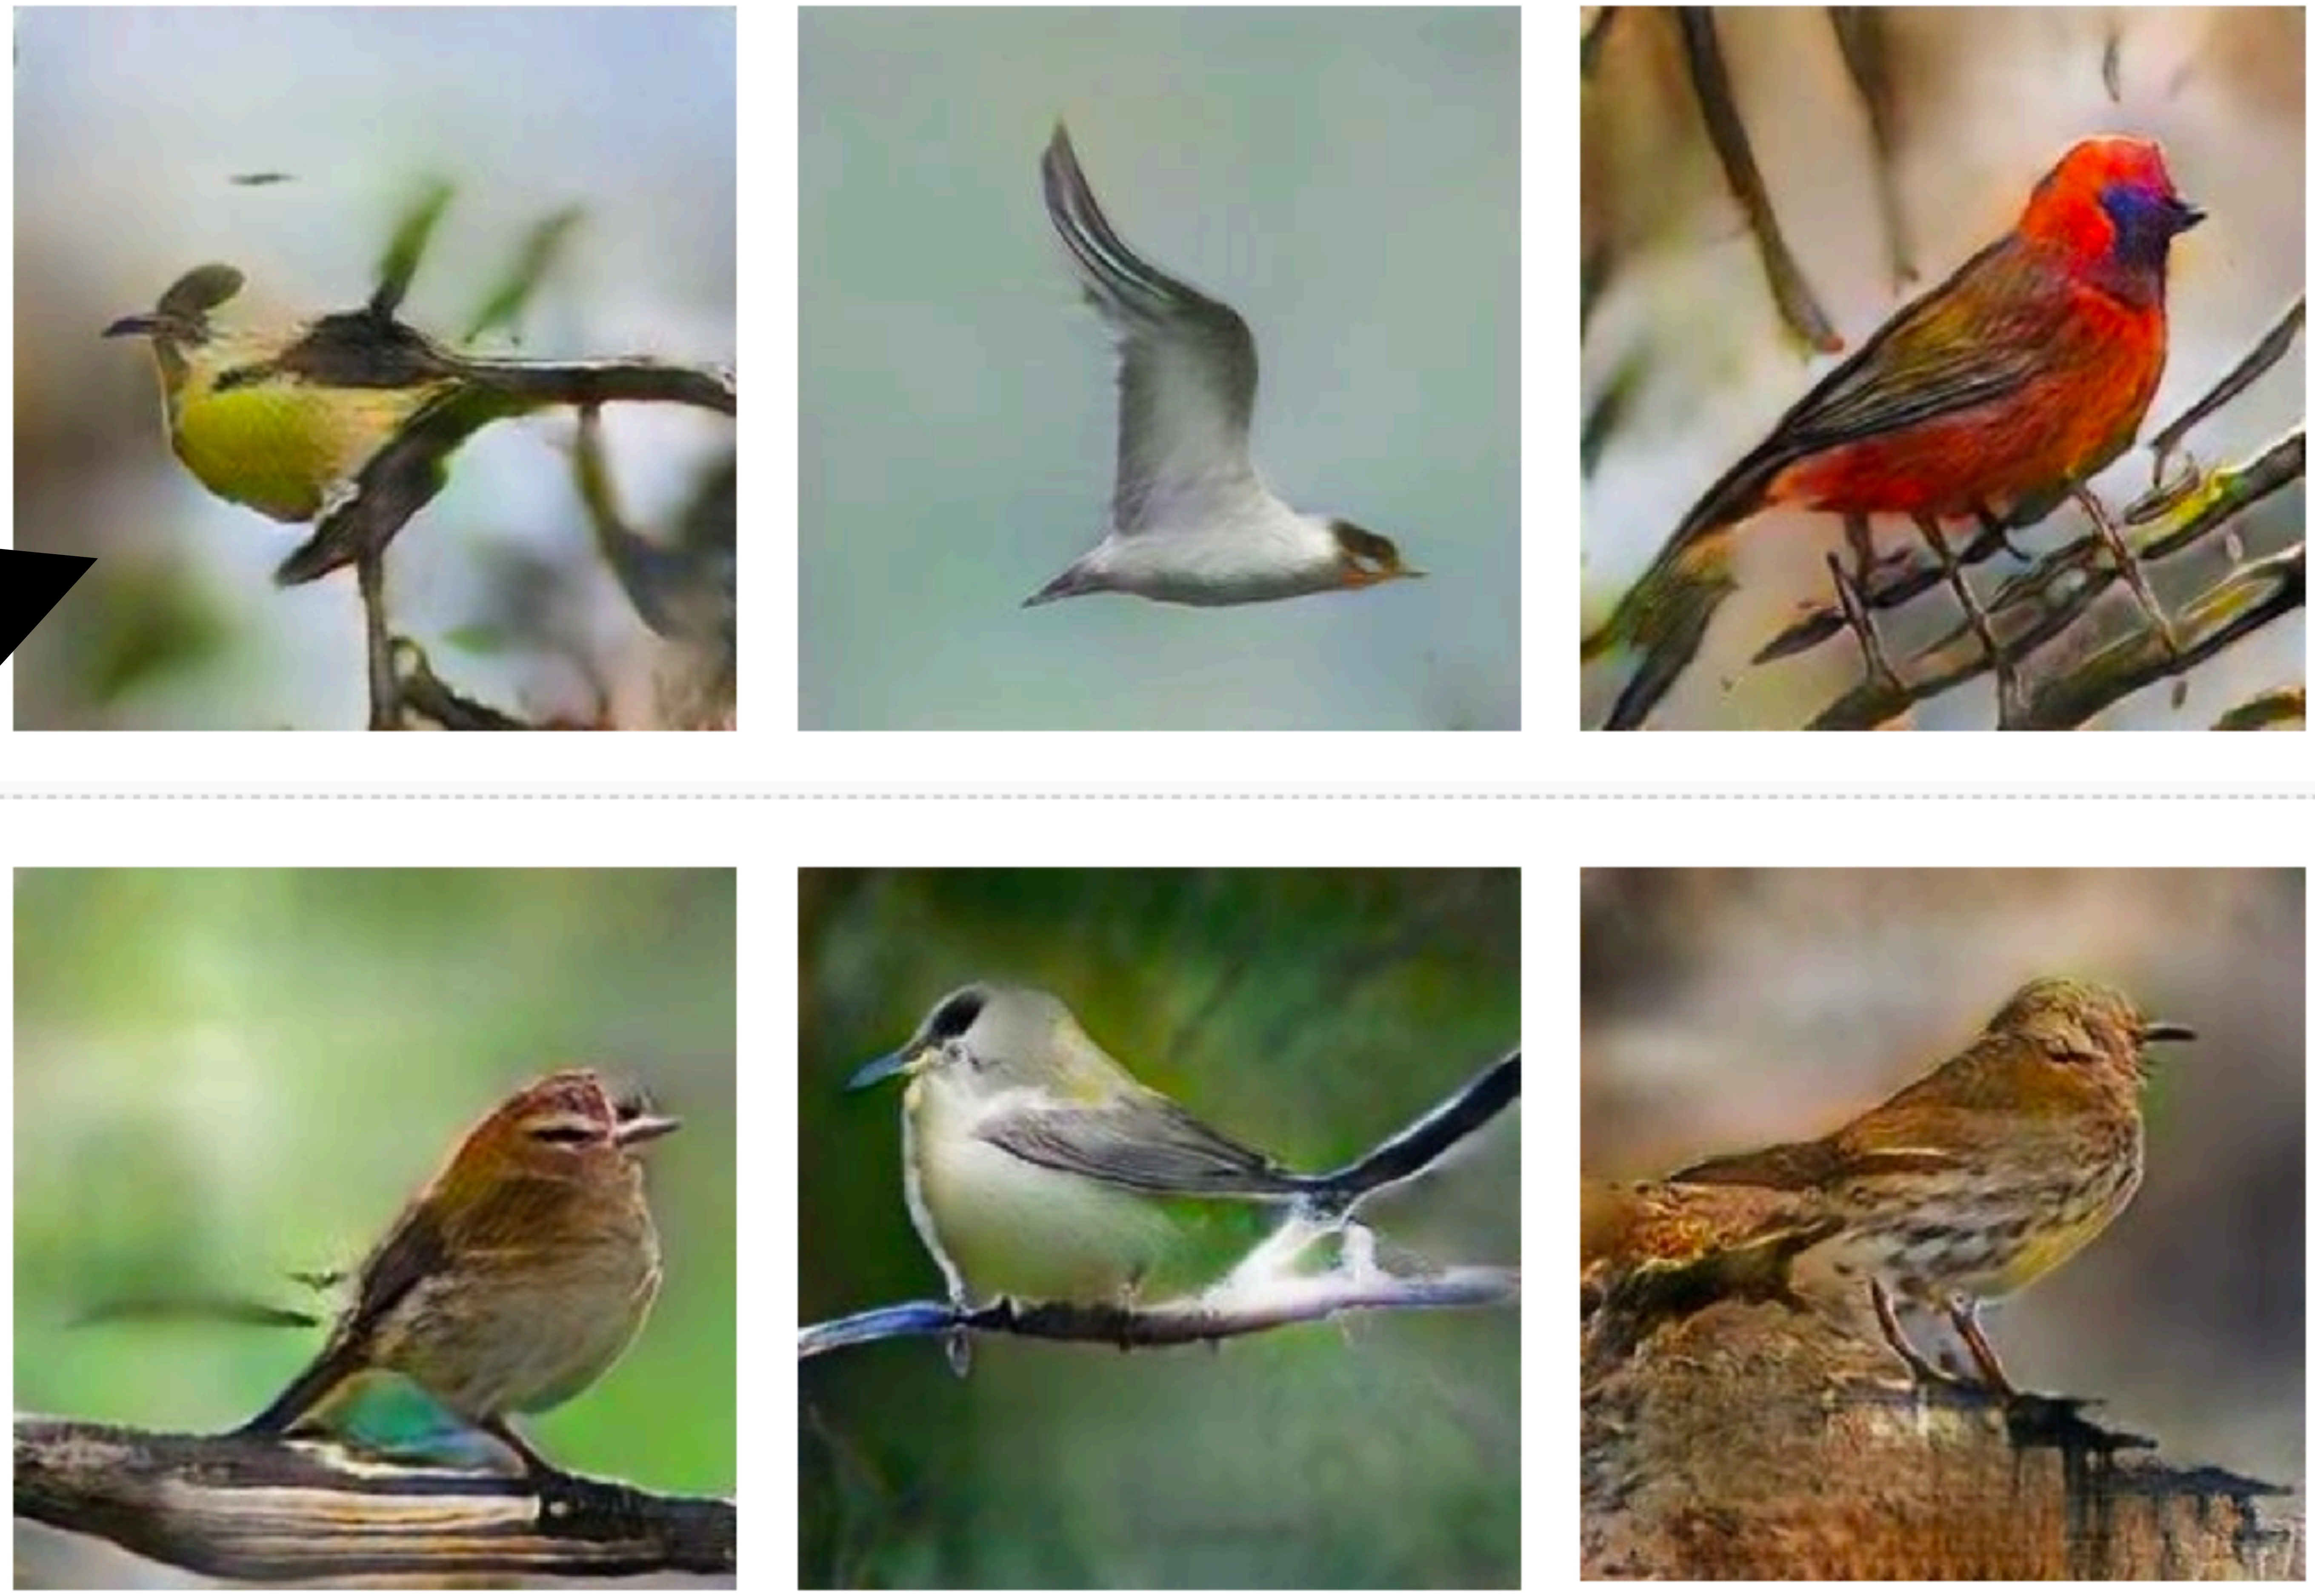

(Zhang et al., 2016)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

Simulating particle physics

Save millions of dollars of CPU time by predicting outcomes of explicit simulations

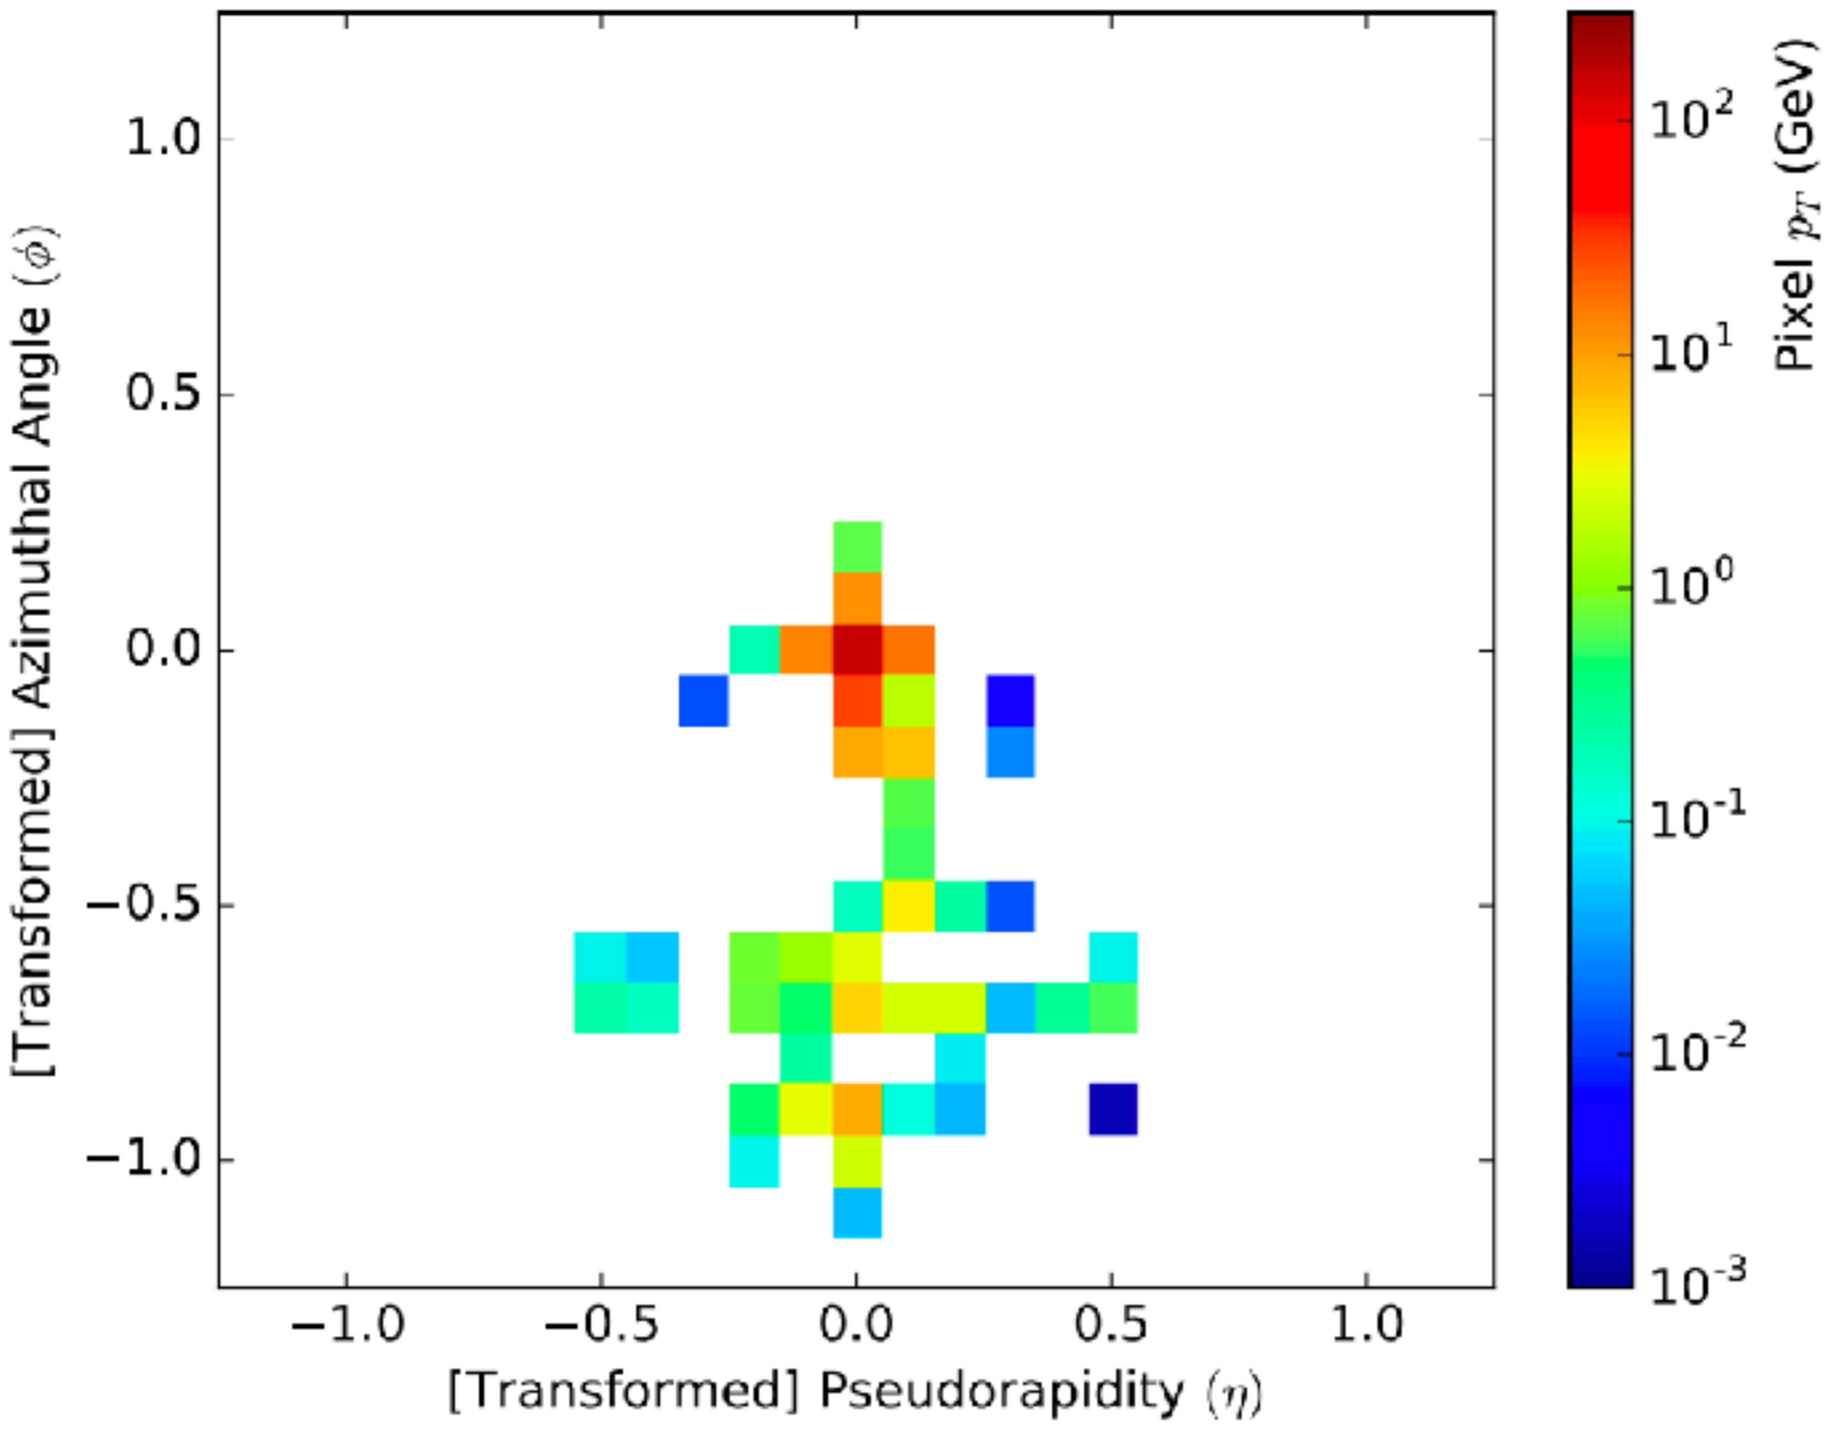

(de Oliveira et al., 2017)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

Adversarial Variational Bayes

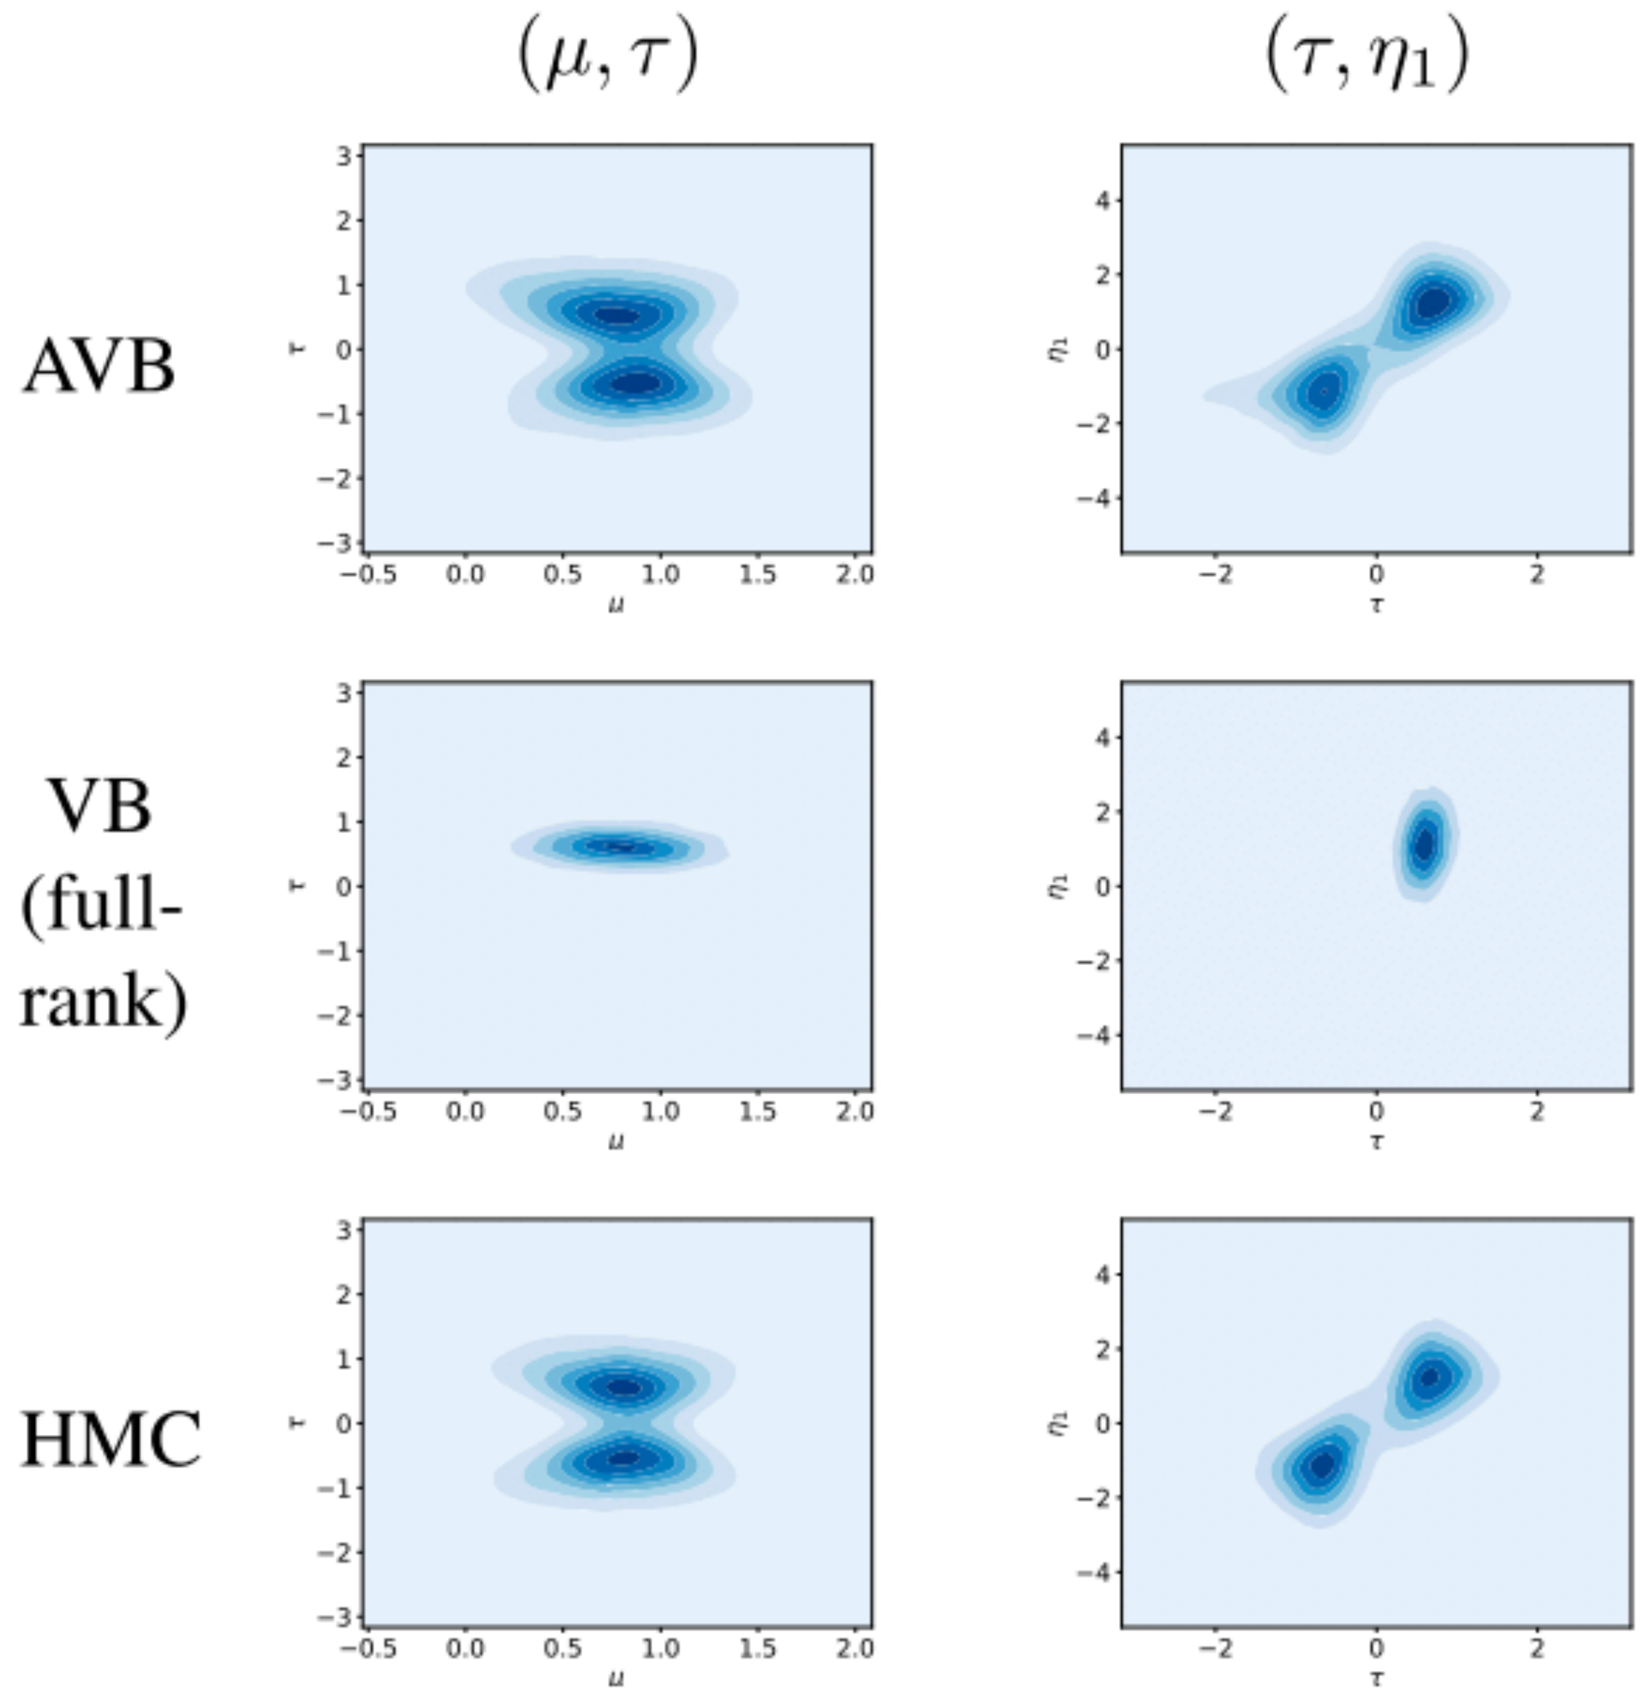

(Mescheder et al, 2017)

What can you do with GANs?

- Simulated environments and training data
- Missing data
  - Semi-supervised learning
- Multiple correct answers
- Realistic generation tasks
- Simulation by prediction
- Solve inference problems
- Learn useful embeddings

Vector Space Arithmetic

Man with glasses

-

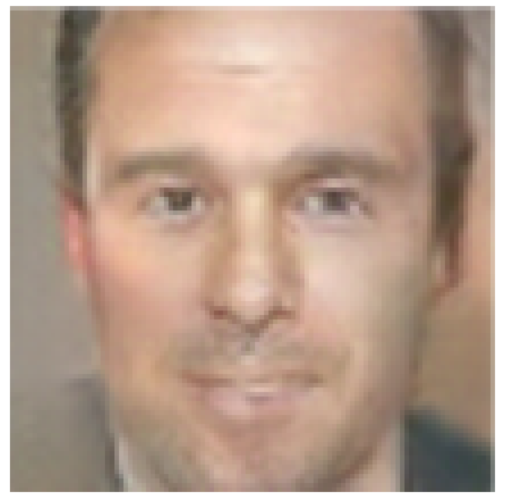

Man

+

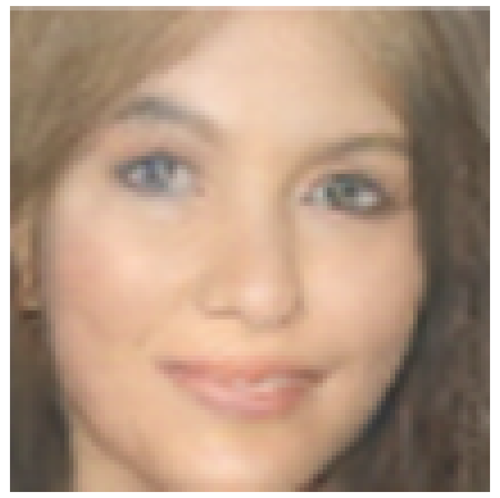

Woman

=

Woman with Glasses

(Radford et al, 2015)

Learning interpretable latent codes / controlling the generation process

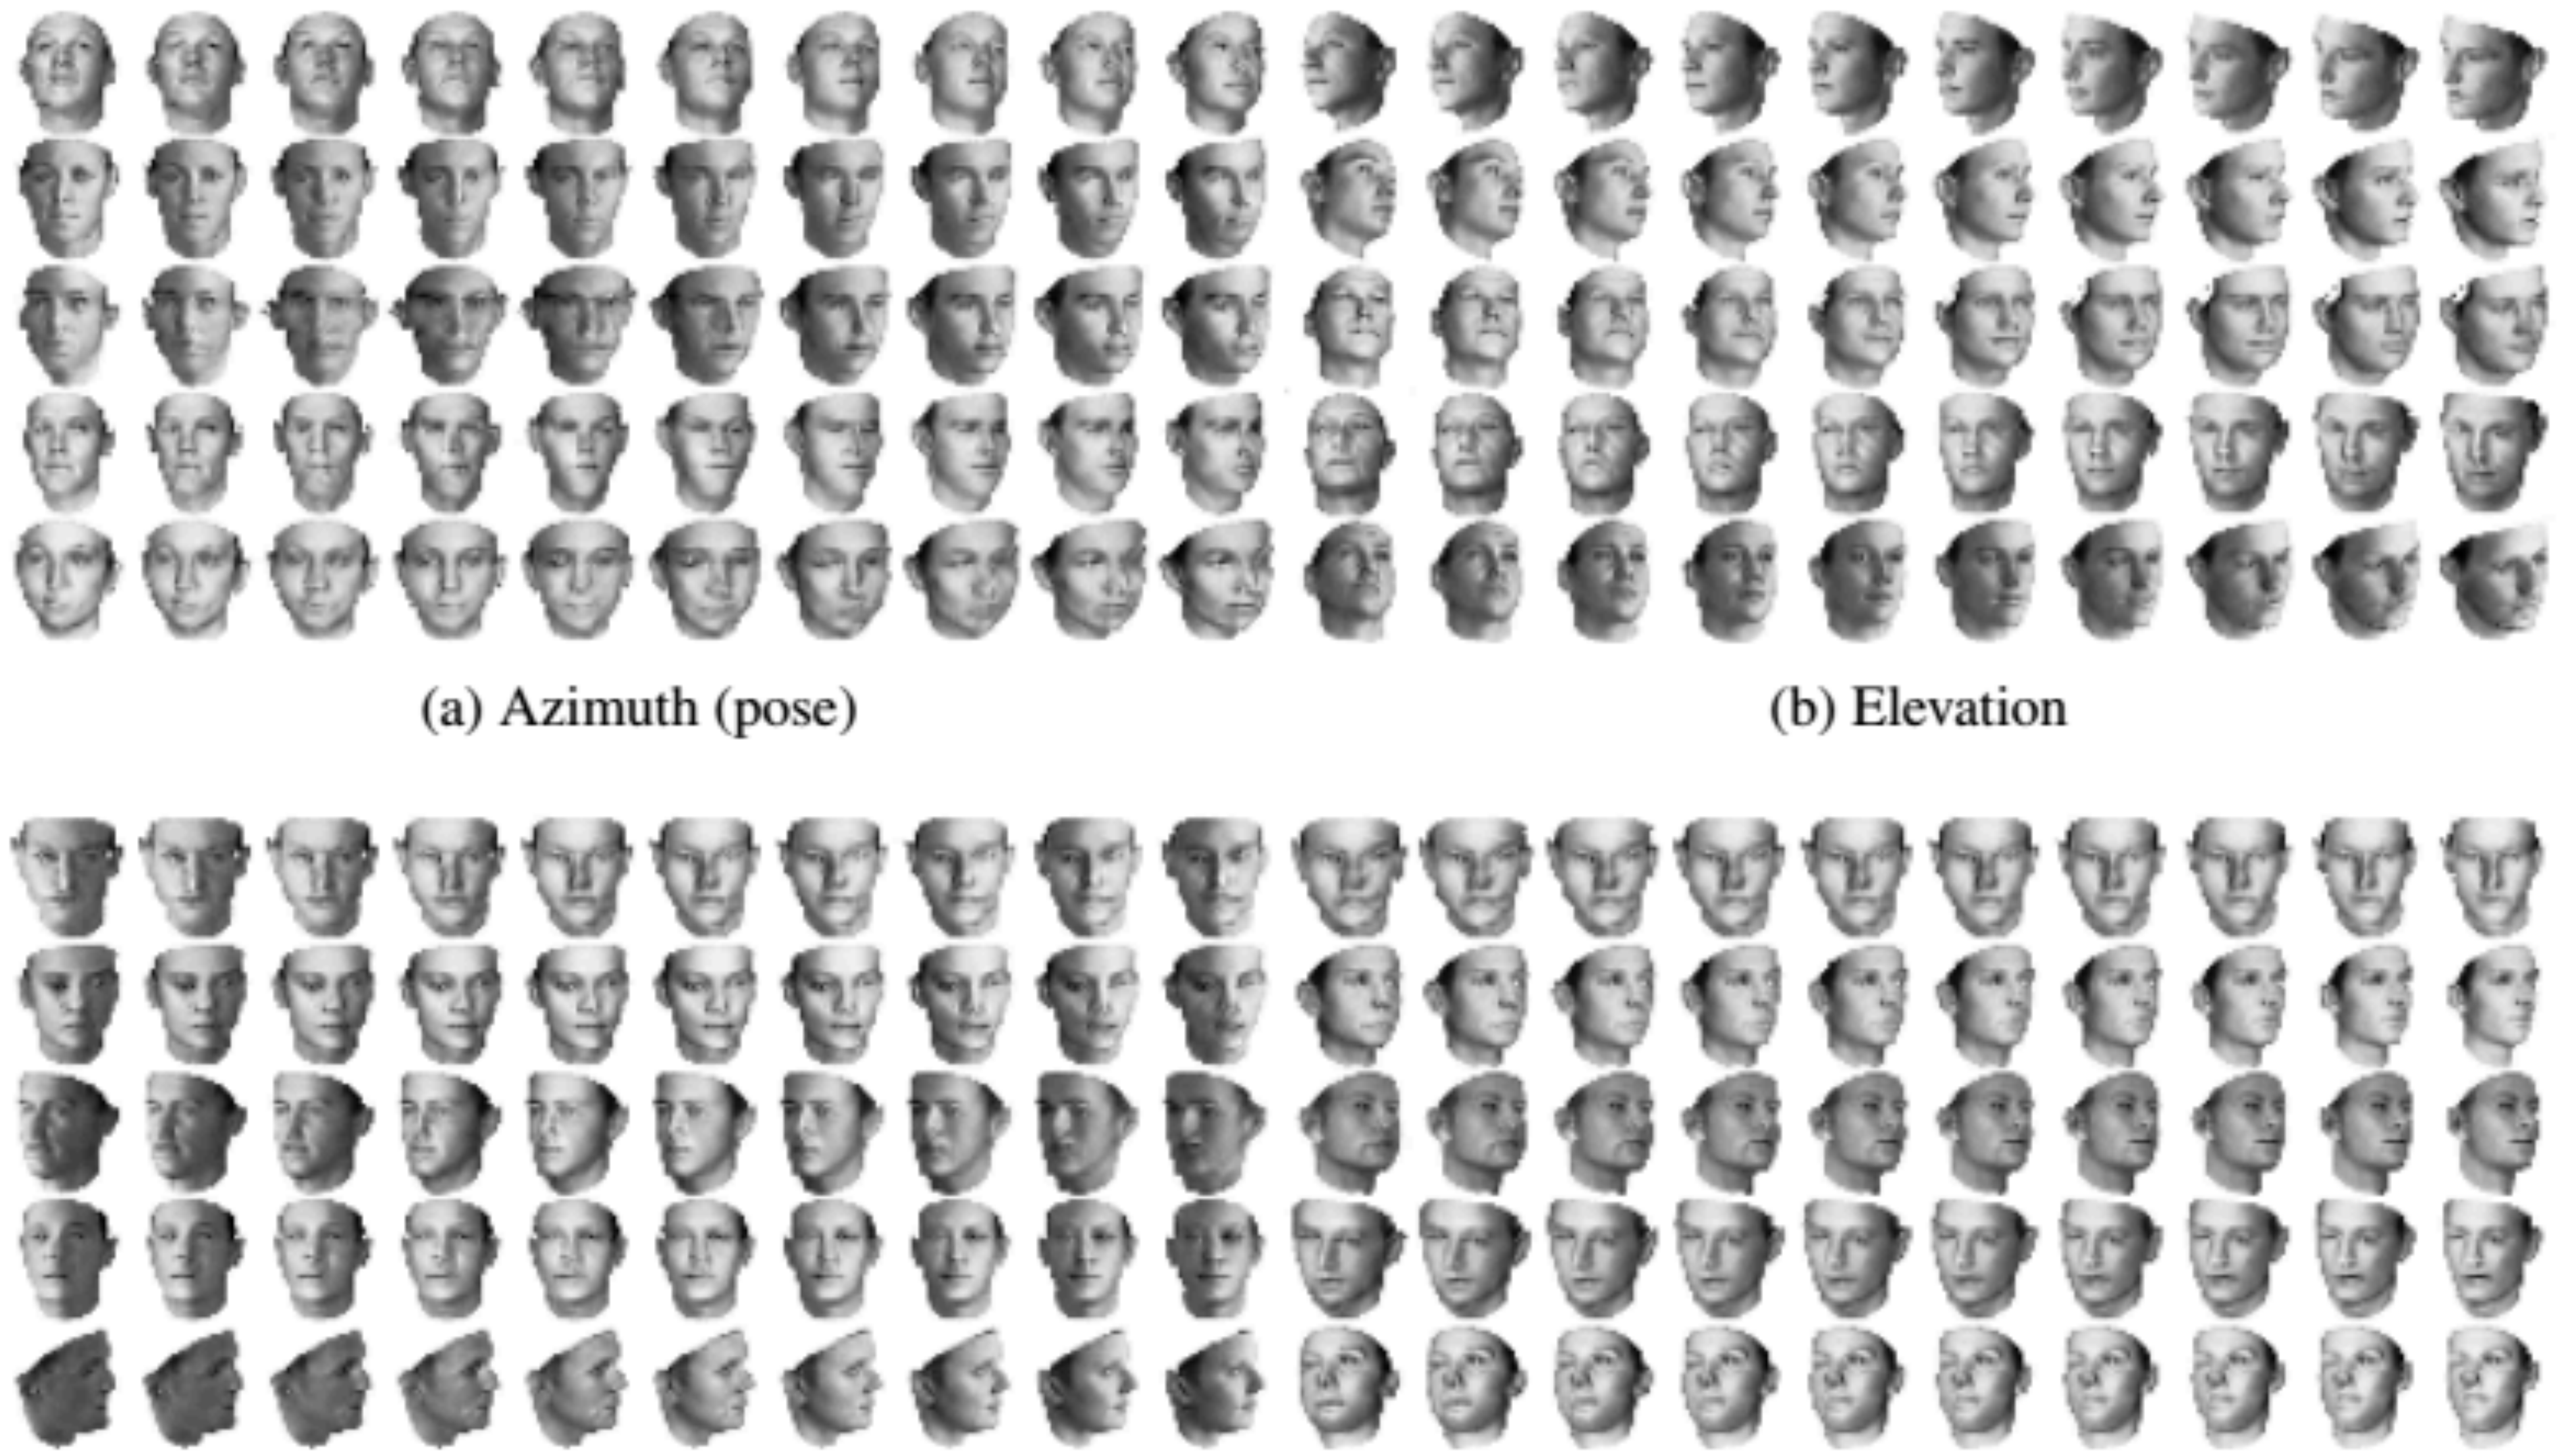

(a) Azimuth (pose) (b) Elevation

(c) Lighting (d) Wide or Narrow

InfoGAN (Chen et al 2016)

How long until GANs can do this?

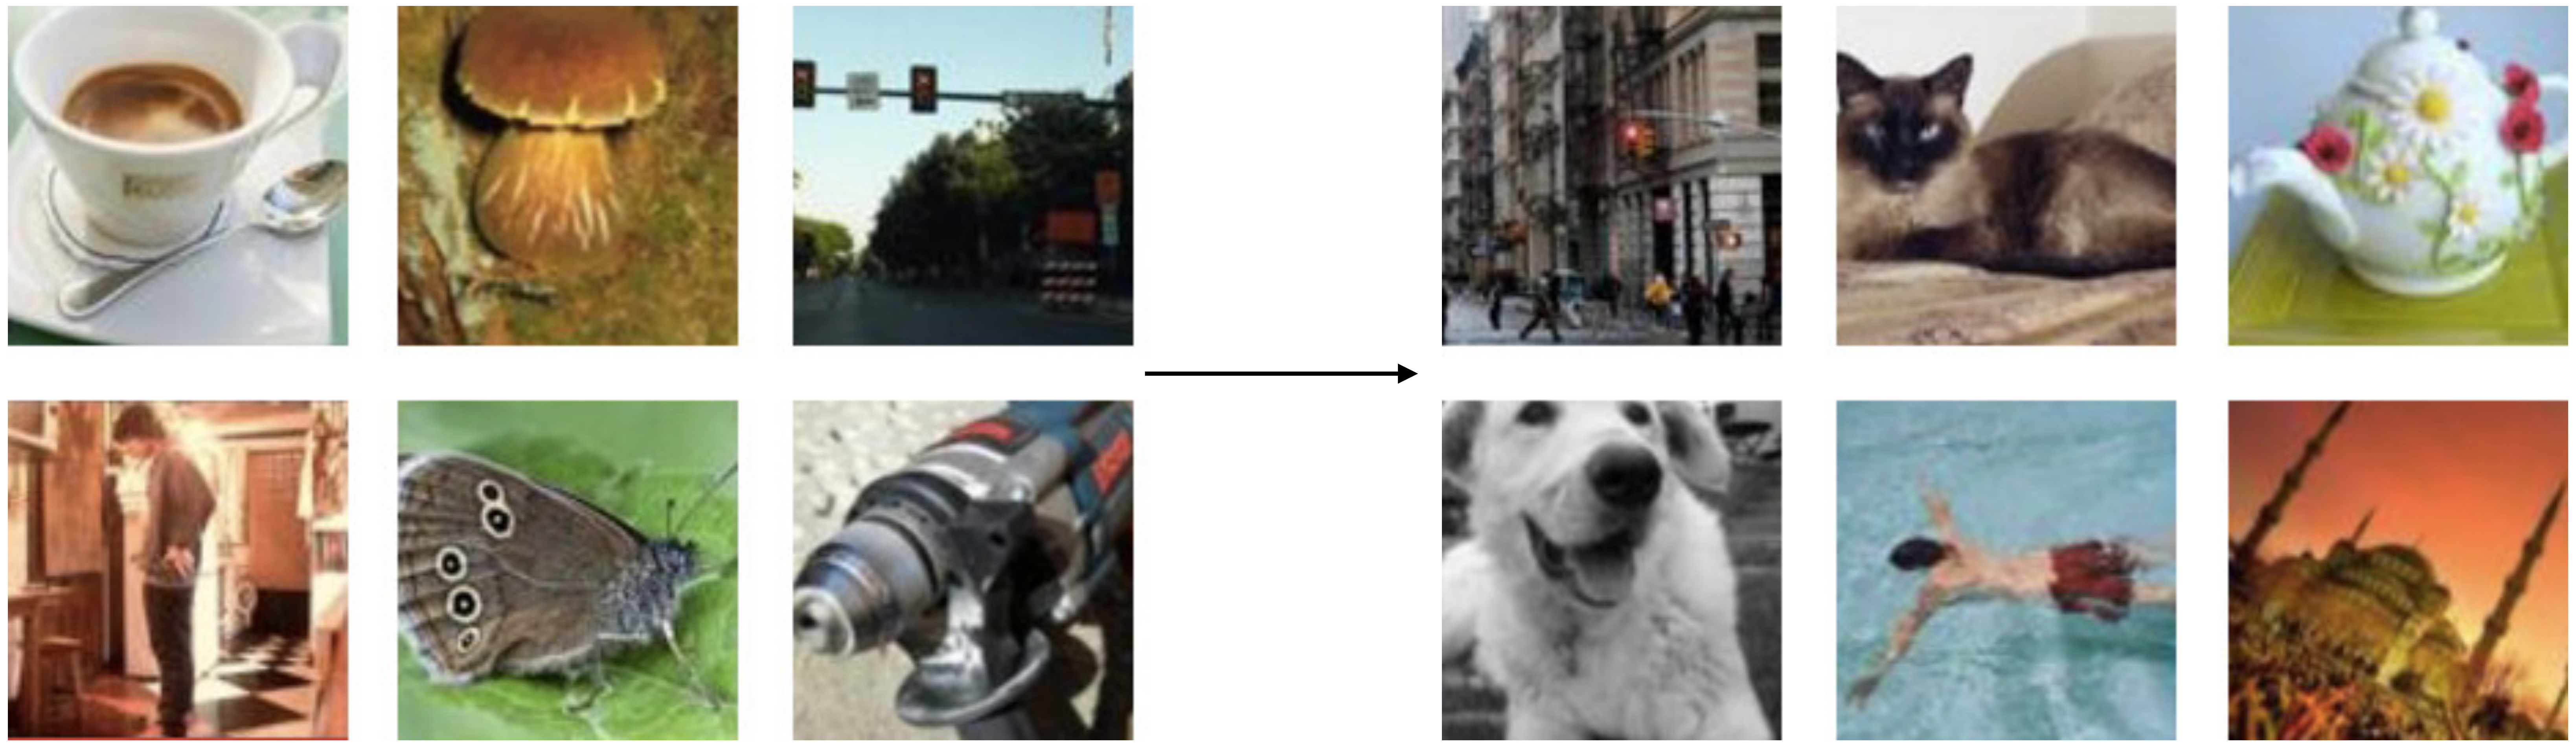

Training examples Model samples

AC-GANs

monarch butterfly goldfinch daisy

(Odena et al., 2016)

Minibatch GAN on ImageNet

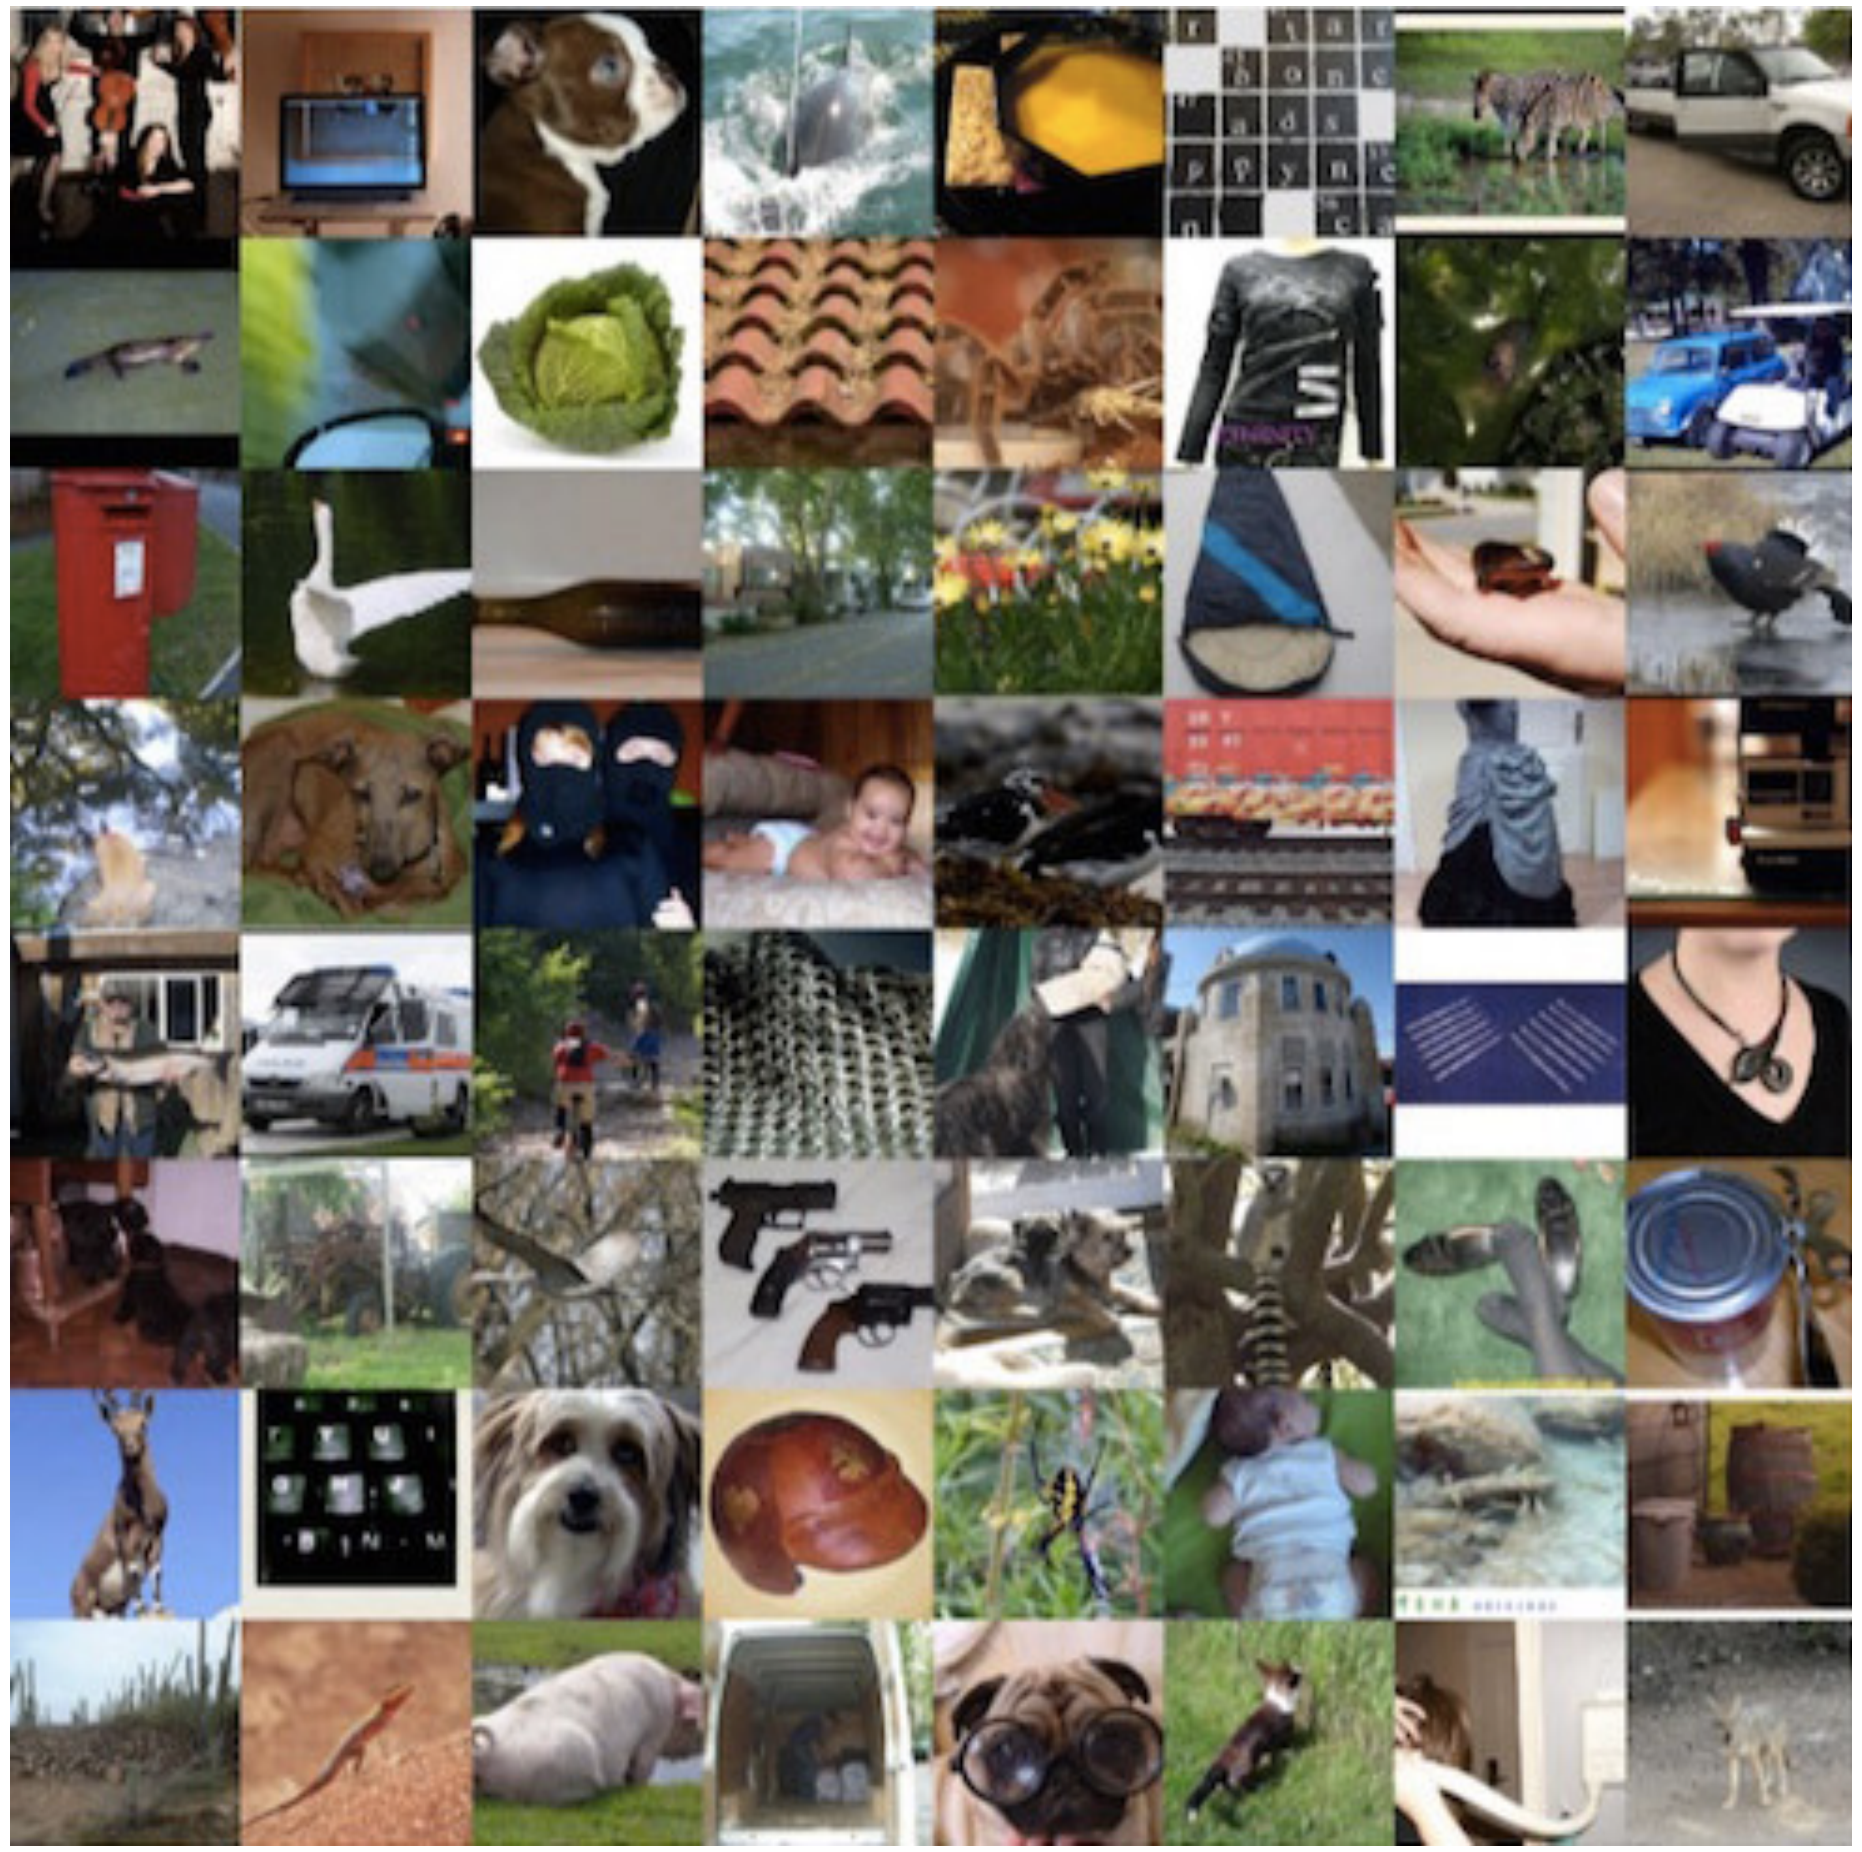

(Salimans et al., 2016)

Cherry-Picked Results





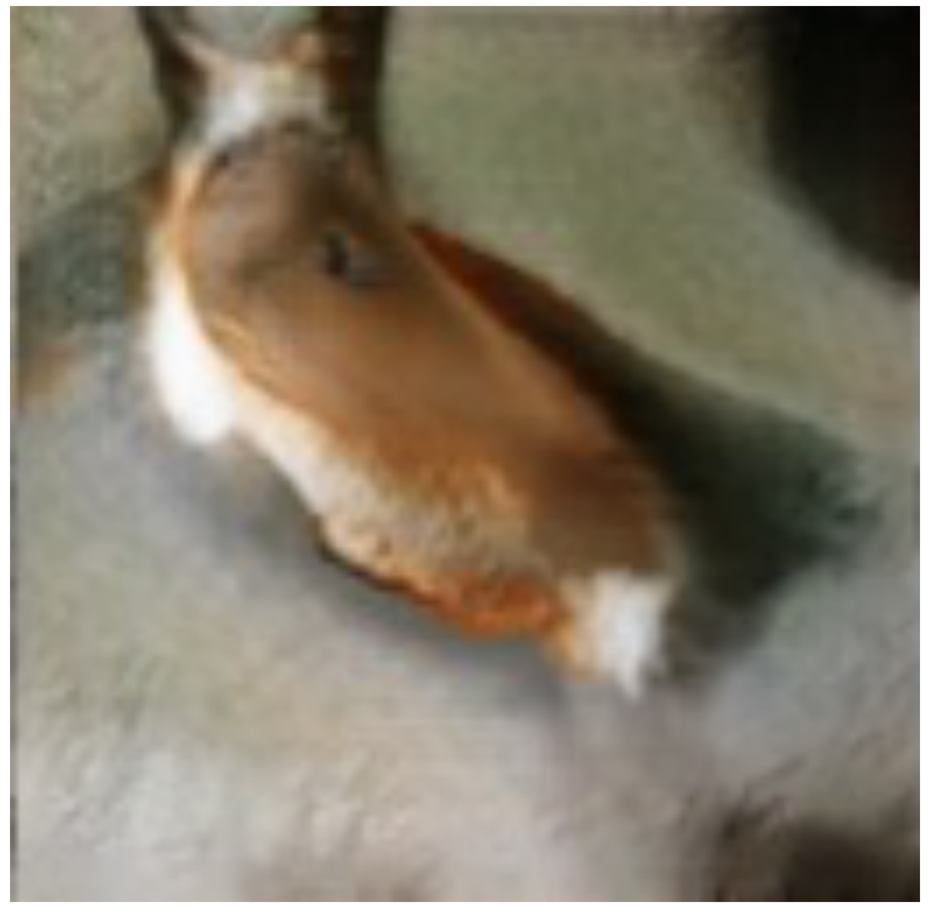

(Goodfellow 2017)

Problems with Counting

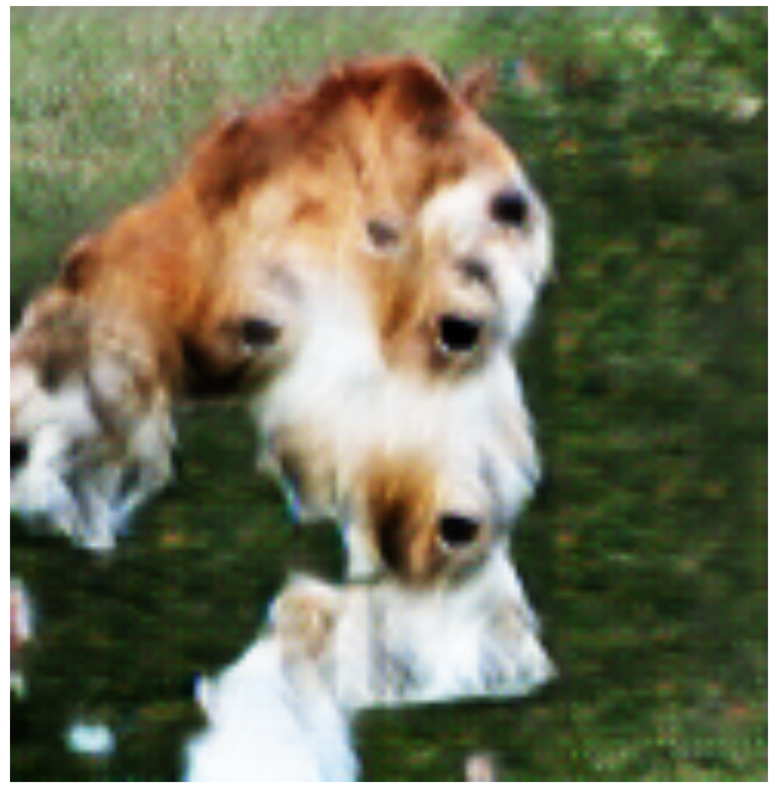




(Goodfellow 2017)

Problems with Perspective

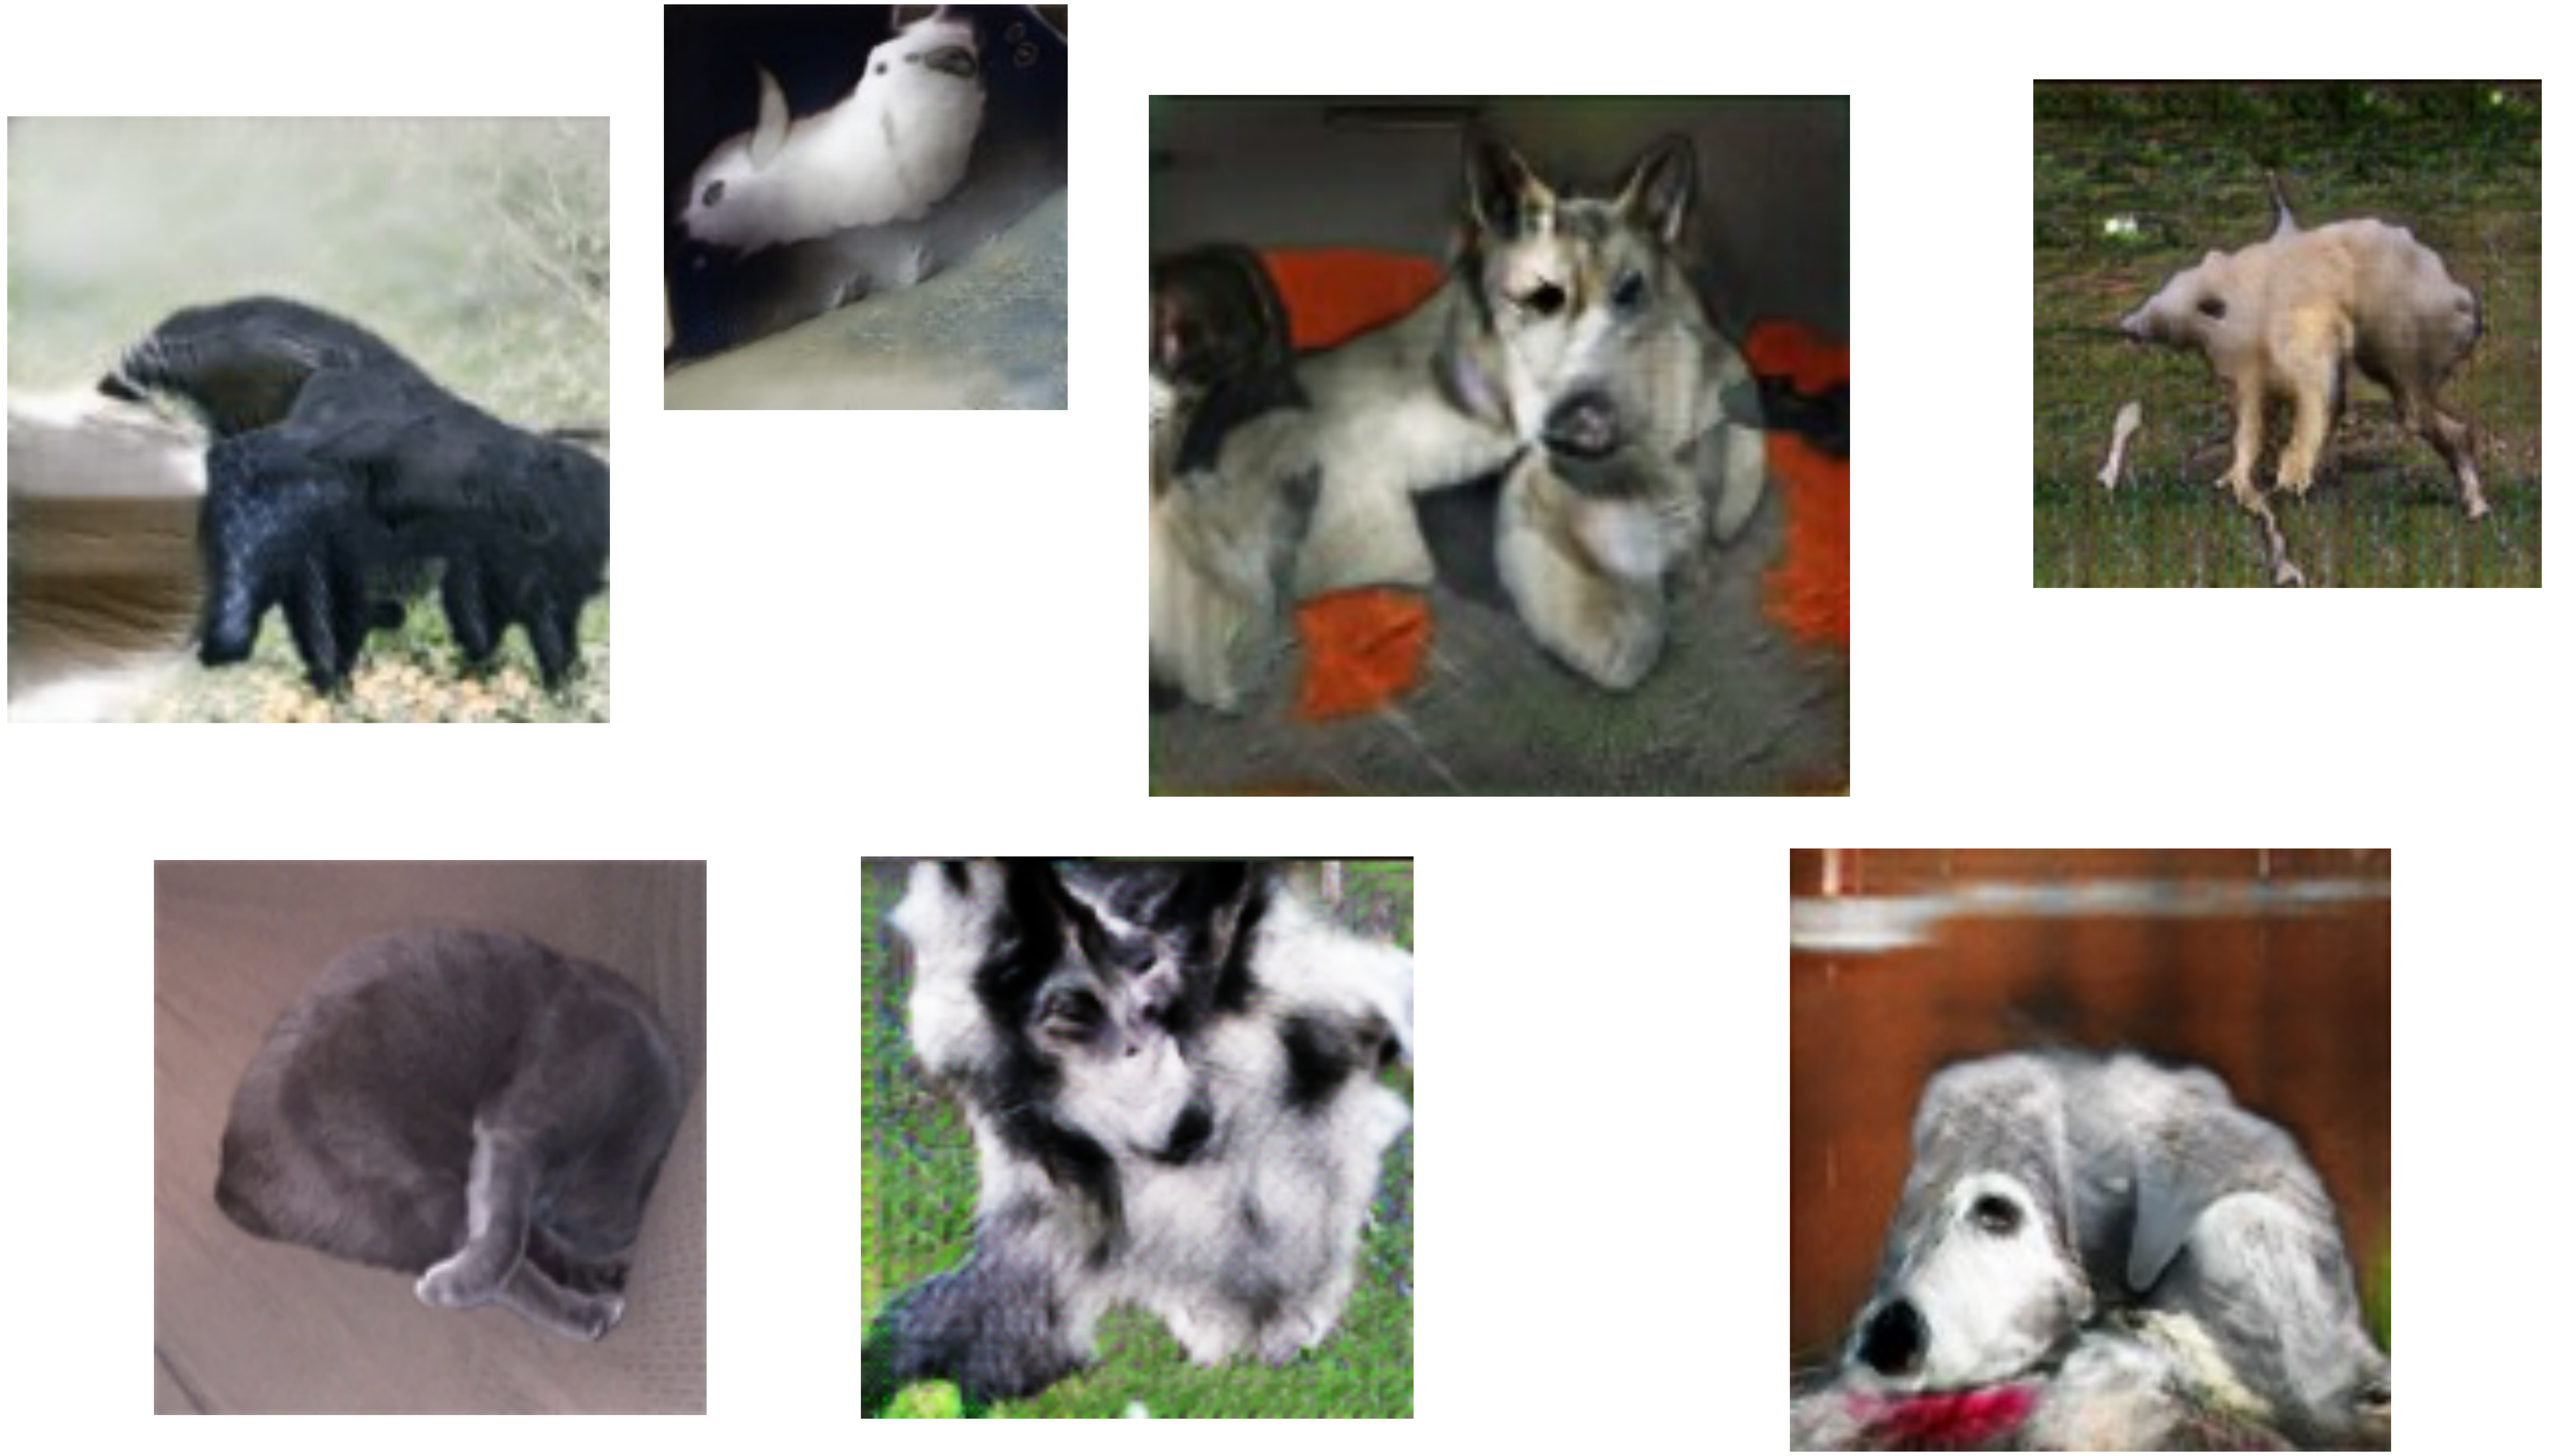

(Goodfellow 2017)

Problems with Global Structure





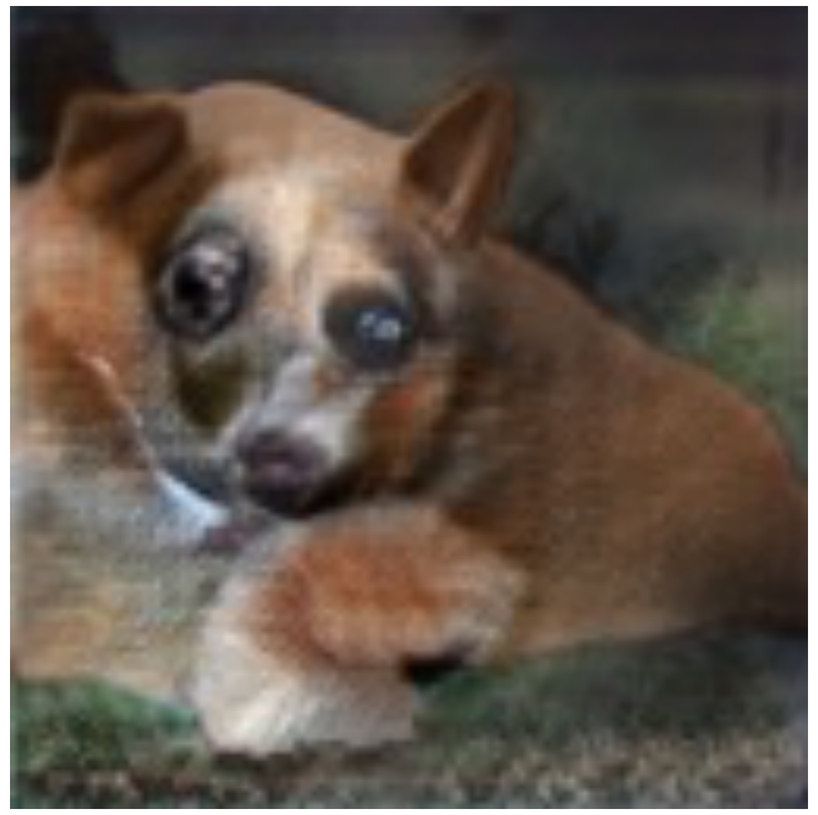
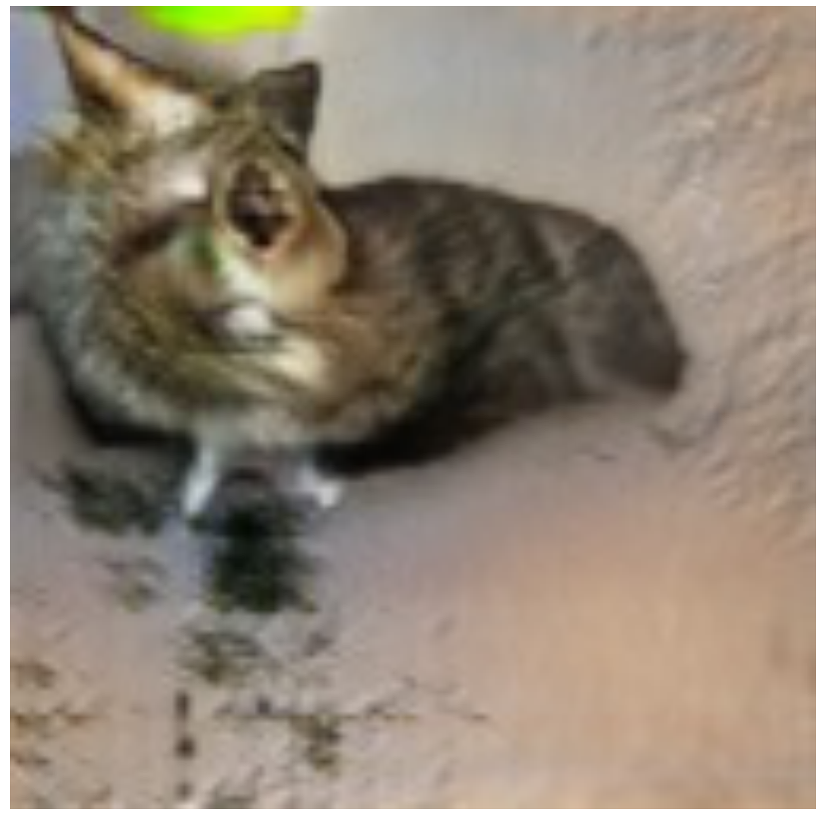

(Goodfellow 2017)

This one is real

(Goodfellow 2017)

Challenges

- Non-convergence, especially mode collapse
- Discrete output variables

(Goodfellow 2017)

Non-convergence

- Recent theoretical work argues that existing GAN training algorithms should converge under some reasonable conditions (Nagarajan and Kolter 2017, Heusel et al 2017)

- The convergence may be very slow because the Jacobian of the player's training gradients with respect to their parameters has unfortunate eigenvalue structure (Roth et al 2017)

- Mode collapse remains poorly understood; no widespread agreement on whether it is primarily a form of non-convergence

Discrete output variables

- Tasks like text generation for machine translation require a generator that produces discrete outputs

- GAN training requires the output to be differentiable with respect to the generator parameters

- Straightforward approaches like Gumbel-Softmax and REINFORCE have so far been disappointing on NLP tasks

(Goodfellow 2017)

Track updates at the GAN Zoo

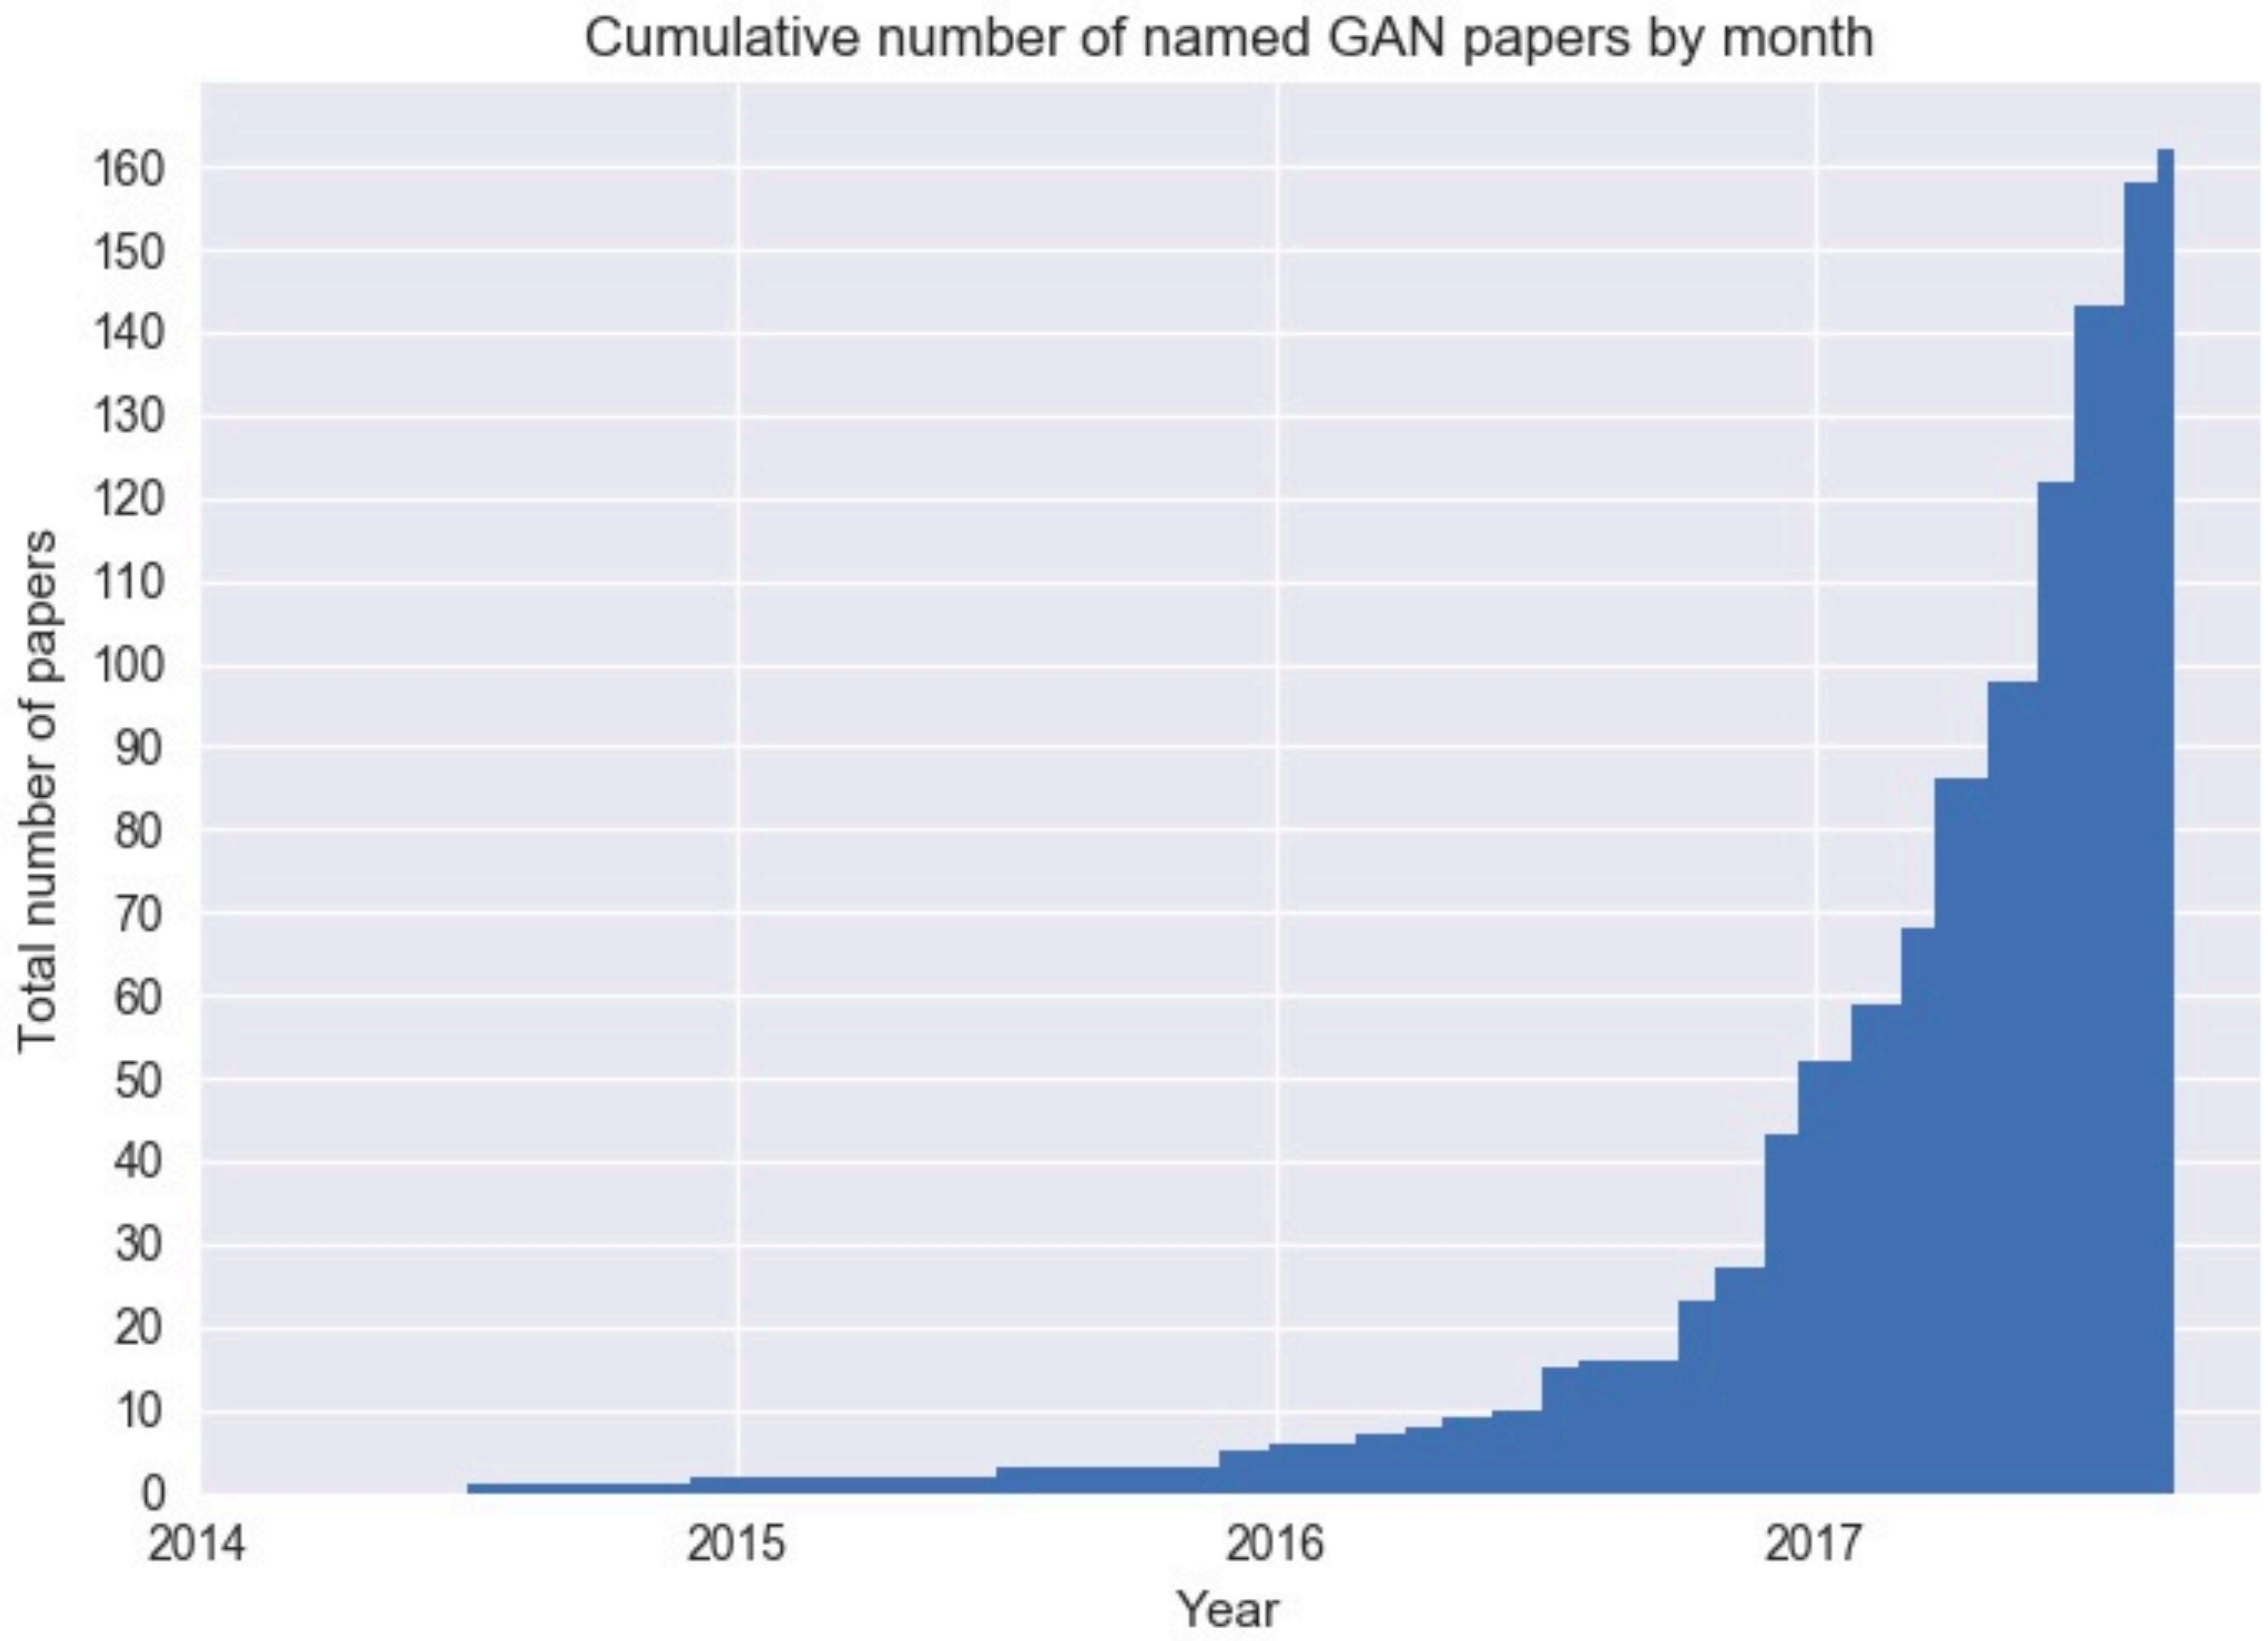

https://github.com/hindupuravinash/the-gan-zoo

(Goodfellow 2017)

Conclusion

- GANs are generative models based on game theory
- GANs open the door to a wide range of engineering tasks
- There are still important research challenges to solve before GANs can generate arbitrary data

(Goodfellow 2017)