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 GPU Technology Conference

San Jose, California 2017-05-09

McGAN MIX+GAN C-RNN-GAN DR-GAN MGAN GoGAN BS-GAN C-VAE-GAN FF-GAN

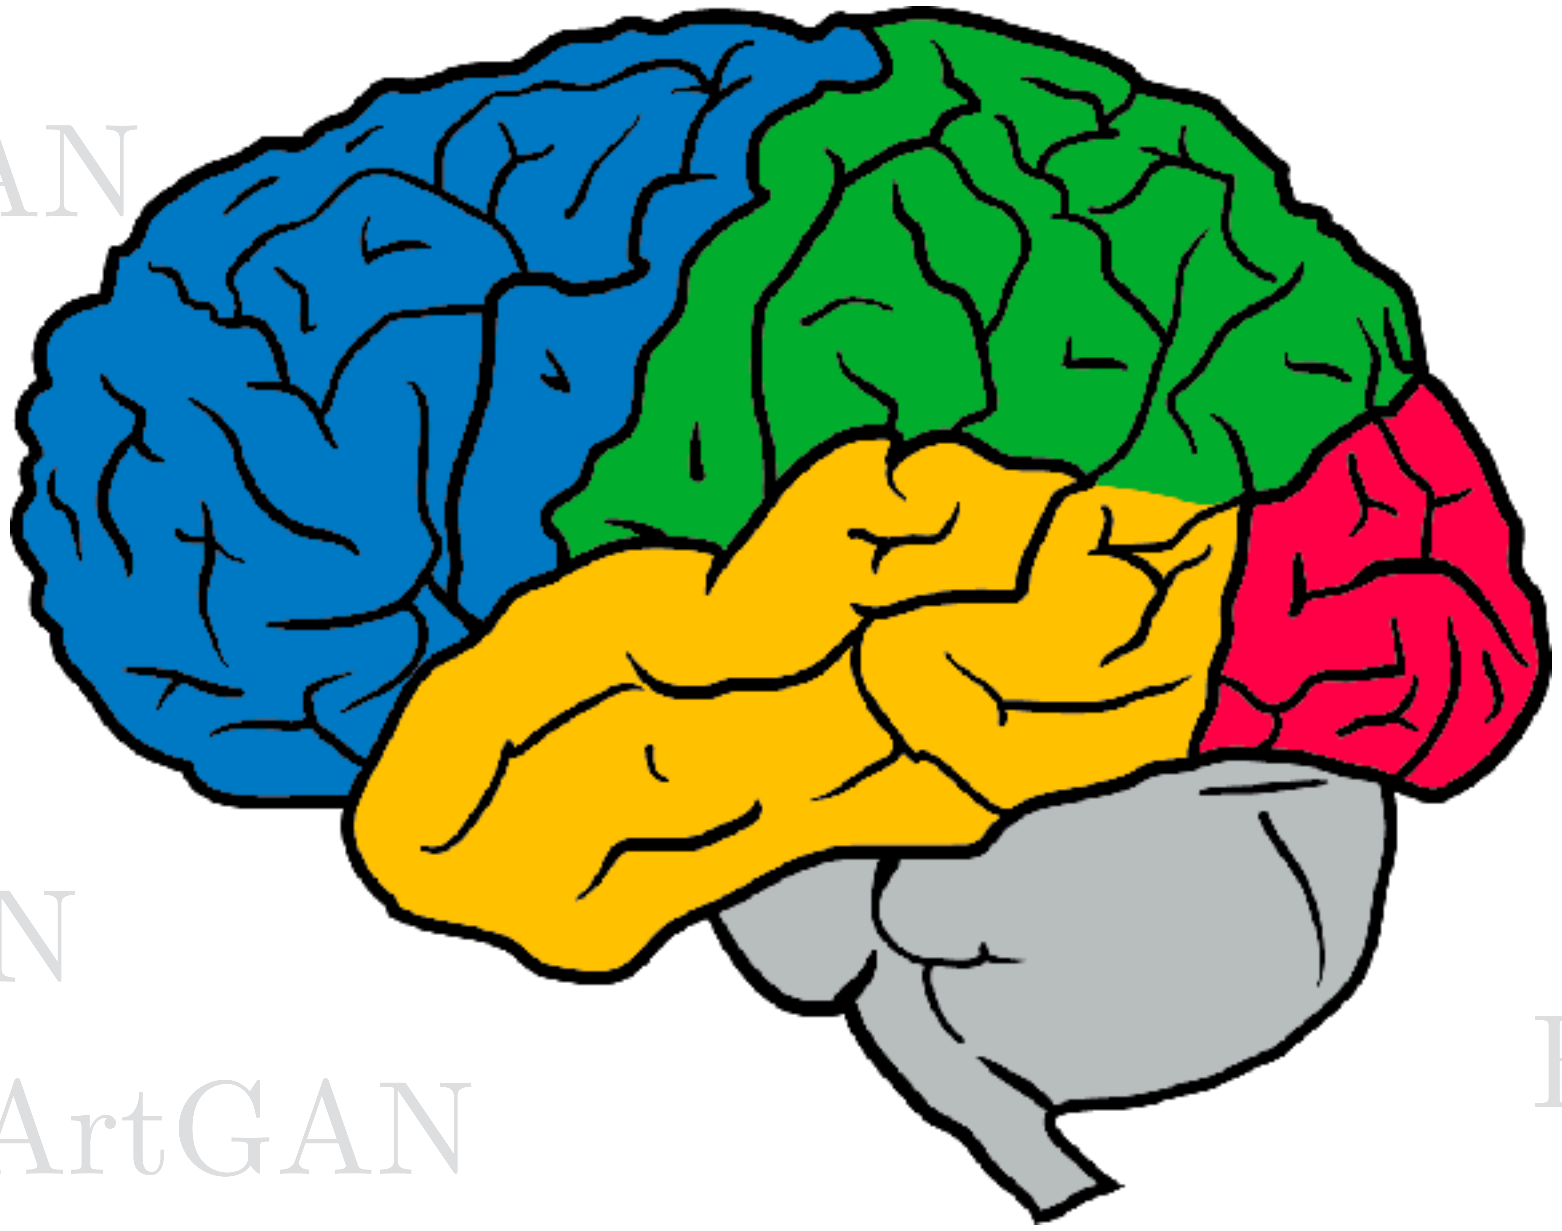

CCGAN MAGAN 3D-GAN AC-GAN DCGAN GAWWN DualGAN CycleGAN BiGAN Bayesian GAN GP-GAN AnoGAN EBGAN DTN Context-RNN-GAN MAD-GAN ALI f-GAN ArtGAN BEGAN AL-CGAN MARTA-GAN MalGAN

Generative Modeling

- Density estimation

- Sample generation

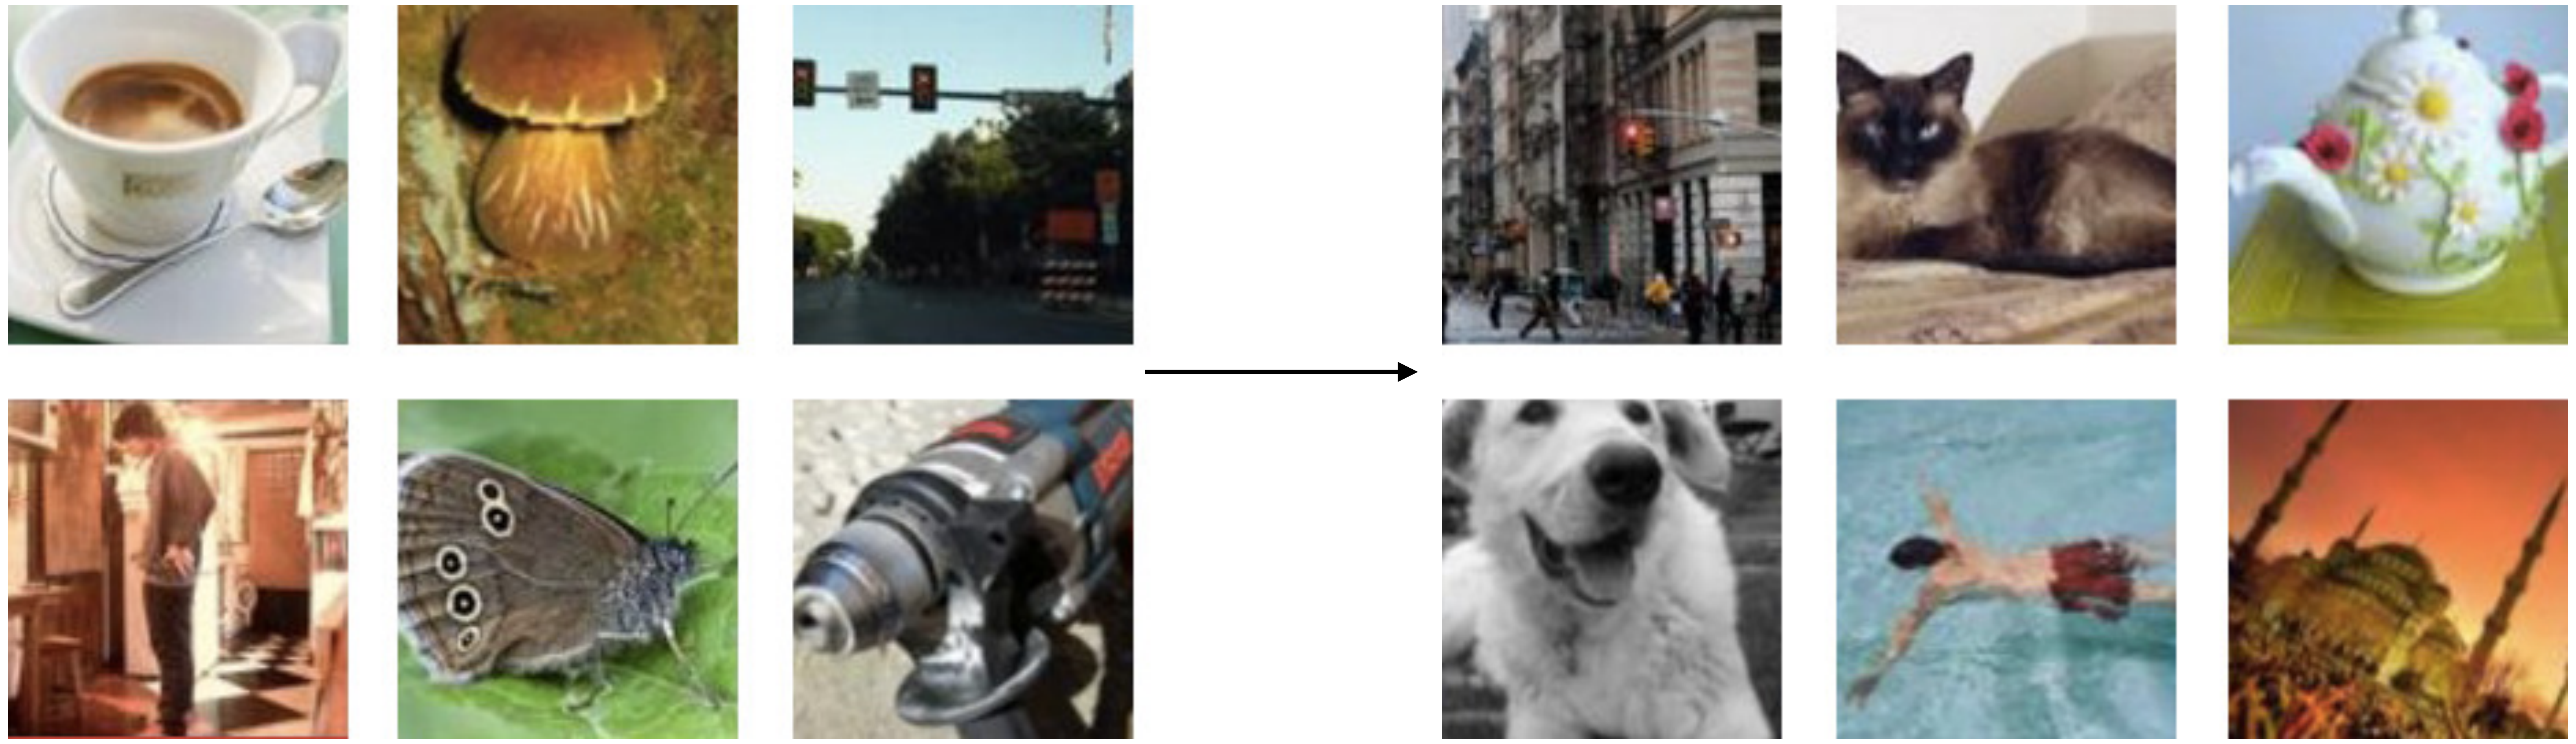

Training examples   Model samples

(Goodfellow 2017)

Maximum Likelihood

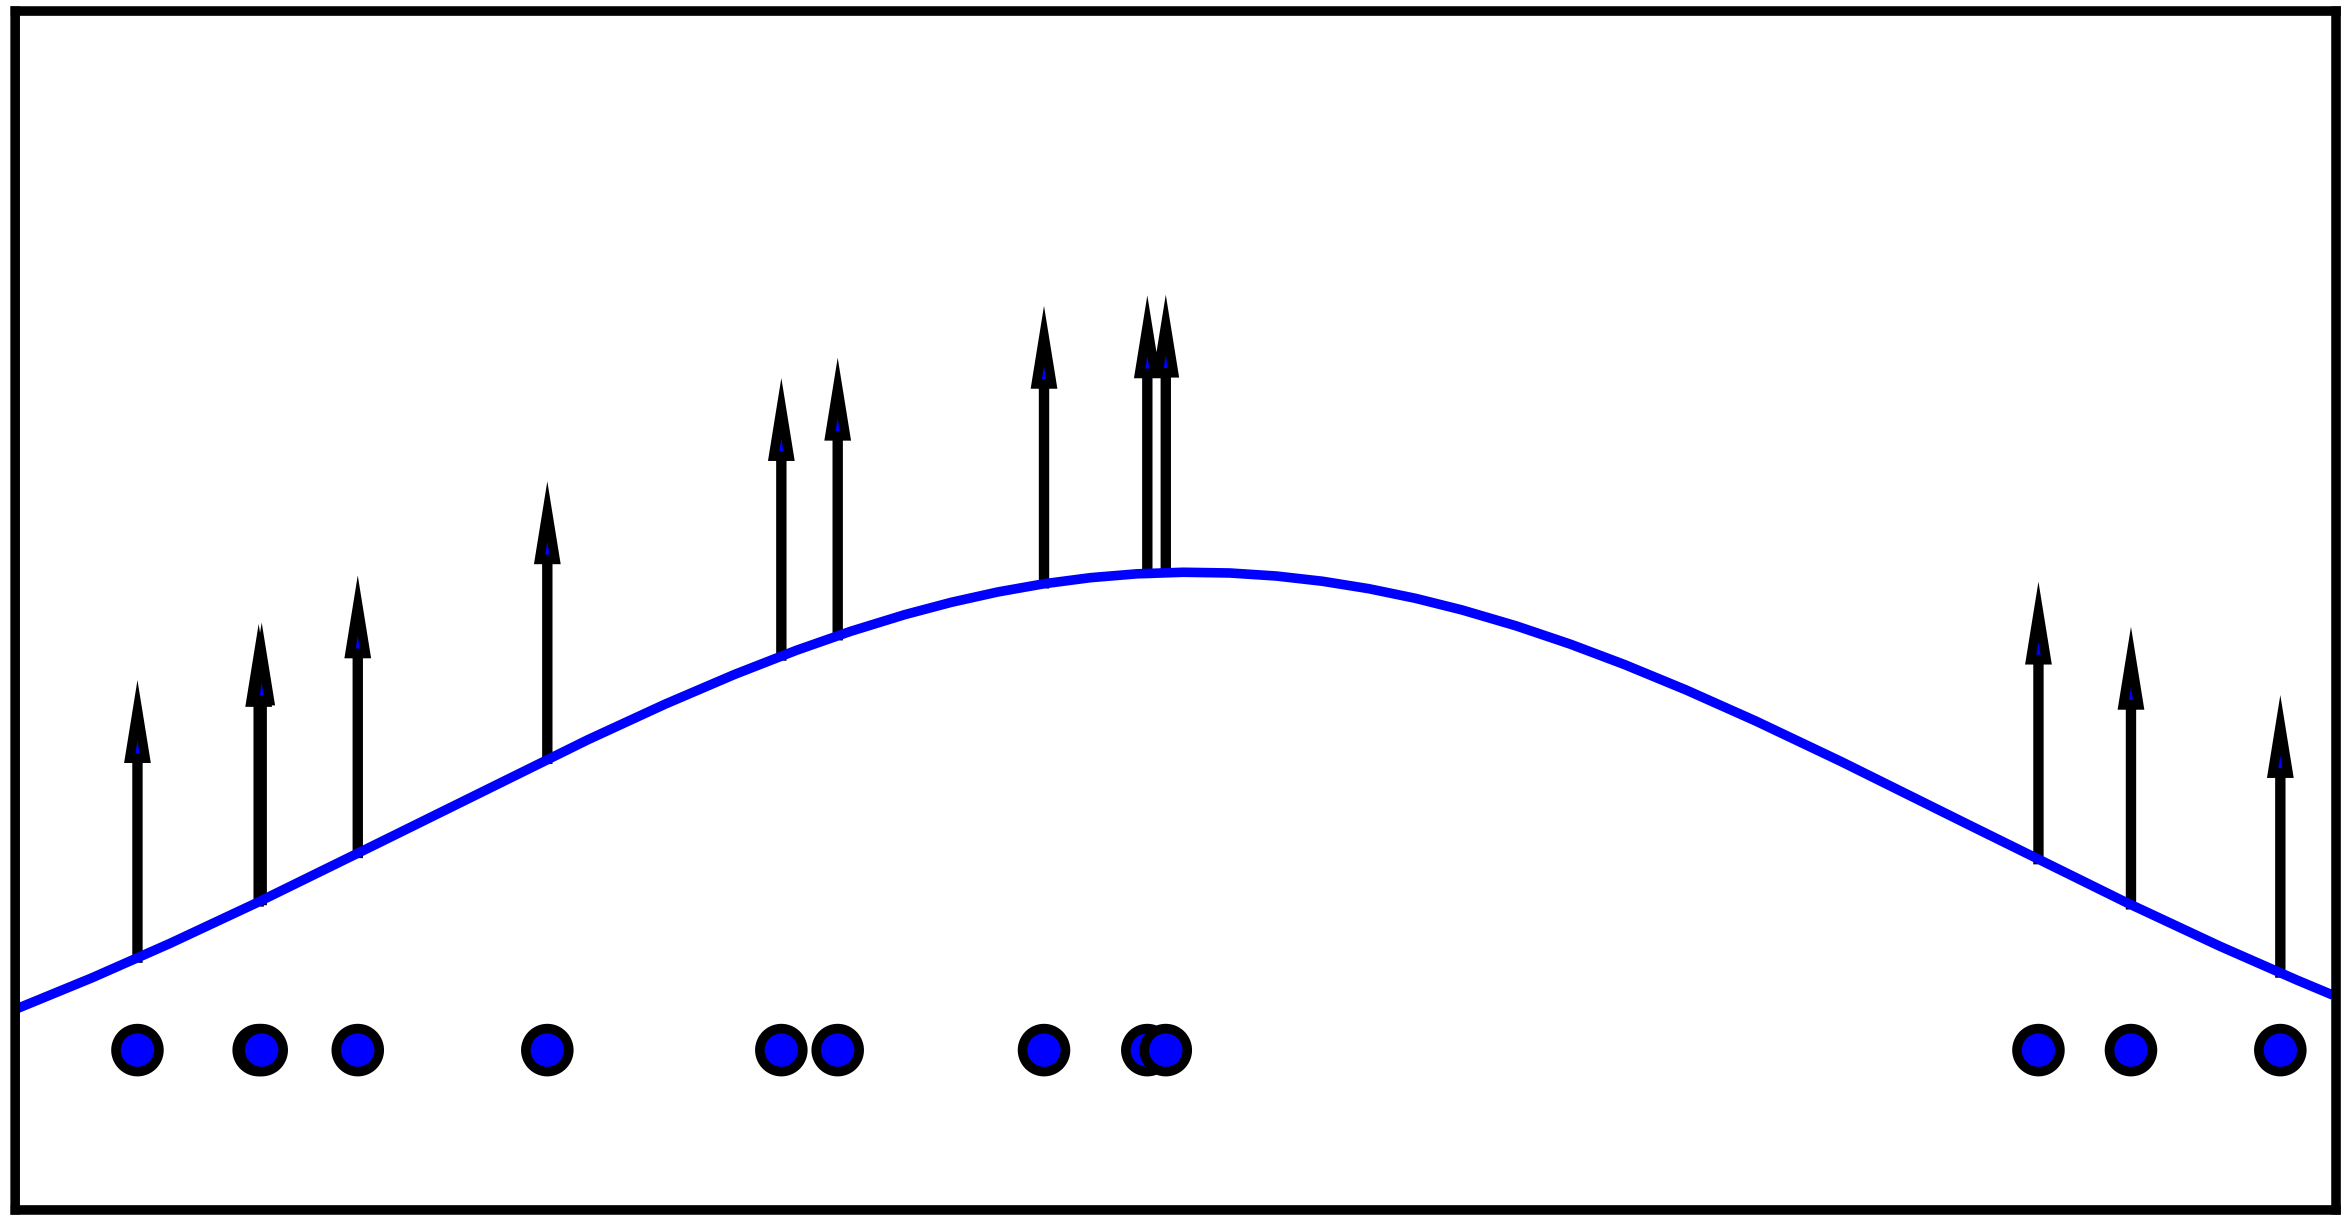

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} \mathbb{E}_{x\sim p_{\text{data}}} \log p_{\text{model}}(\boldsymbol{x} \mid \boldsymbol{\theta})$$

Adversarial Nets Framework

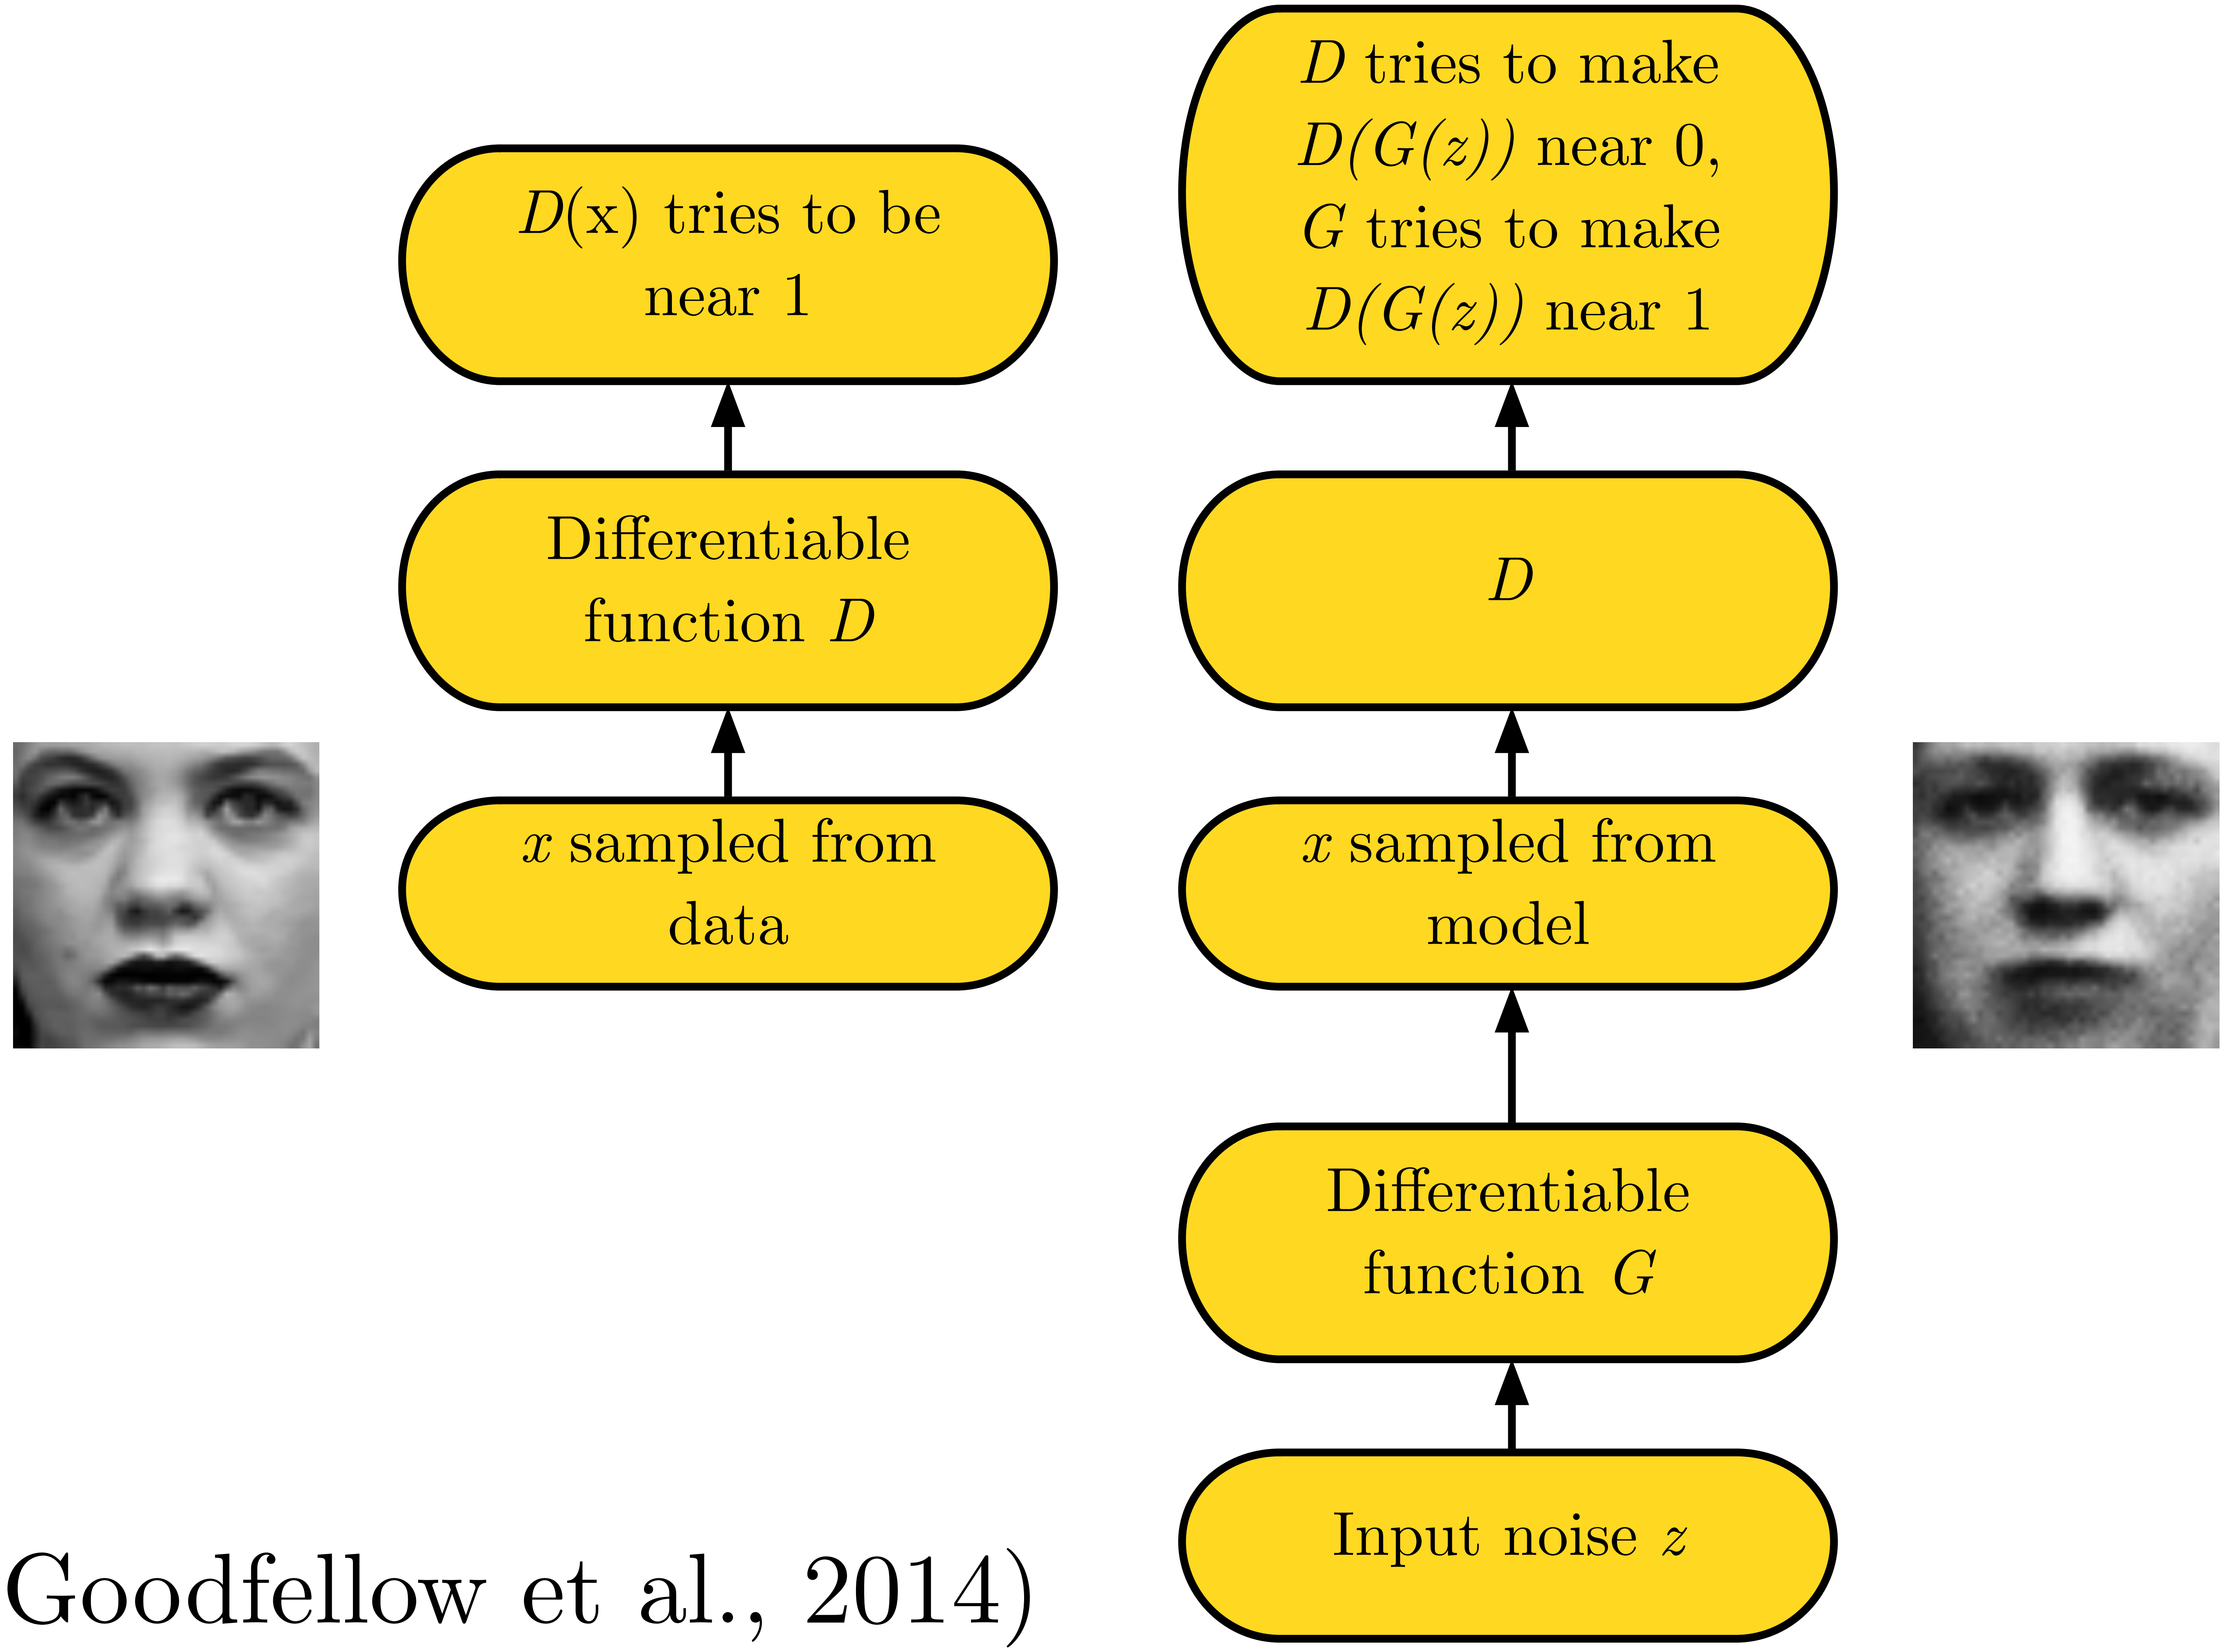

(Goodfellow et al., 2014)

(Goodfellow 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

TEACHING AID

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

GANs for simulated training data

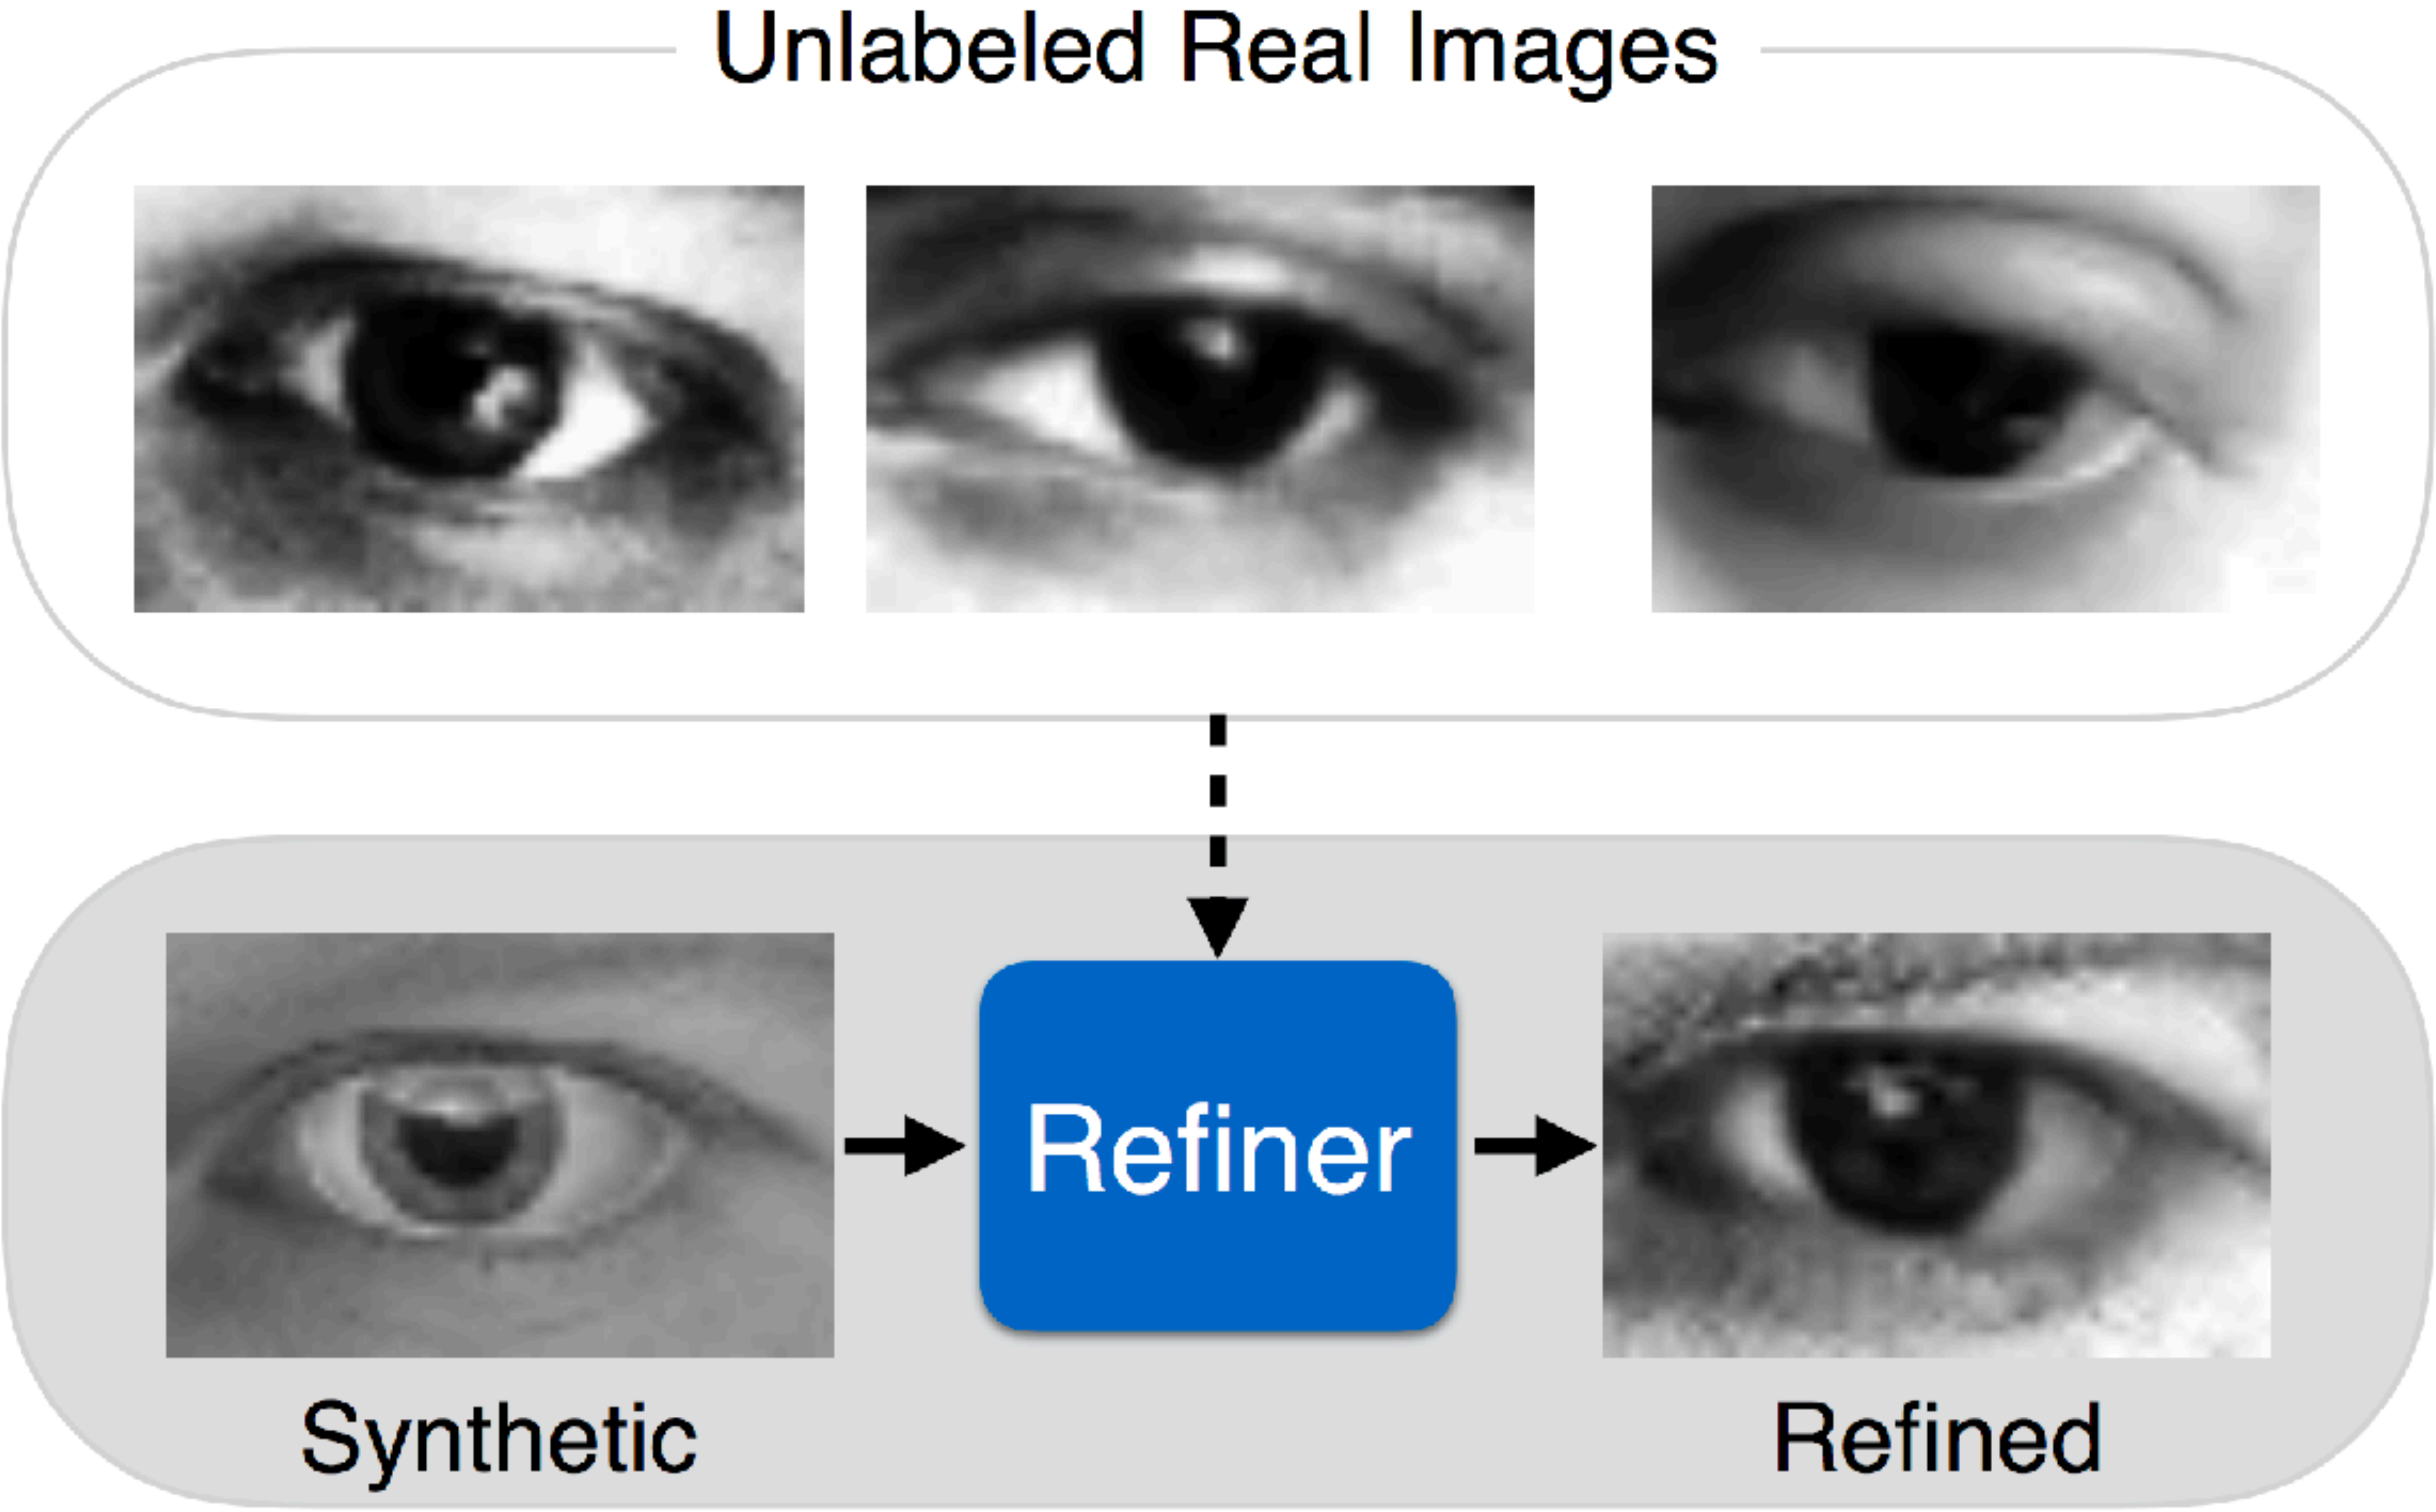

(Shrivastava et al., 2016)

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

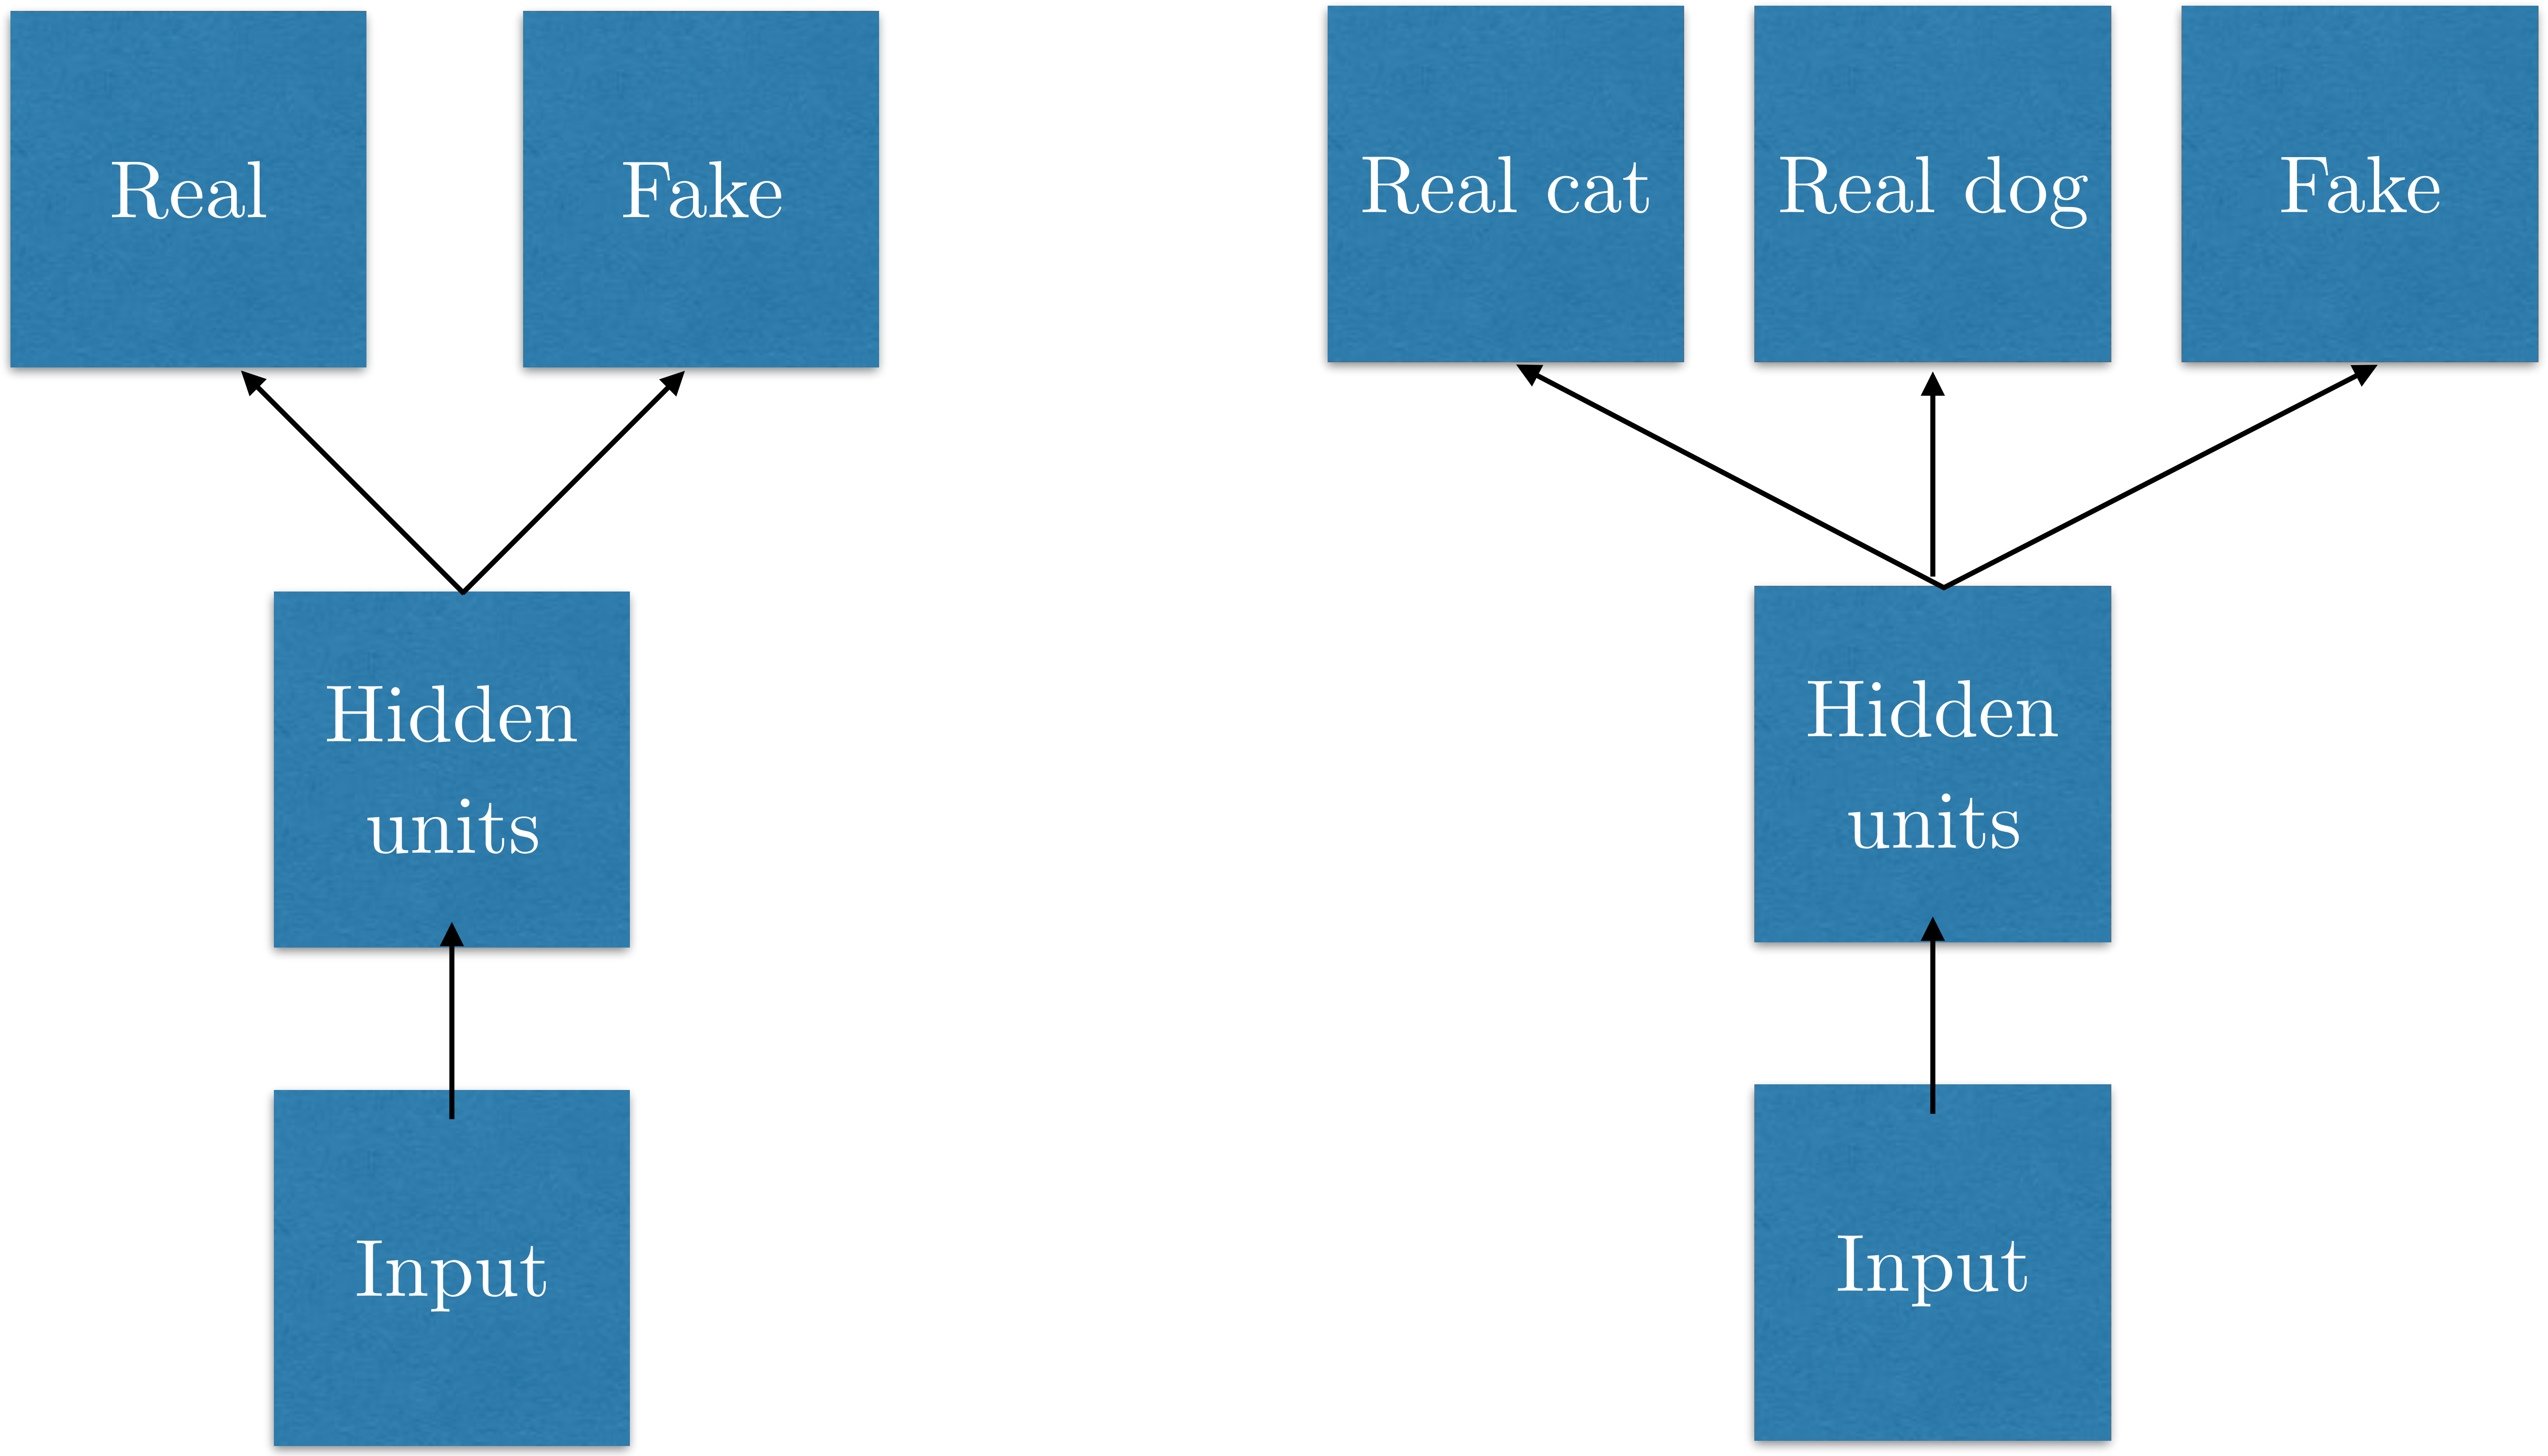

(Odena 2016, Salimans et al 2016)

Semi-Supervised Classification

MNIST (Permutation Invariant)

Model	Number of incorrectly predicted test examples for a given number of labeled samples			
	20	50	100	200
DGN [21]			333 ± 14	
Virtual Adversarial [22]			212	
CatGAN [14]			191 ± 10	
Skip Deep Generative Model [23]			132 ± 7	
Ladder network [24]			106 ± 37	
Auxiliary Deep Generative Model [23]			96 ± 2	
Our model	1677 ± 452	221 ± 136	93 ± 6.5	90 ± 4.2
Ensemble of 10 of our models	1134 ± 445	142 ± 96	86 ± 5.6	81 ± 4.3

(Salimans et al 2016)

Semi-Supervised Classification

CIFAR-10

Model	Test error rate for a given number of labeled samples			
	1000	2000	4000	8000
Ladder network [24]			20.40±0.47	
CatGAN [14]			19.58±0.46	
Our model	21.83±2.01	19.61±2.09	18.63±2.32	17.72±1.82
Ensemble of 10 of our models	19.22±0.54	17.25±0.66	15.59±0.47	14.87±0.89

SVHN

Model	Percentage of incorrectly predicted test examples for a given number of labeled samples		
	500	1000	2000
DGN [21]		36.02±0.10	
Virtual Adversarial [22]		24.63	
Auxiliary Deep Generative Model [23]		22.86	
Skip Deep Generative Model [23]		16.61±0.24	
Our model	18.44 ± 4.8	8.11 ± 1.3	6.16 ± 0.58
Ensemble of 10 of our models		5.88 ± 1.0	

(Salimans 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

Next Video Frame Prediction

Ground Truth

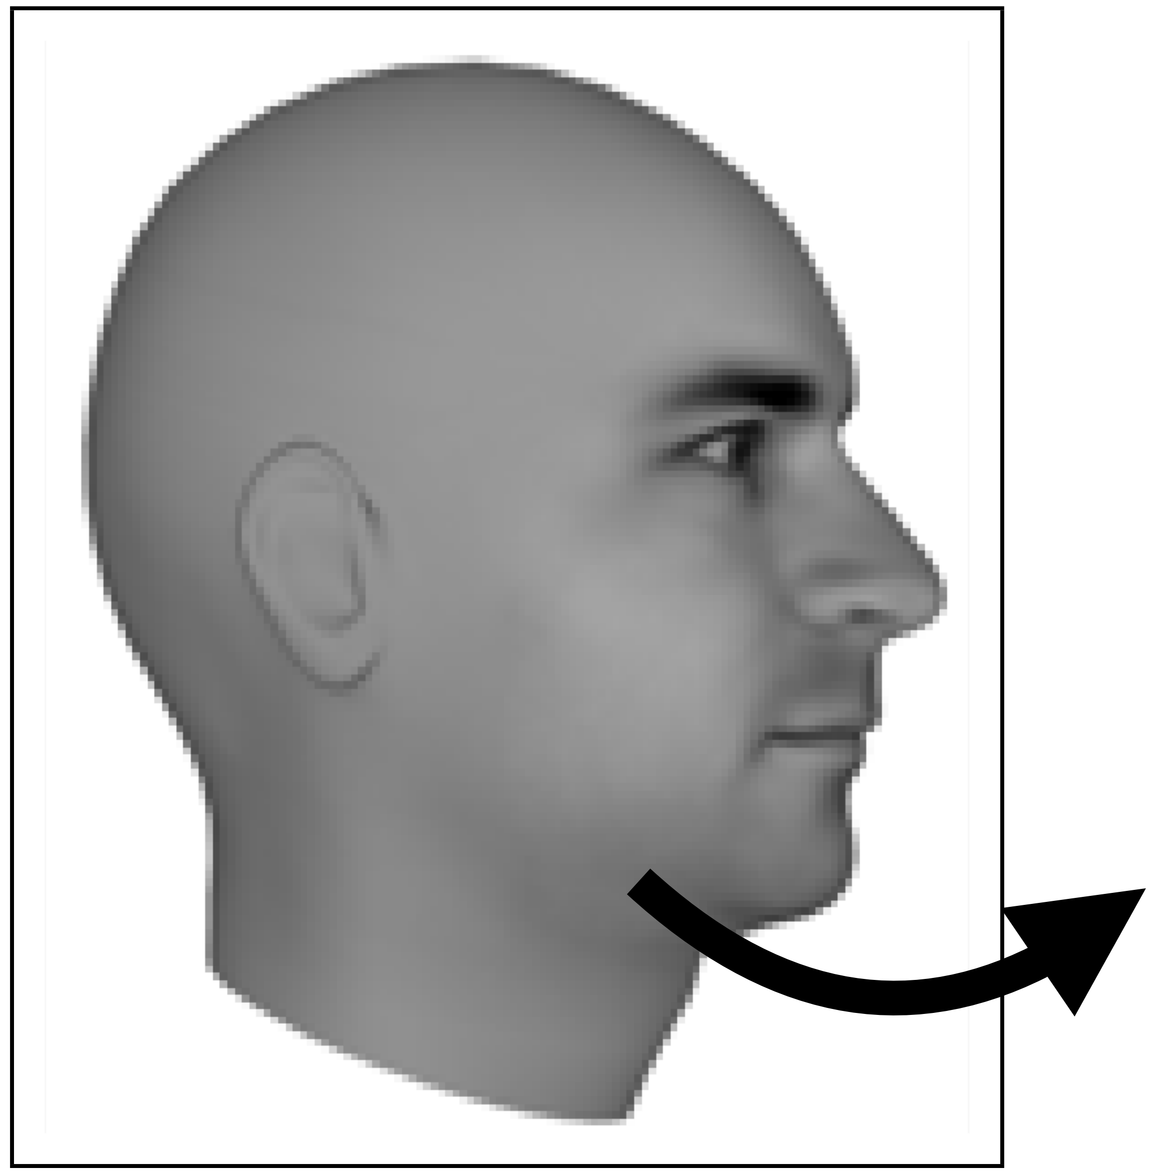

What happens next?

(Lotter et al 2016)

Next Video Frame Prediction

Ground Truth MSE Adversarial

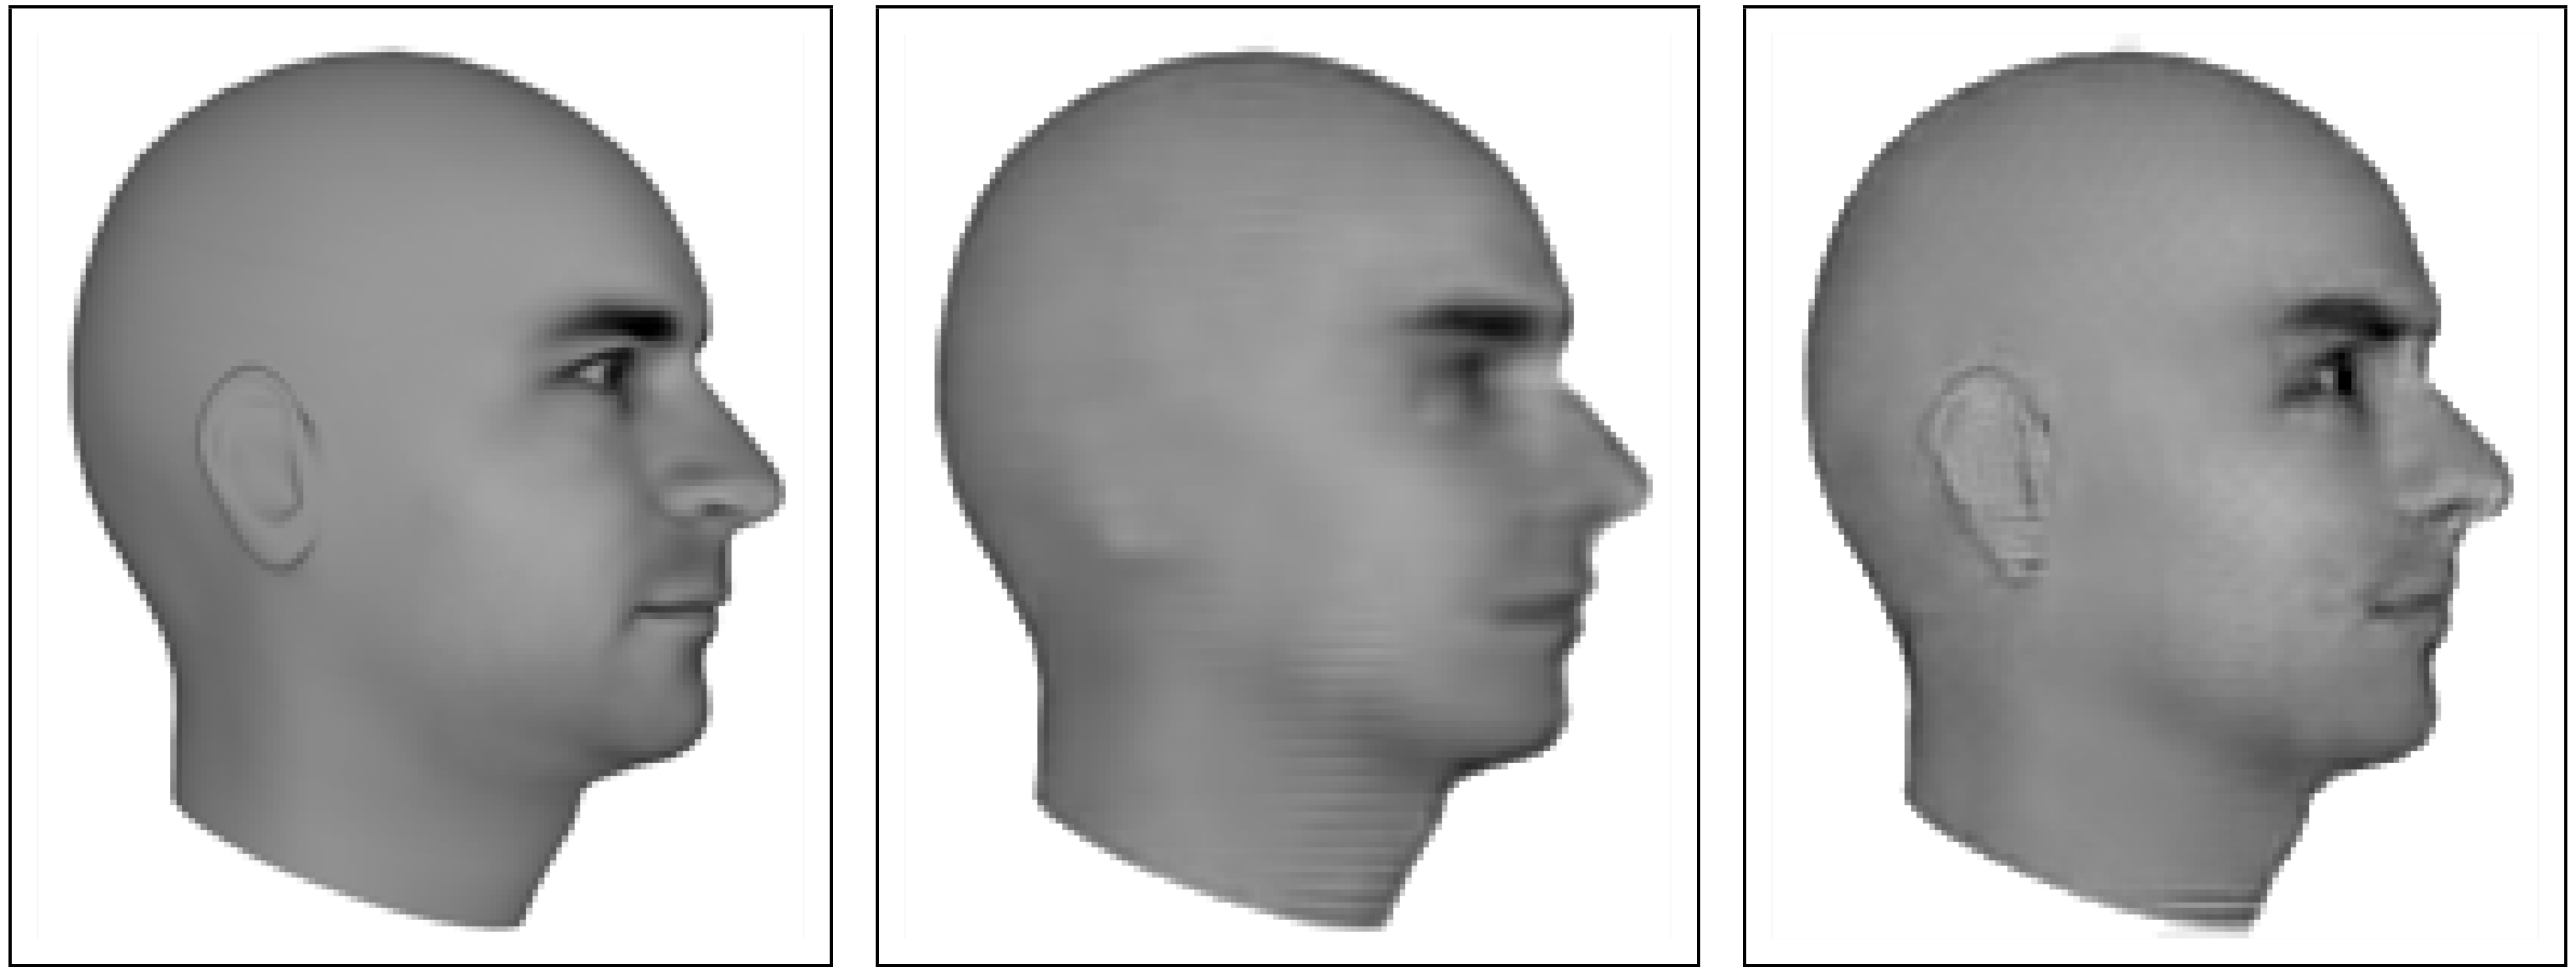

(Lotter et al 2016)

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

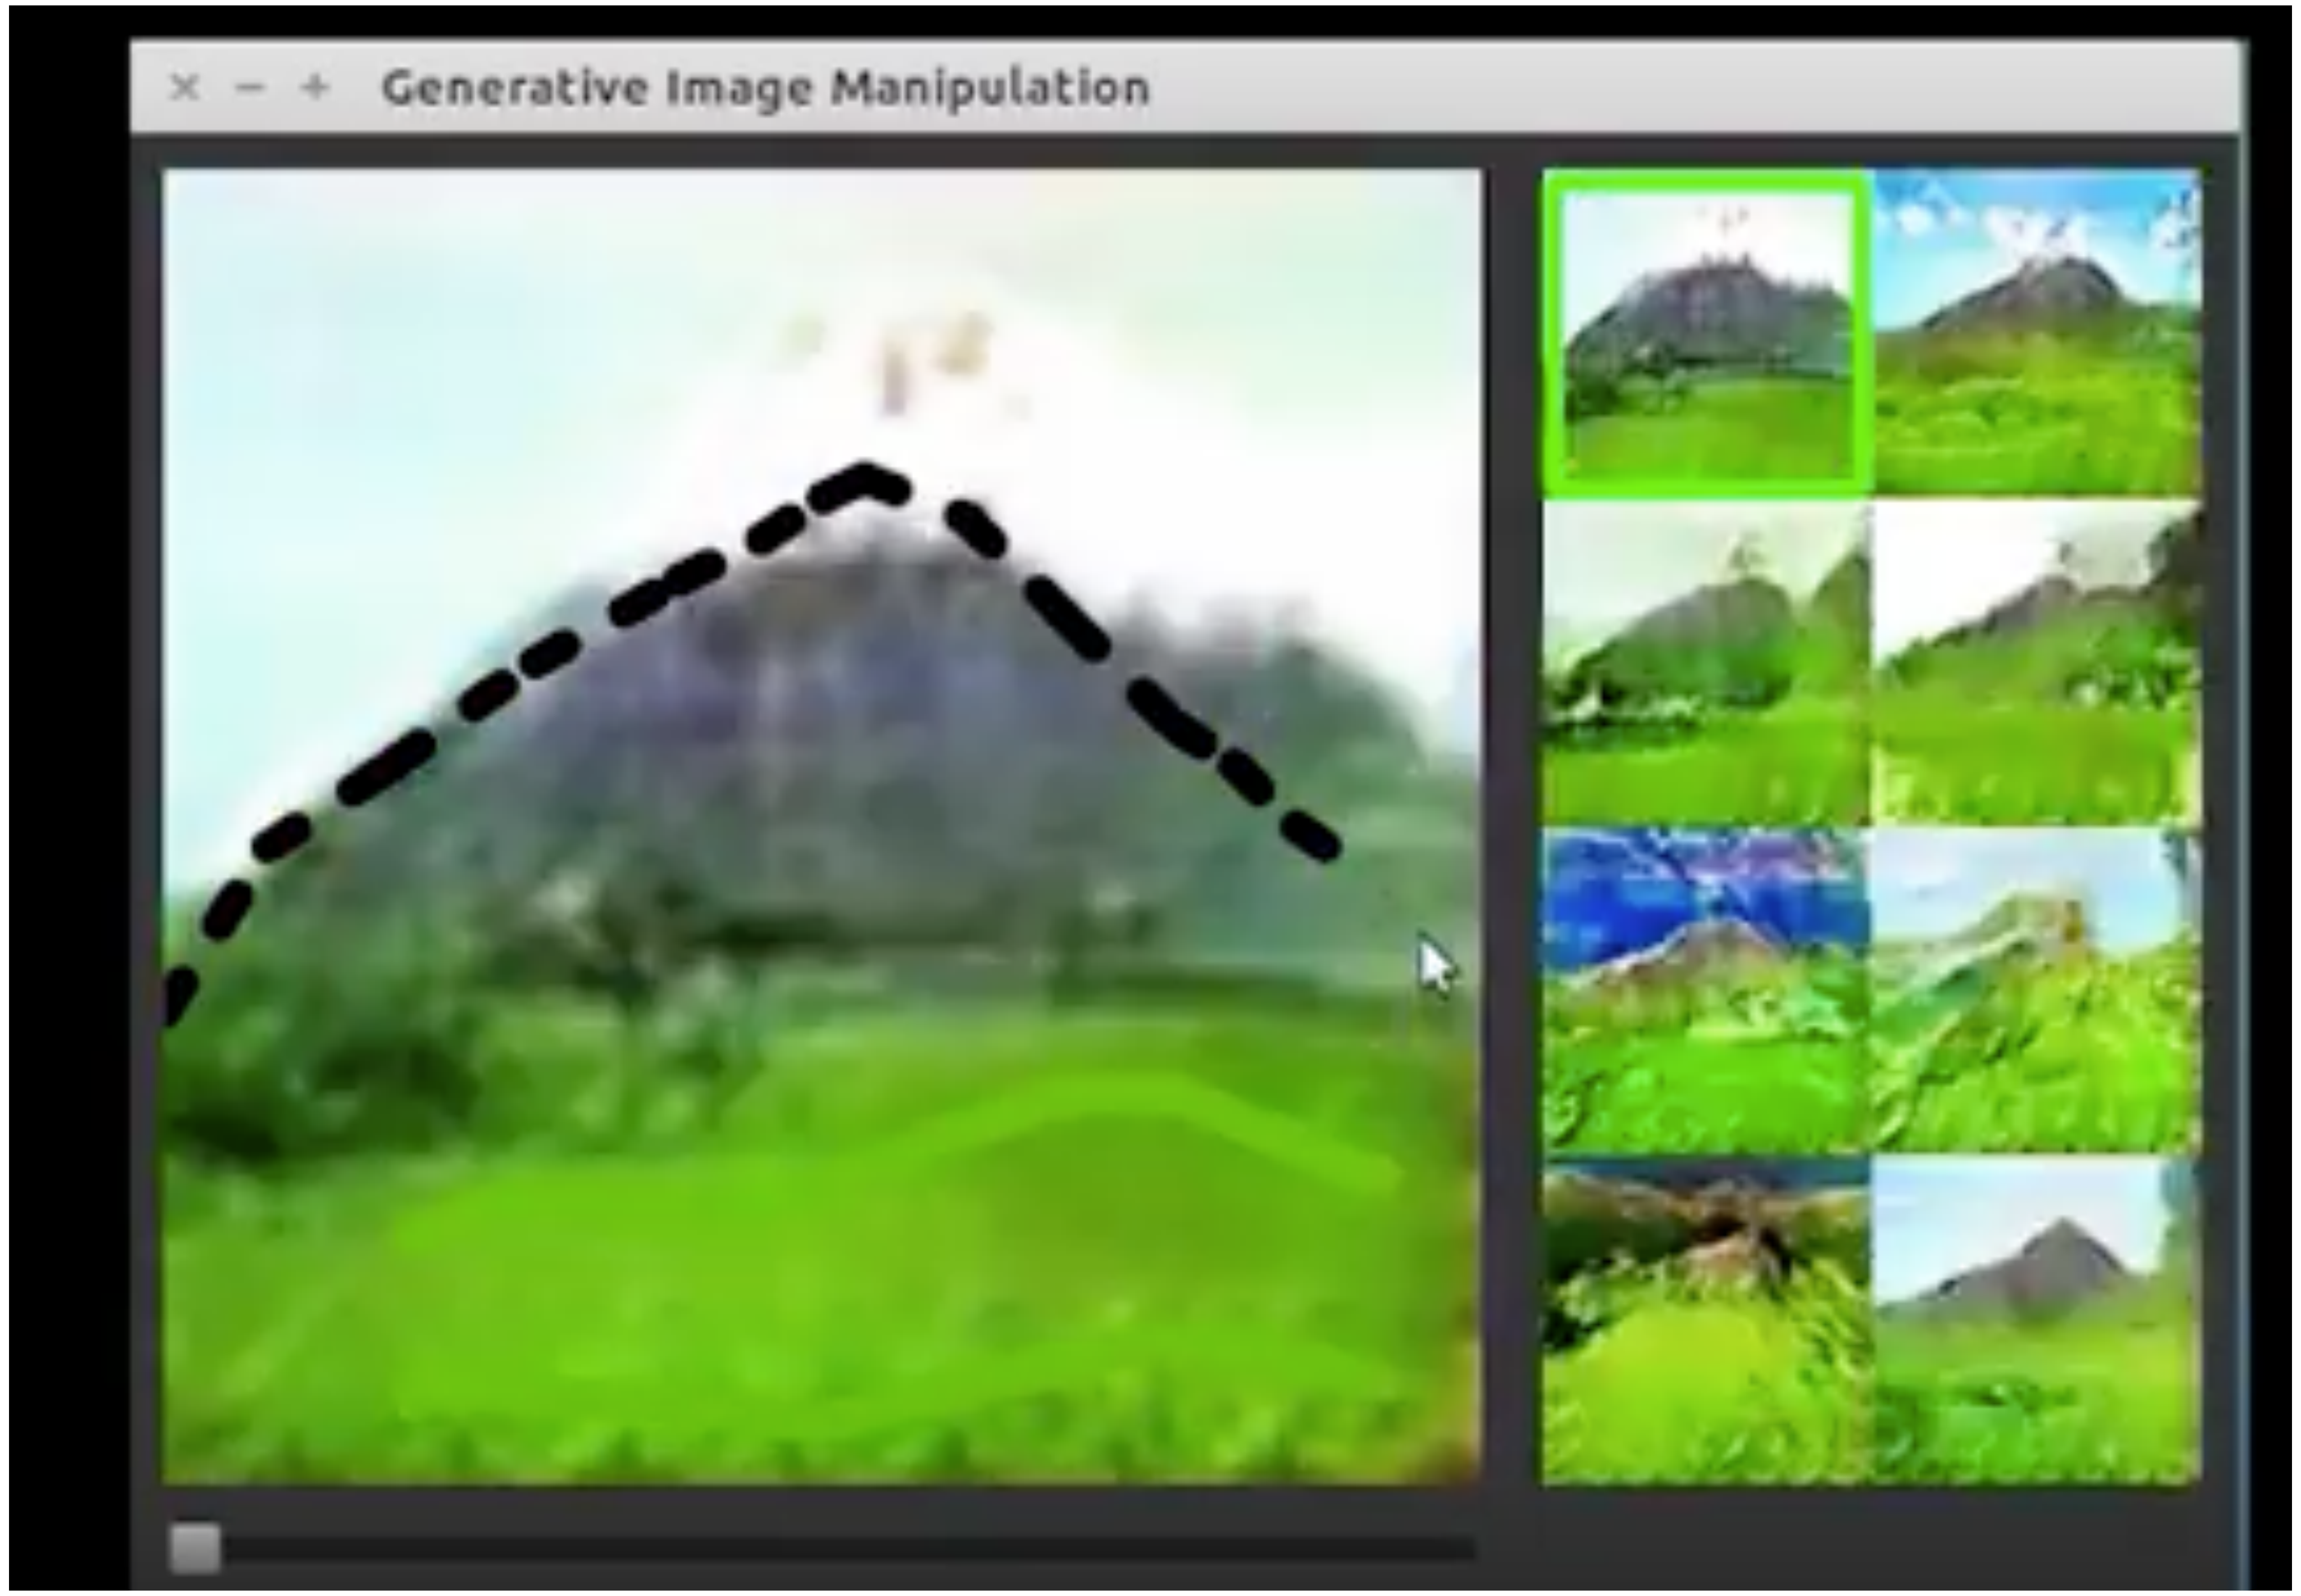

youtube

(Zhu et al., 2016)

Introspective Adversarial Networks

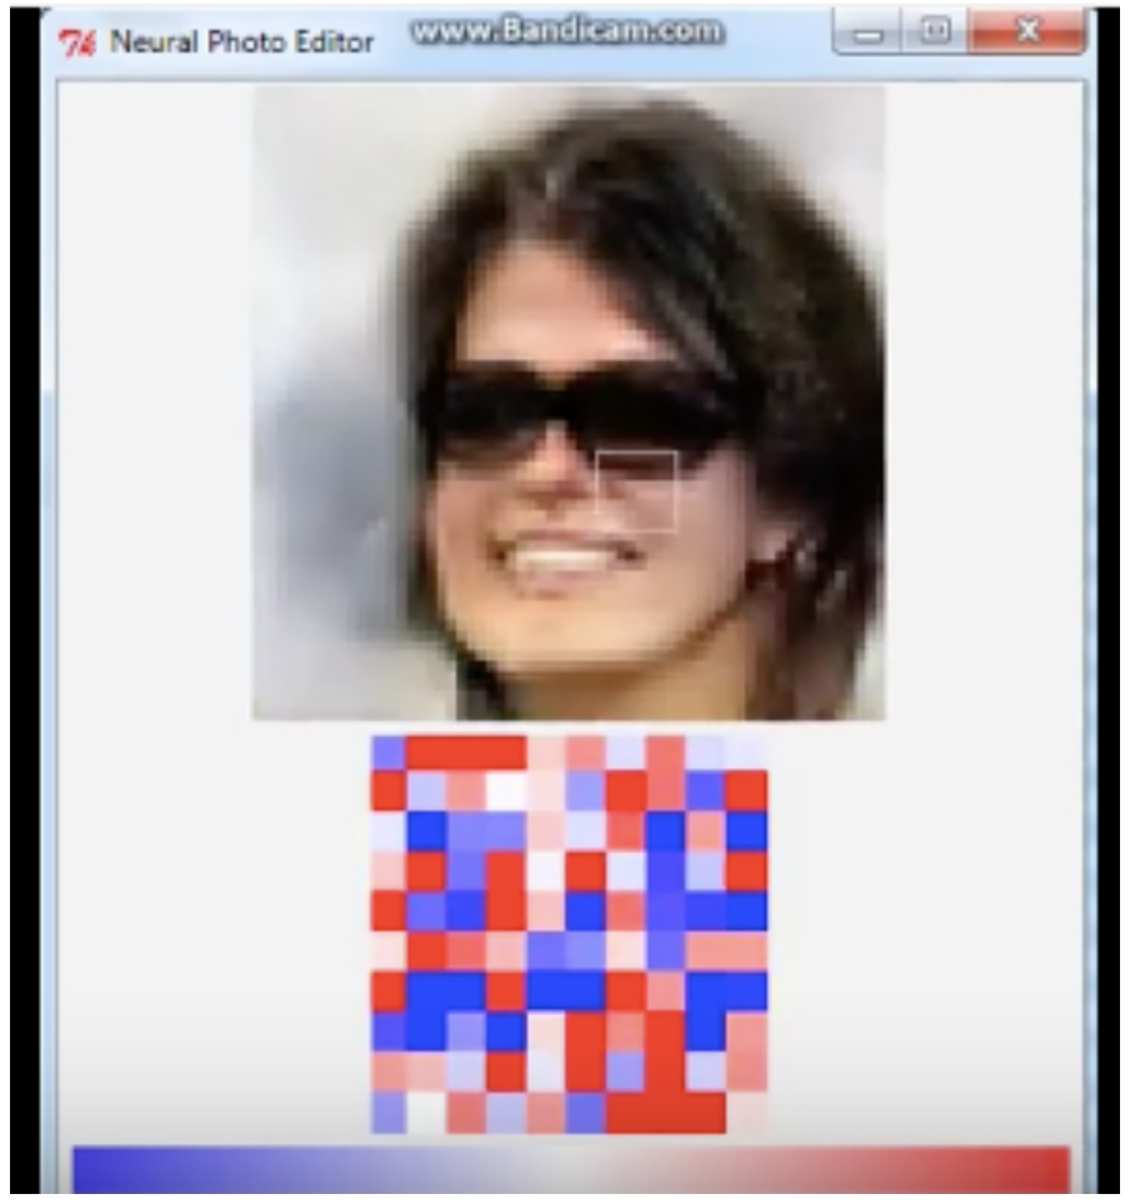

youtube

(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

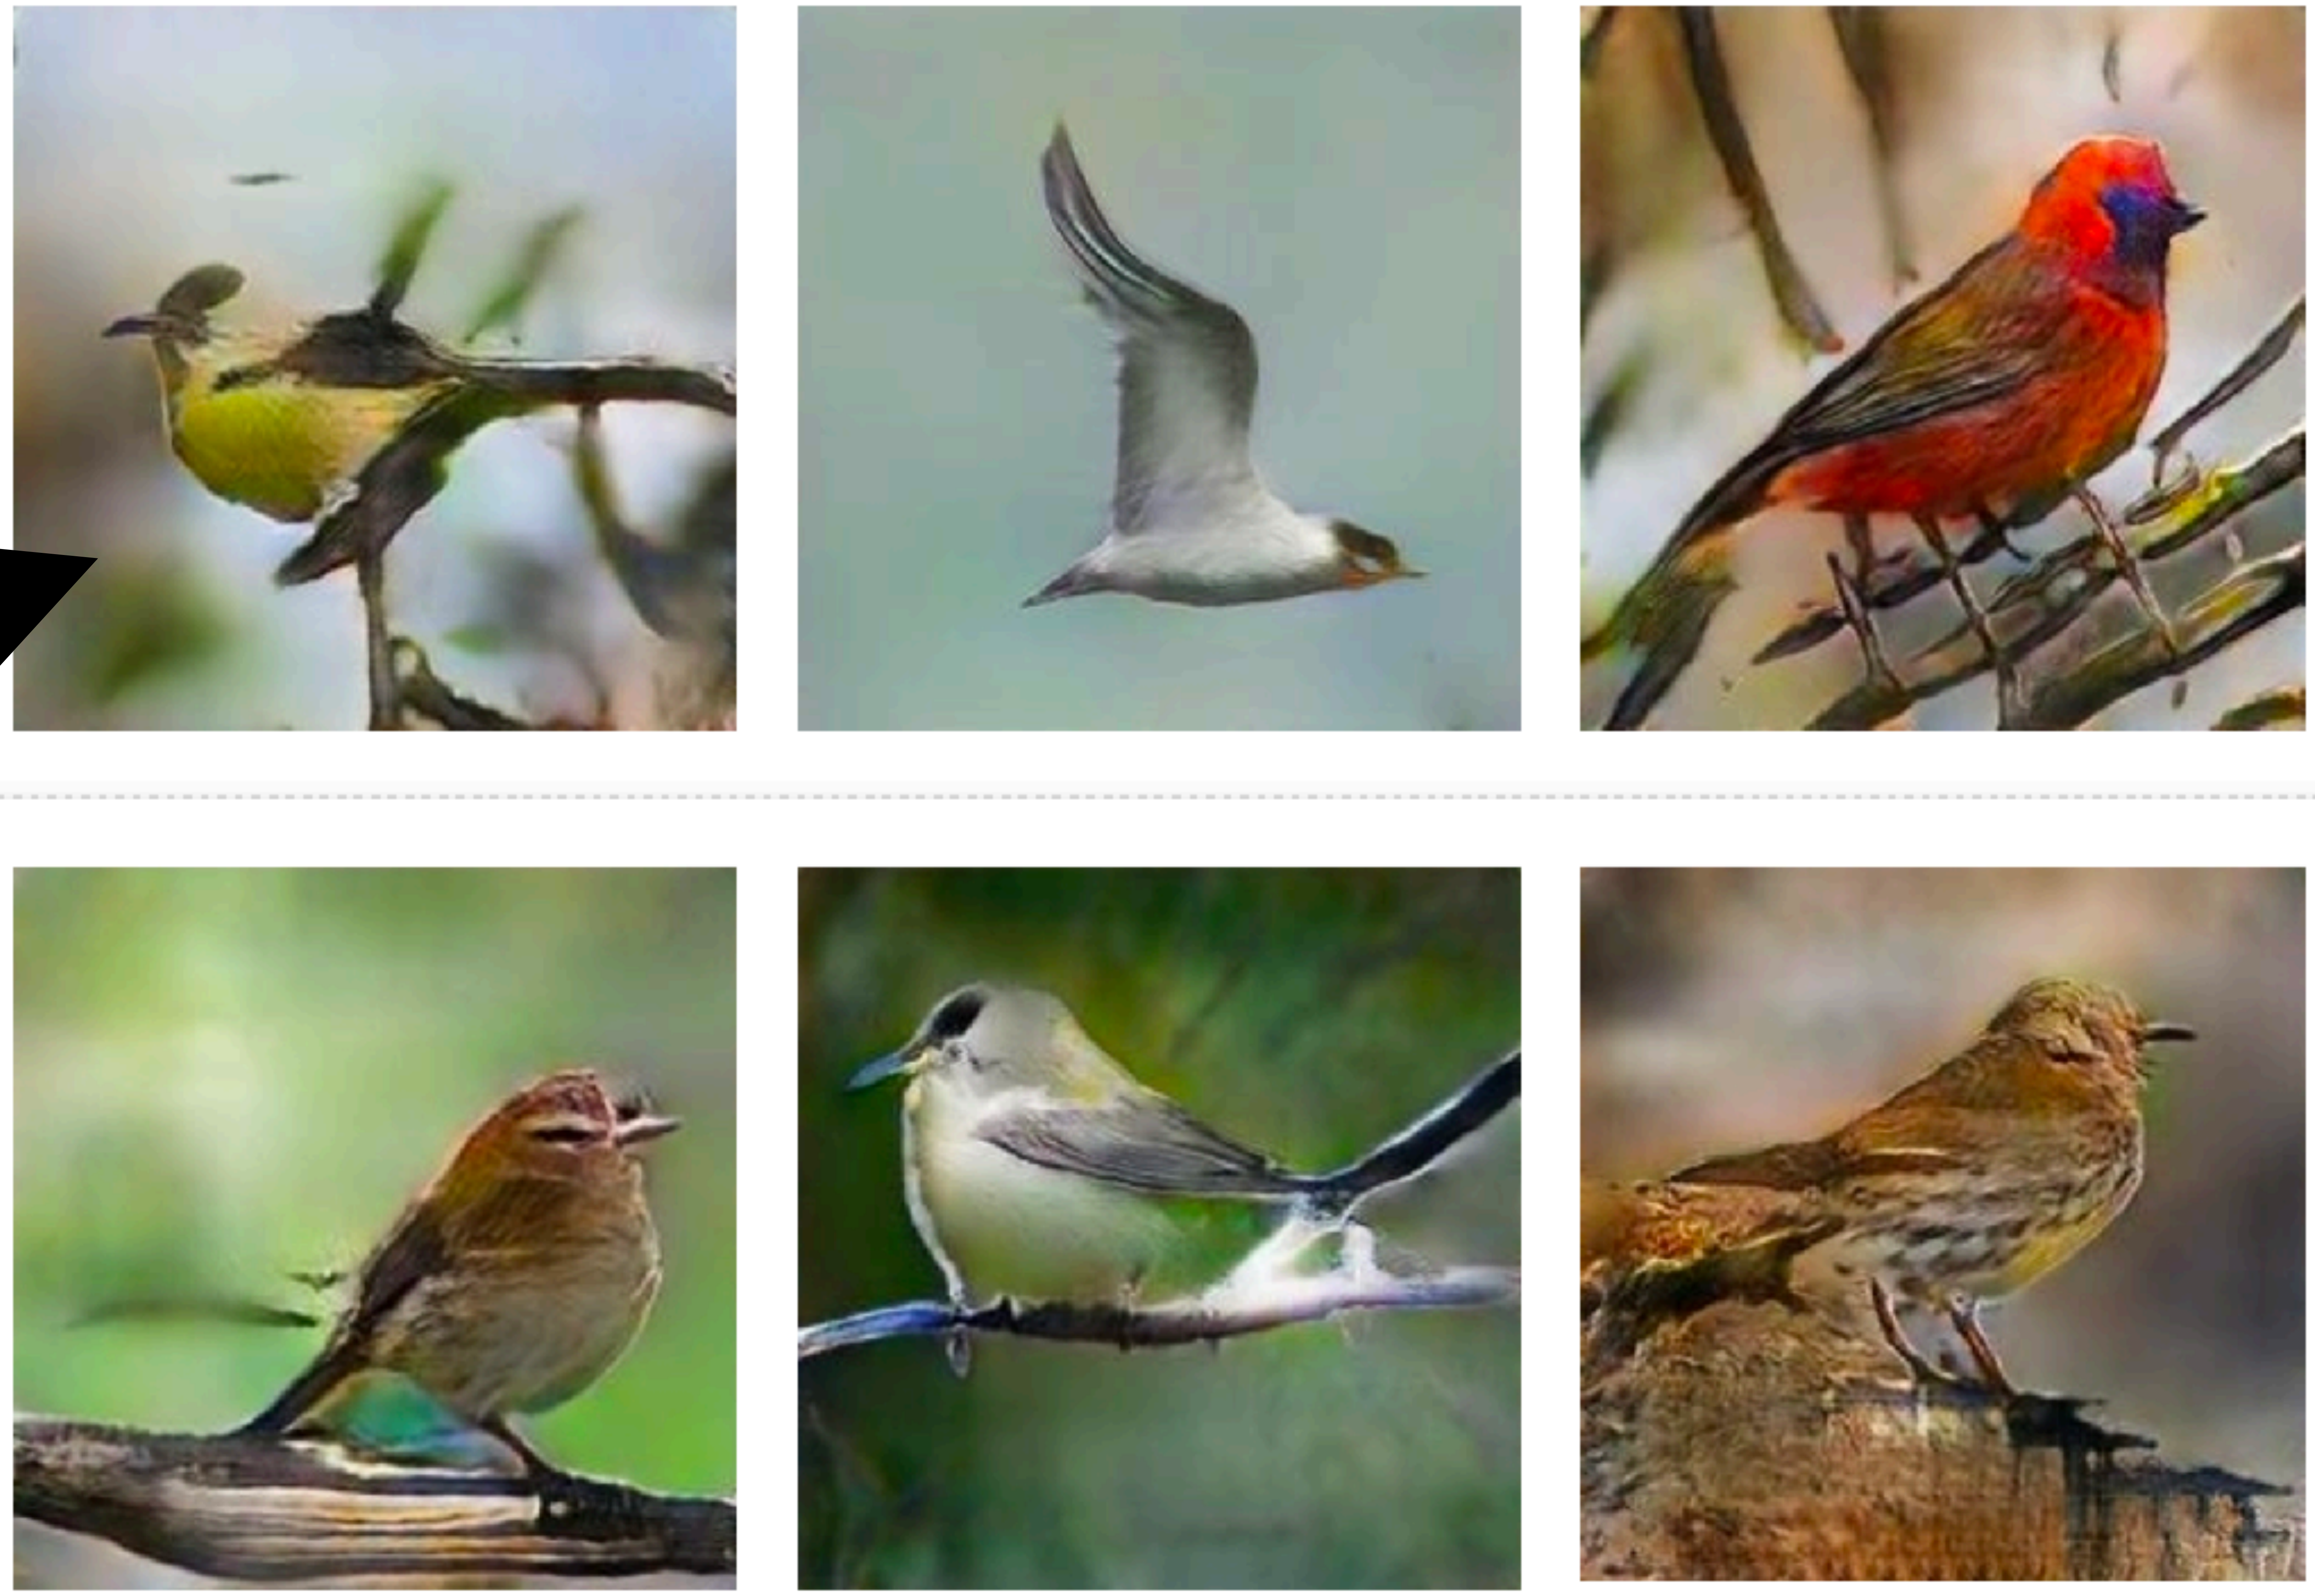

(Zhang et al., 2016)

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

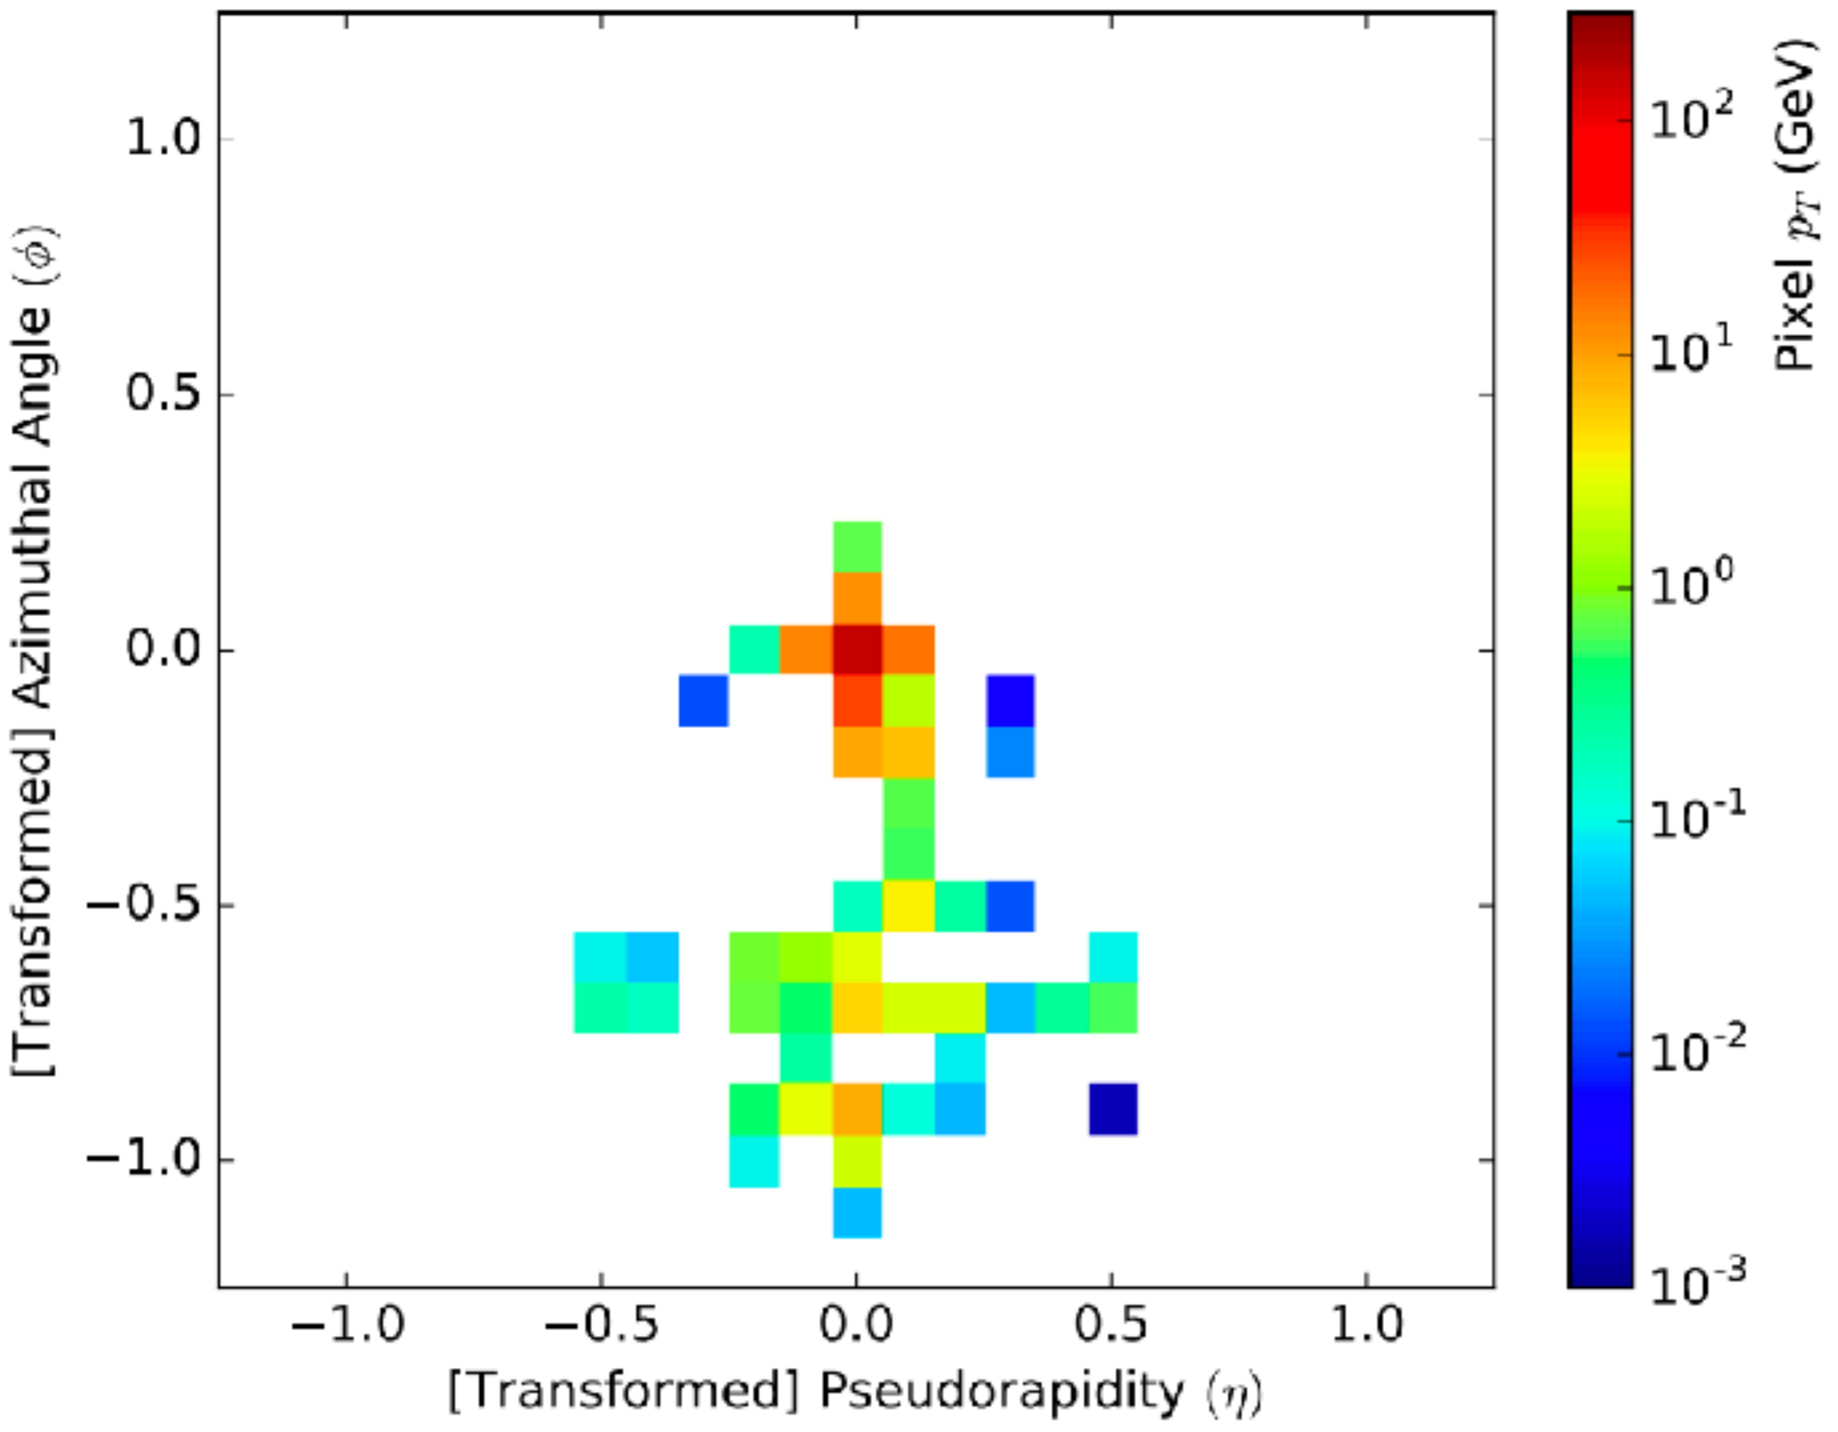

(de Oliveira et al., 2017)

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

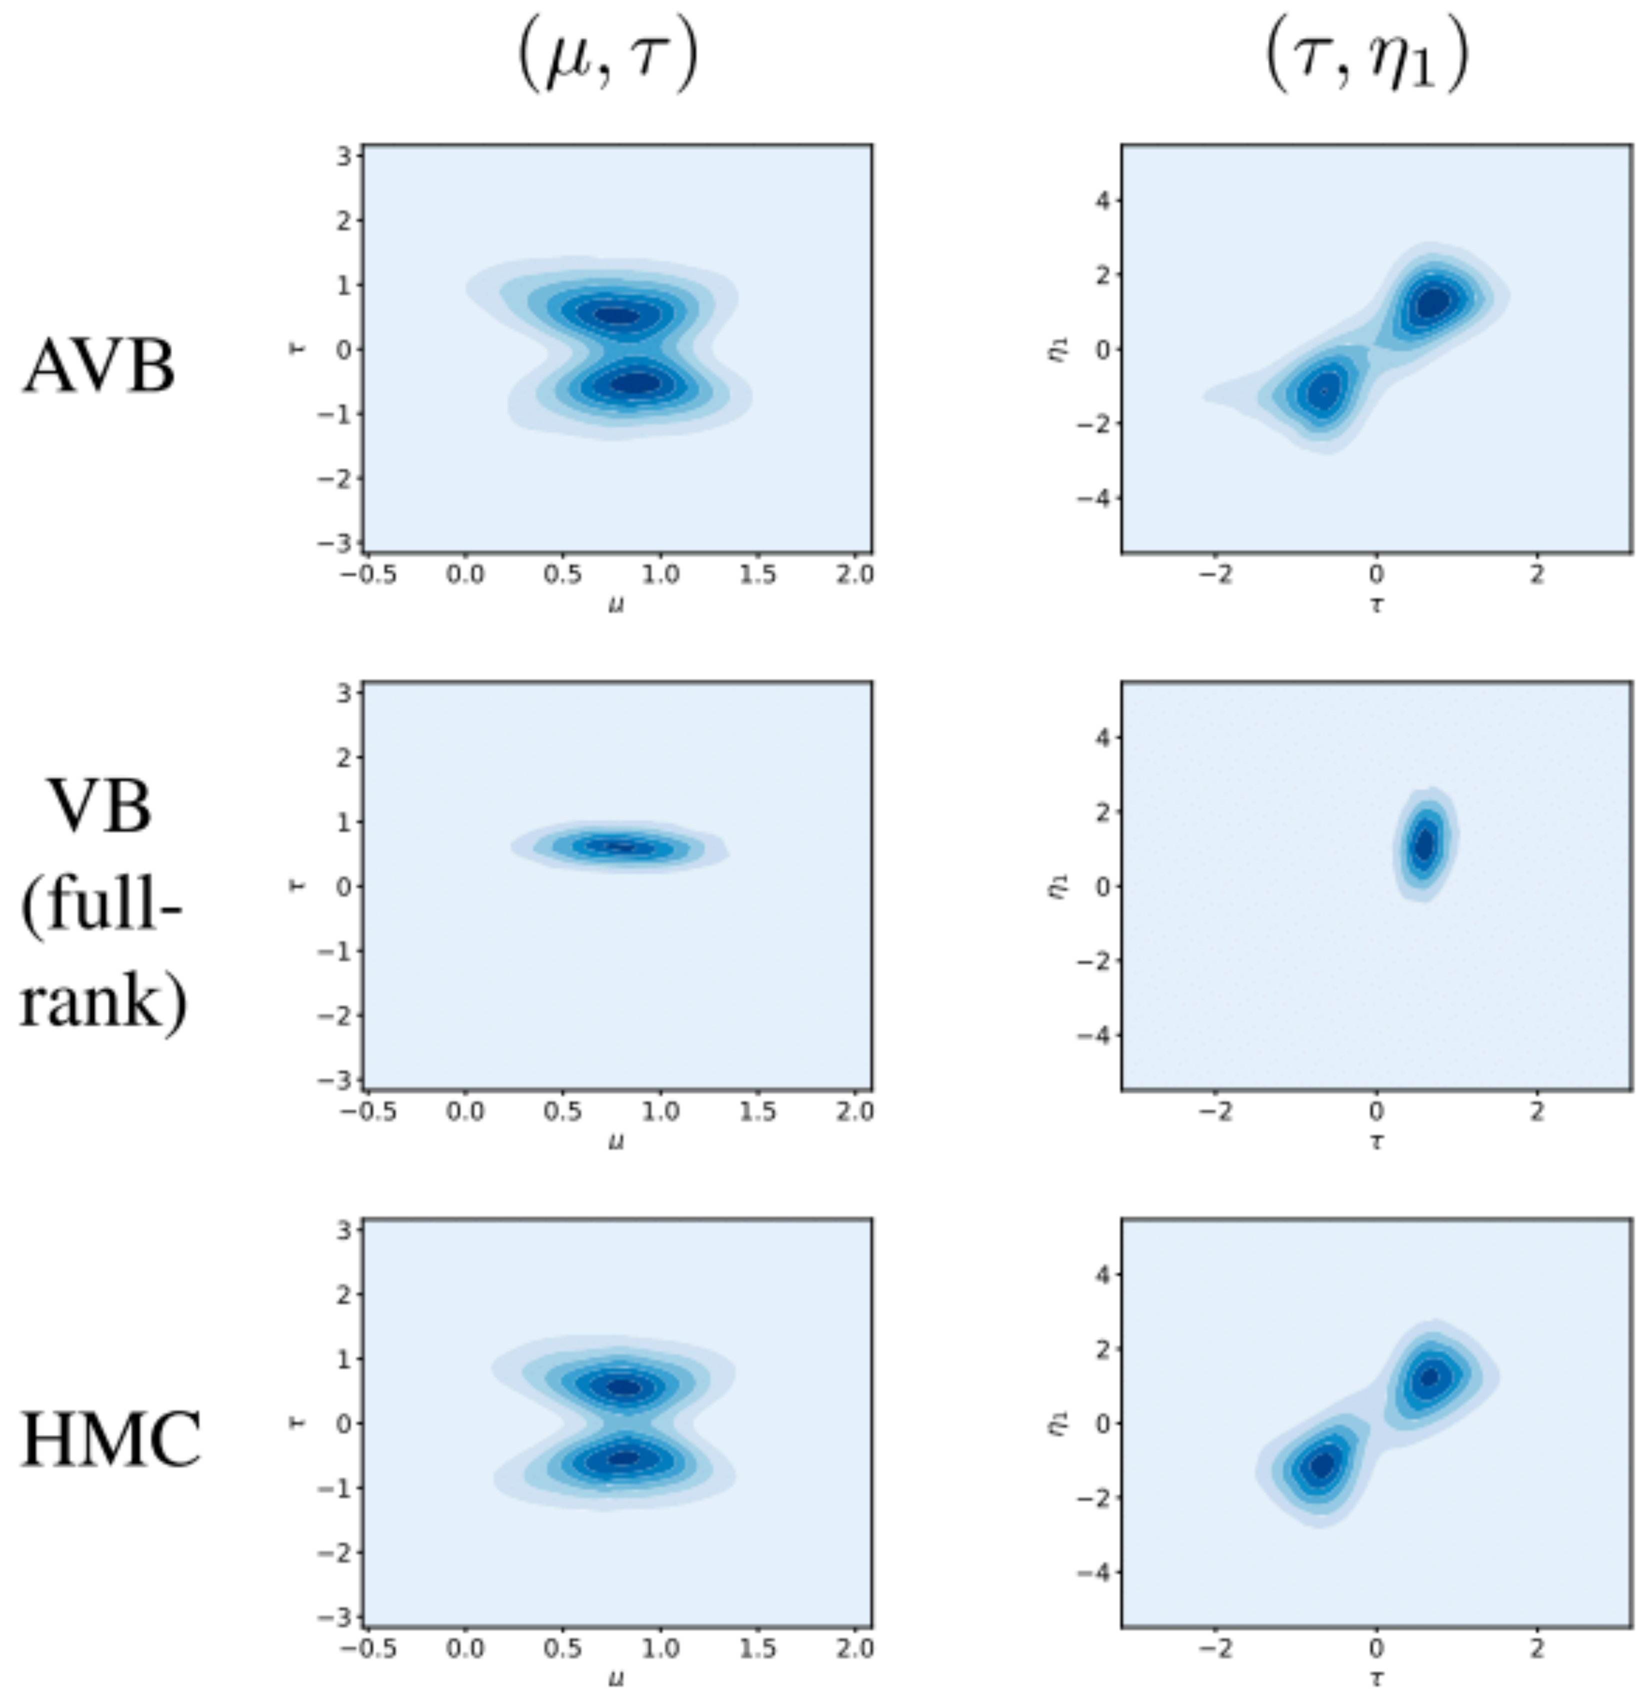

(Mescheder et al, 2017)

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

 - 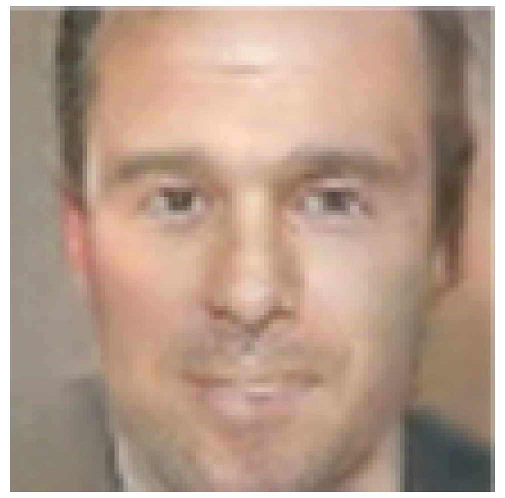 + 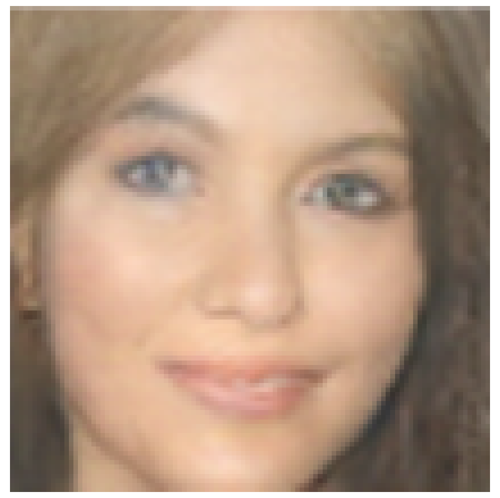 = 

Man with glasses | Man | Woman | Woman with Glasses

(Radford et al, 2015)

Learning interpretable latent codes / controlling the generation process

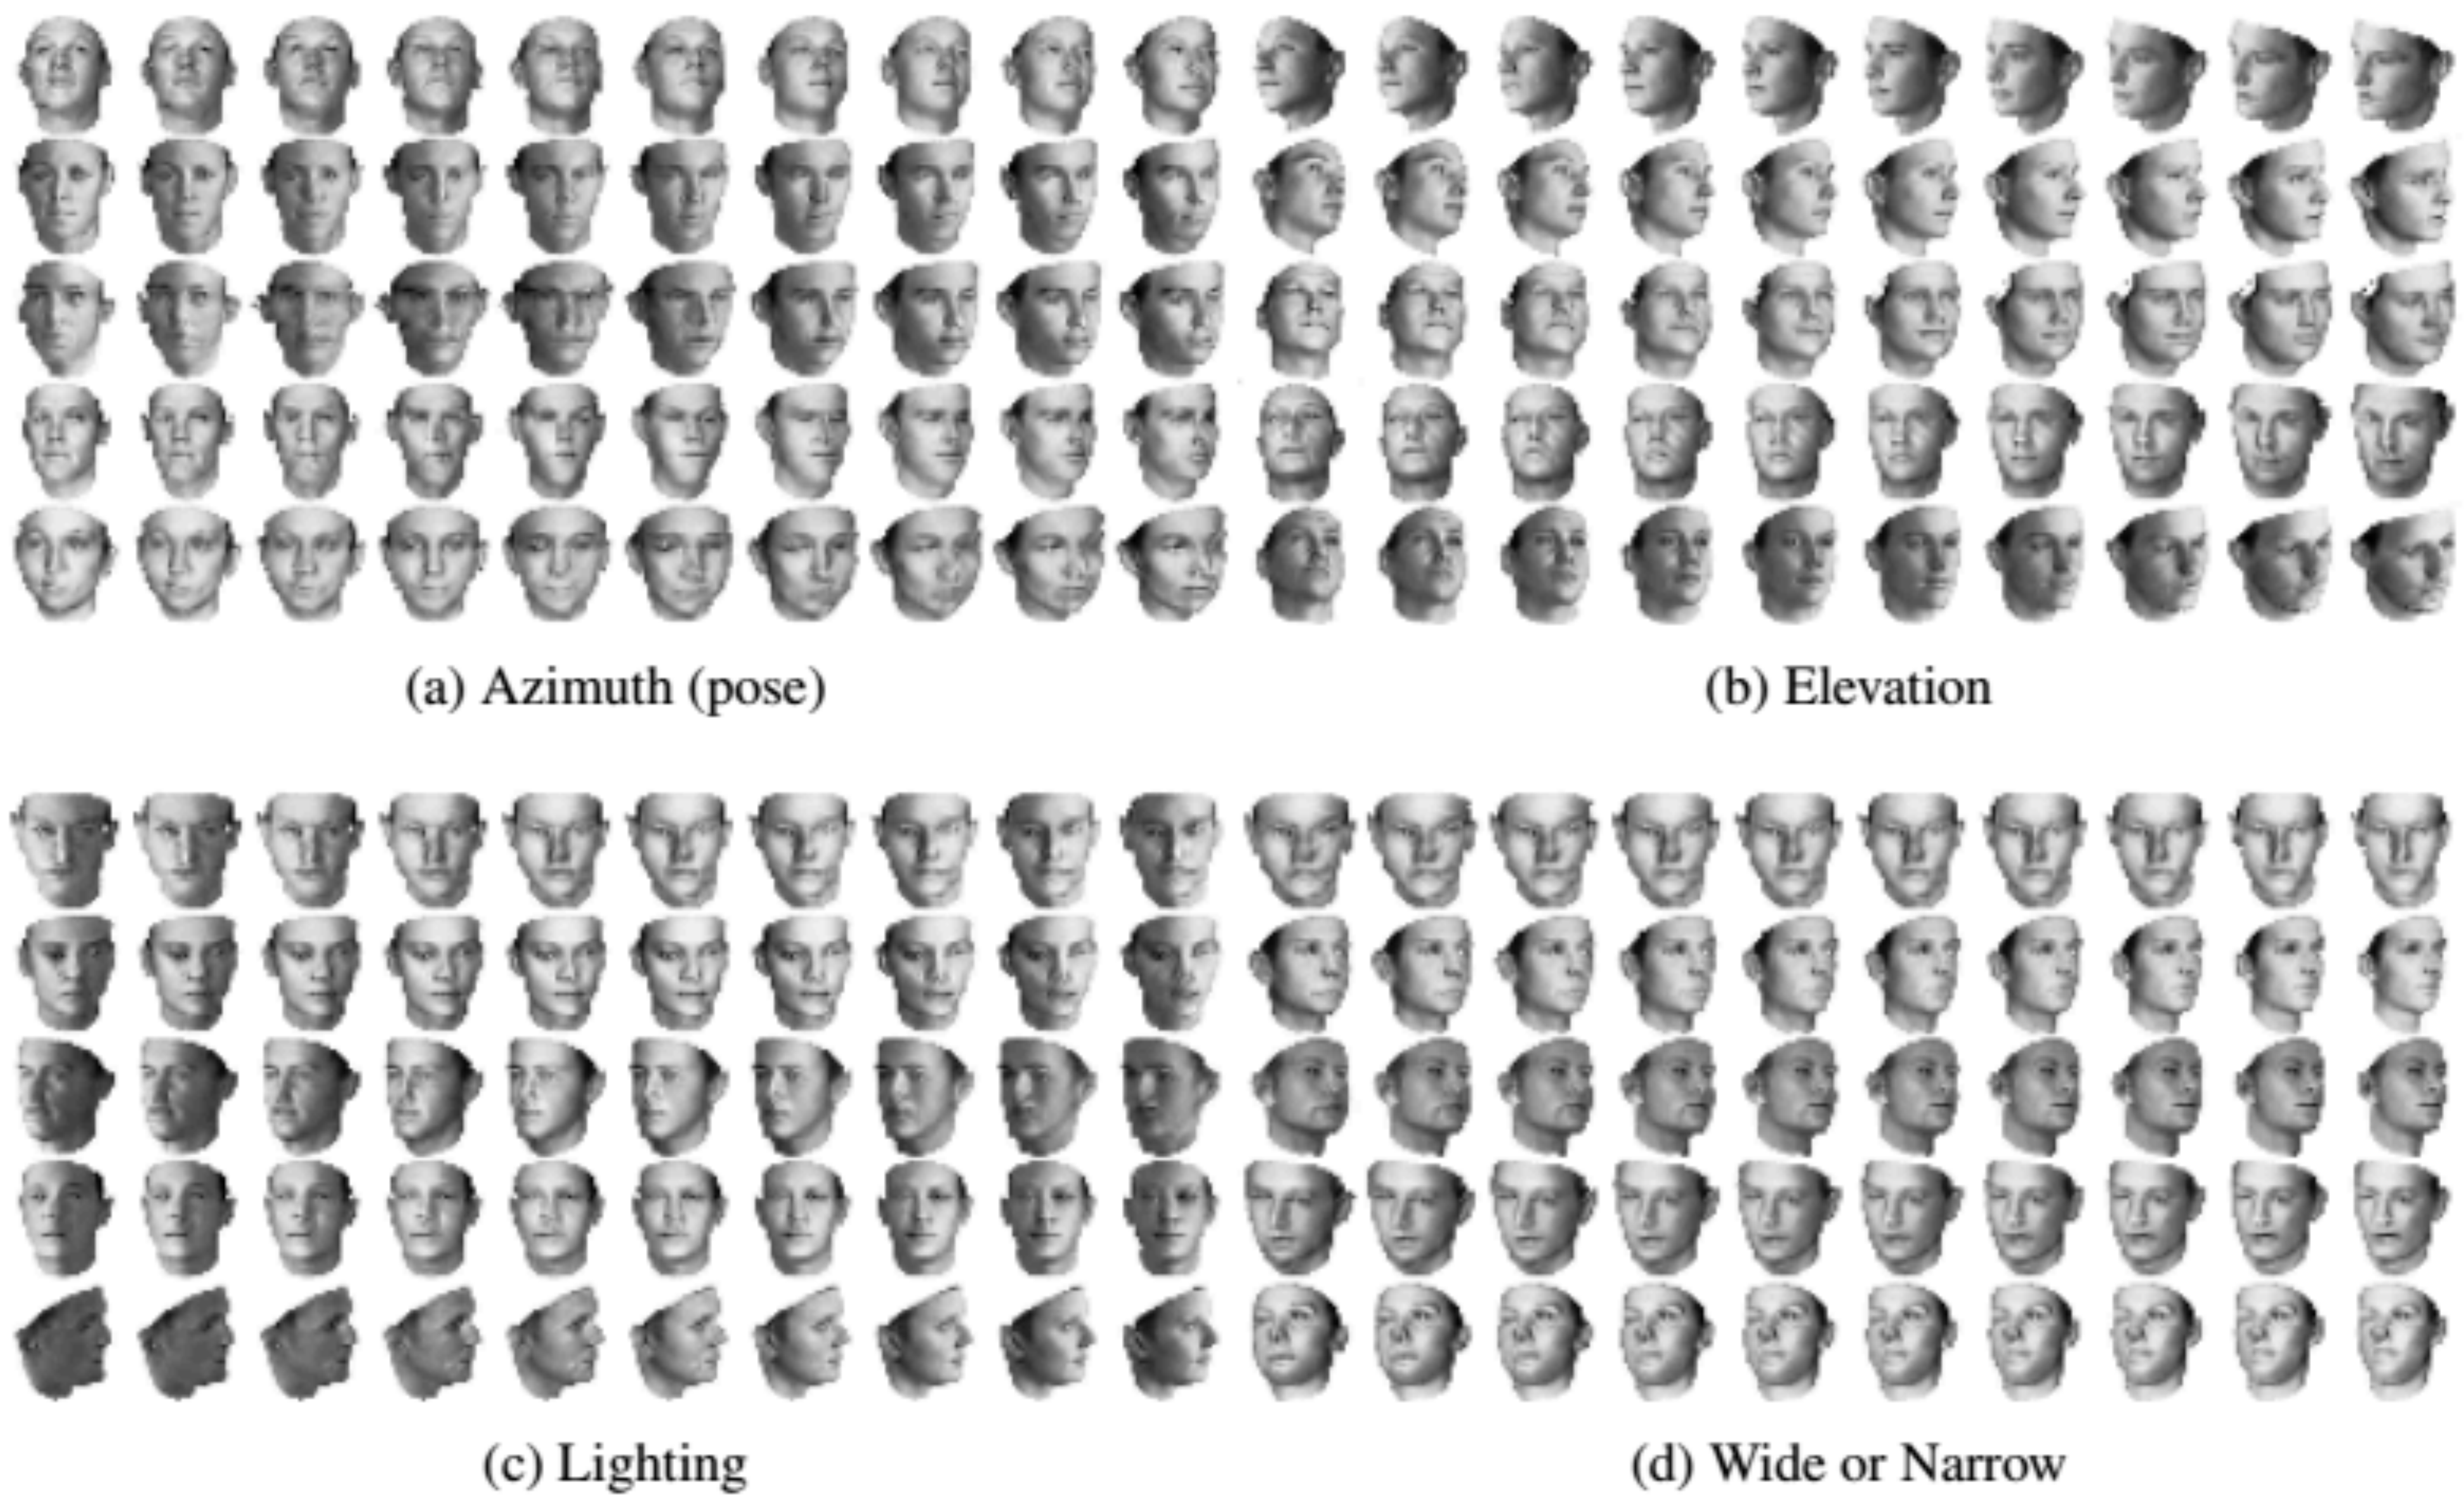

InfoGAN (Chen et al 2016)

How long until GANs can do this?

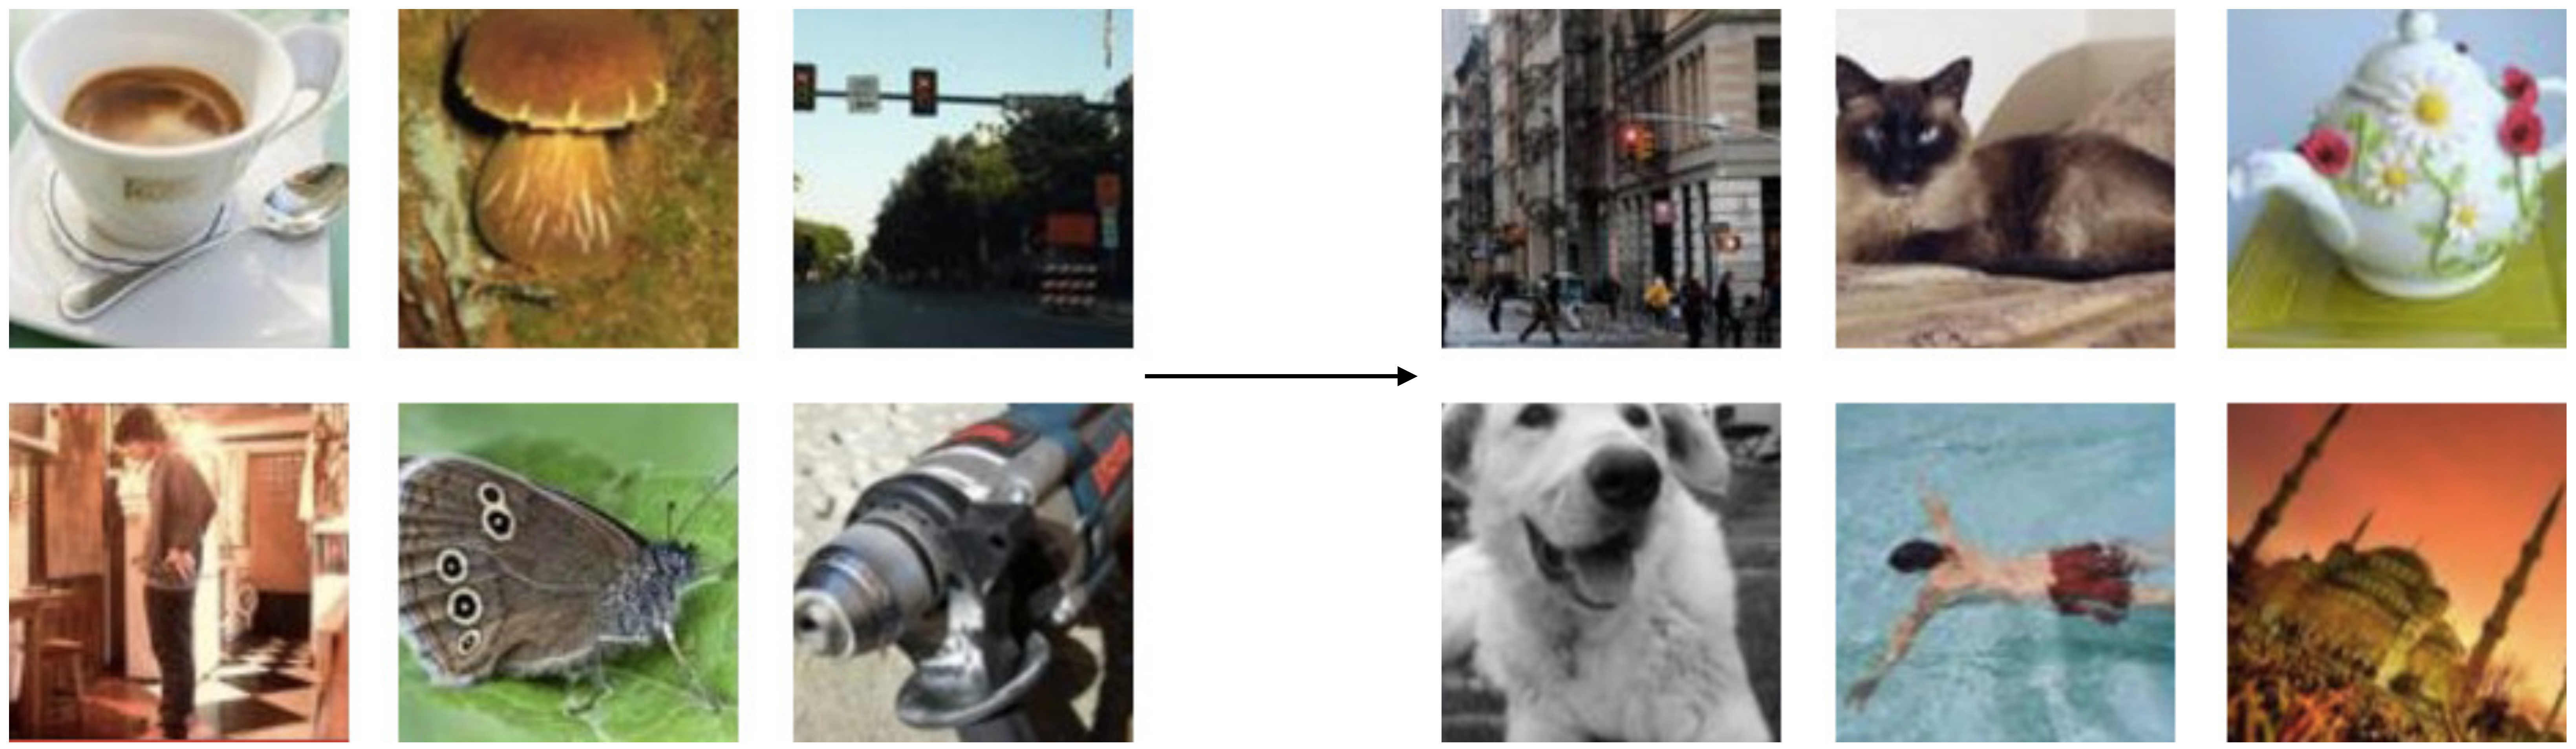

Training examples Model samples

(Goodfellow 2017)

AC-GANs

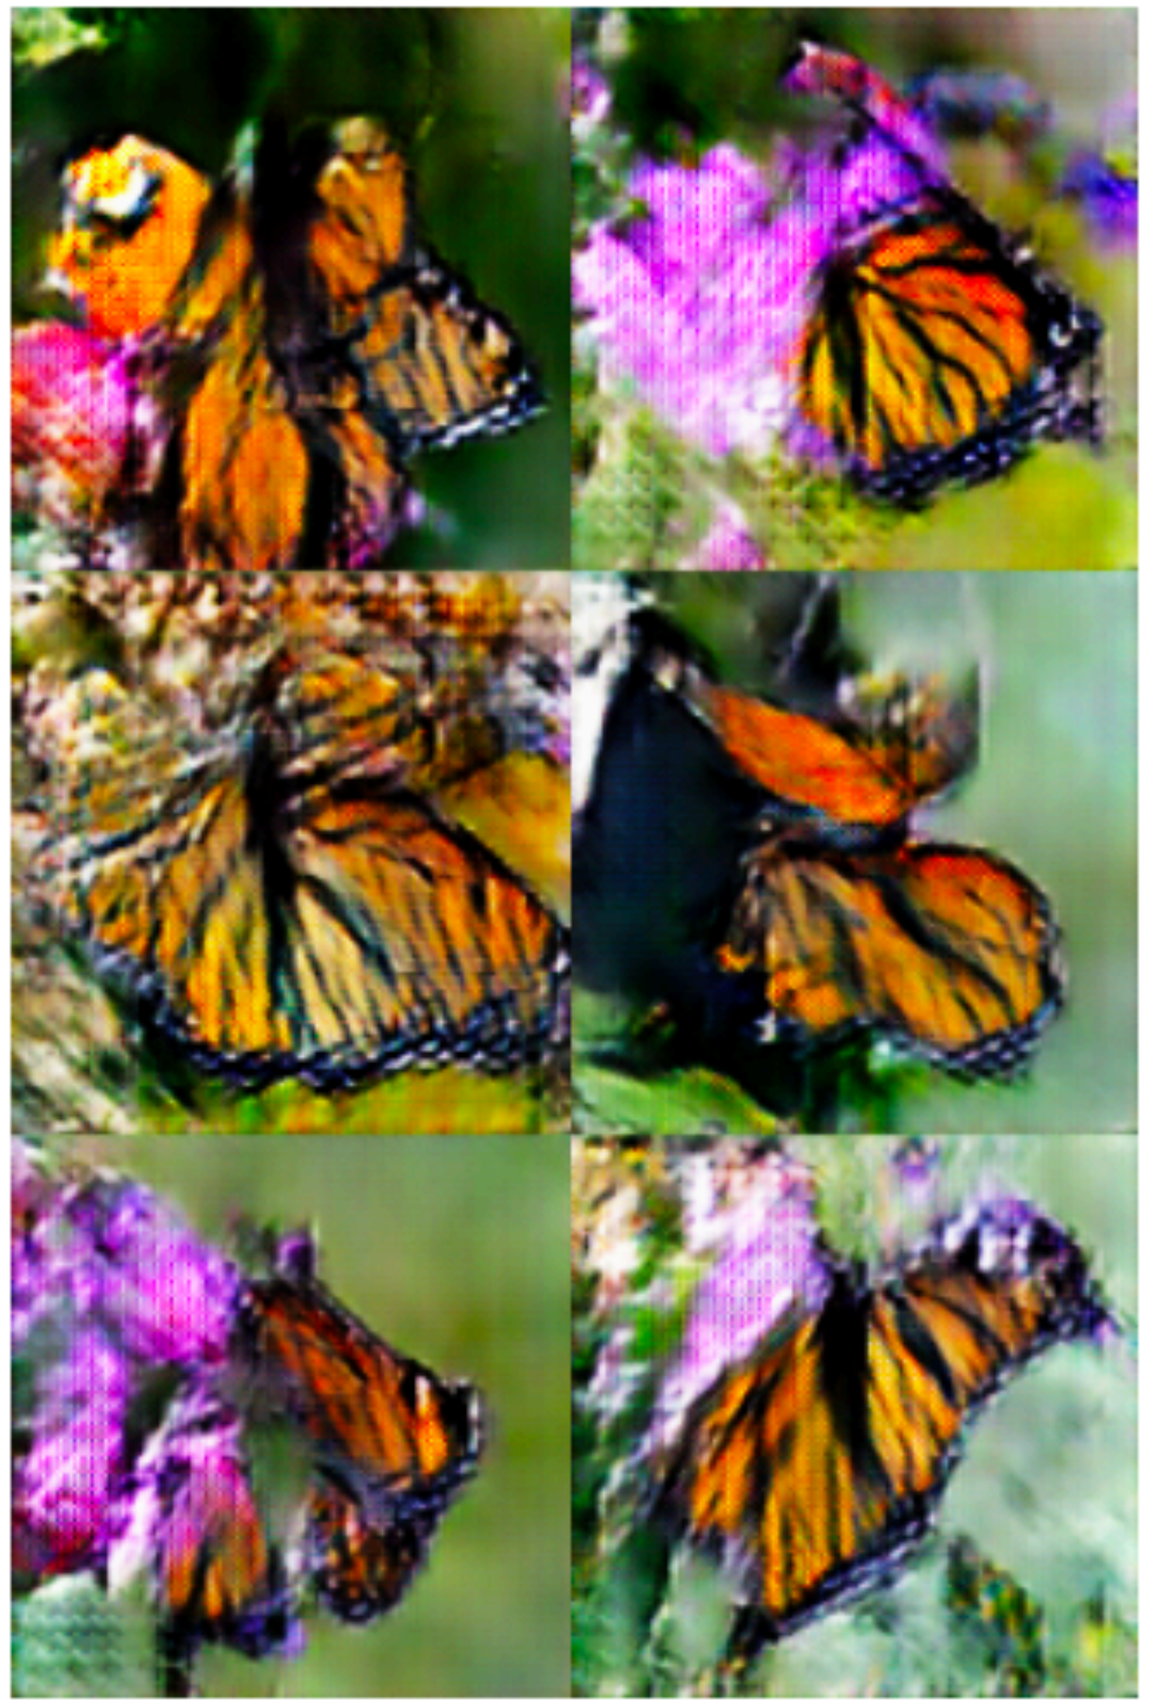

monarch butterfly

goldfinch

daisy

(Odena et al., 2016)

Minibatch GAN on ImageNet

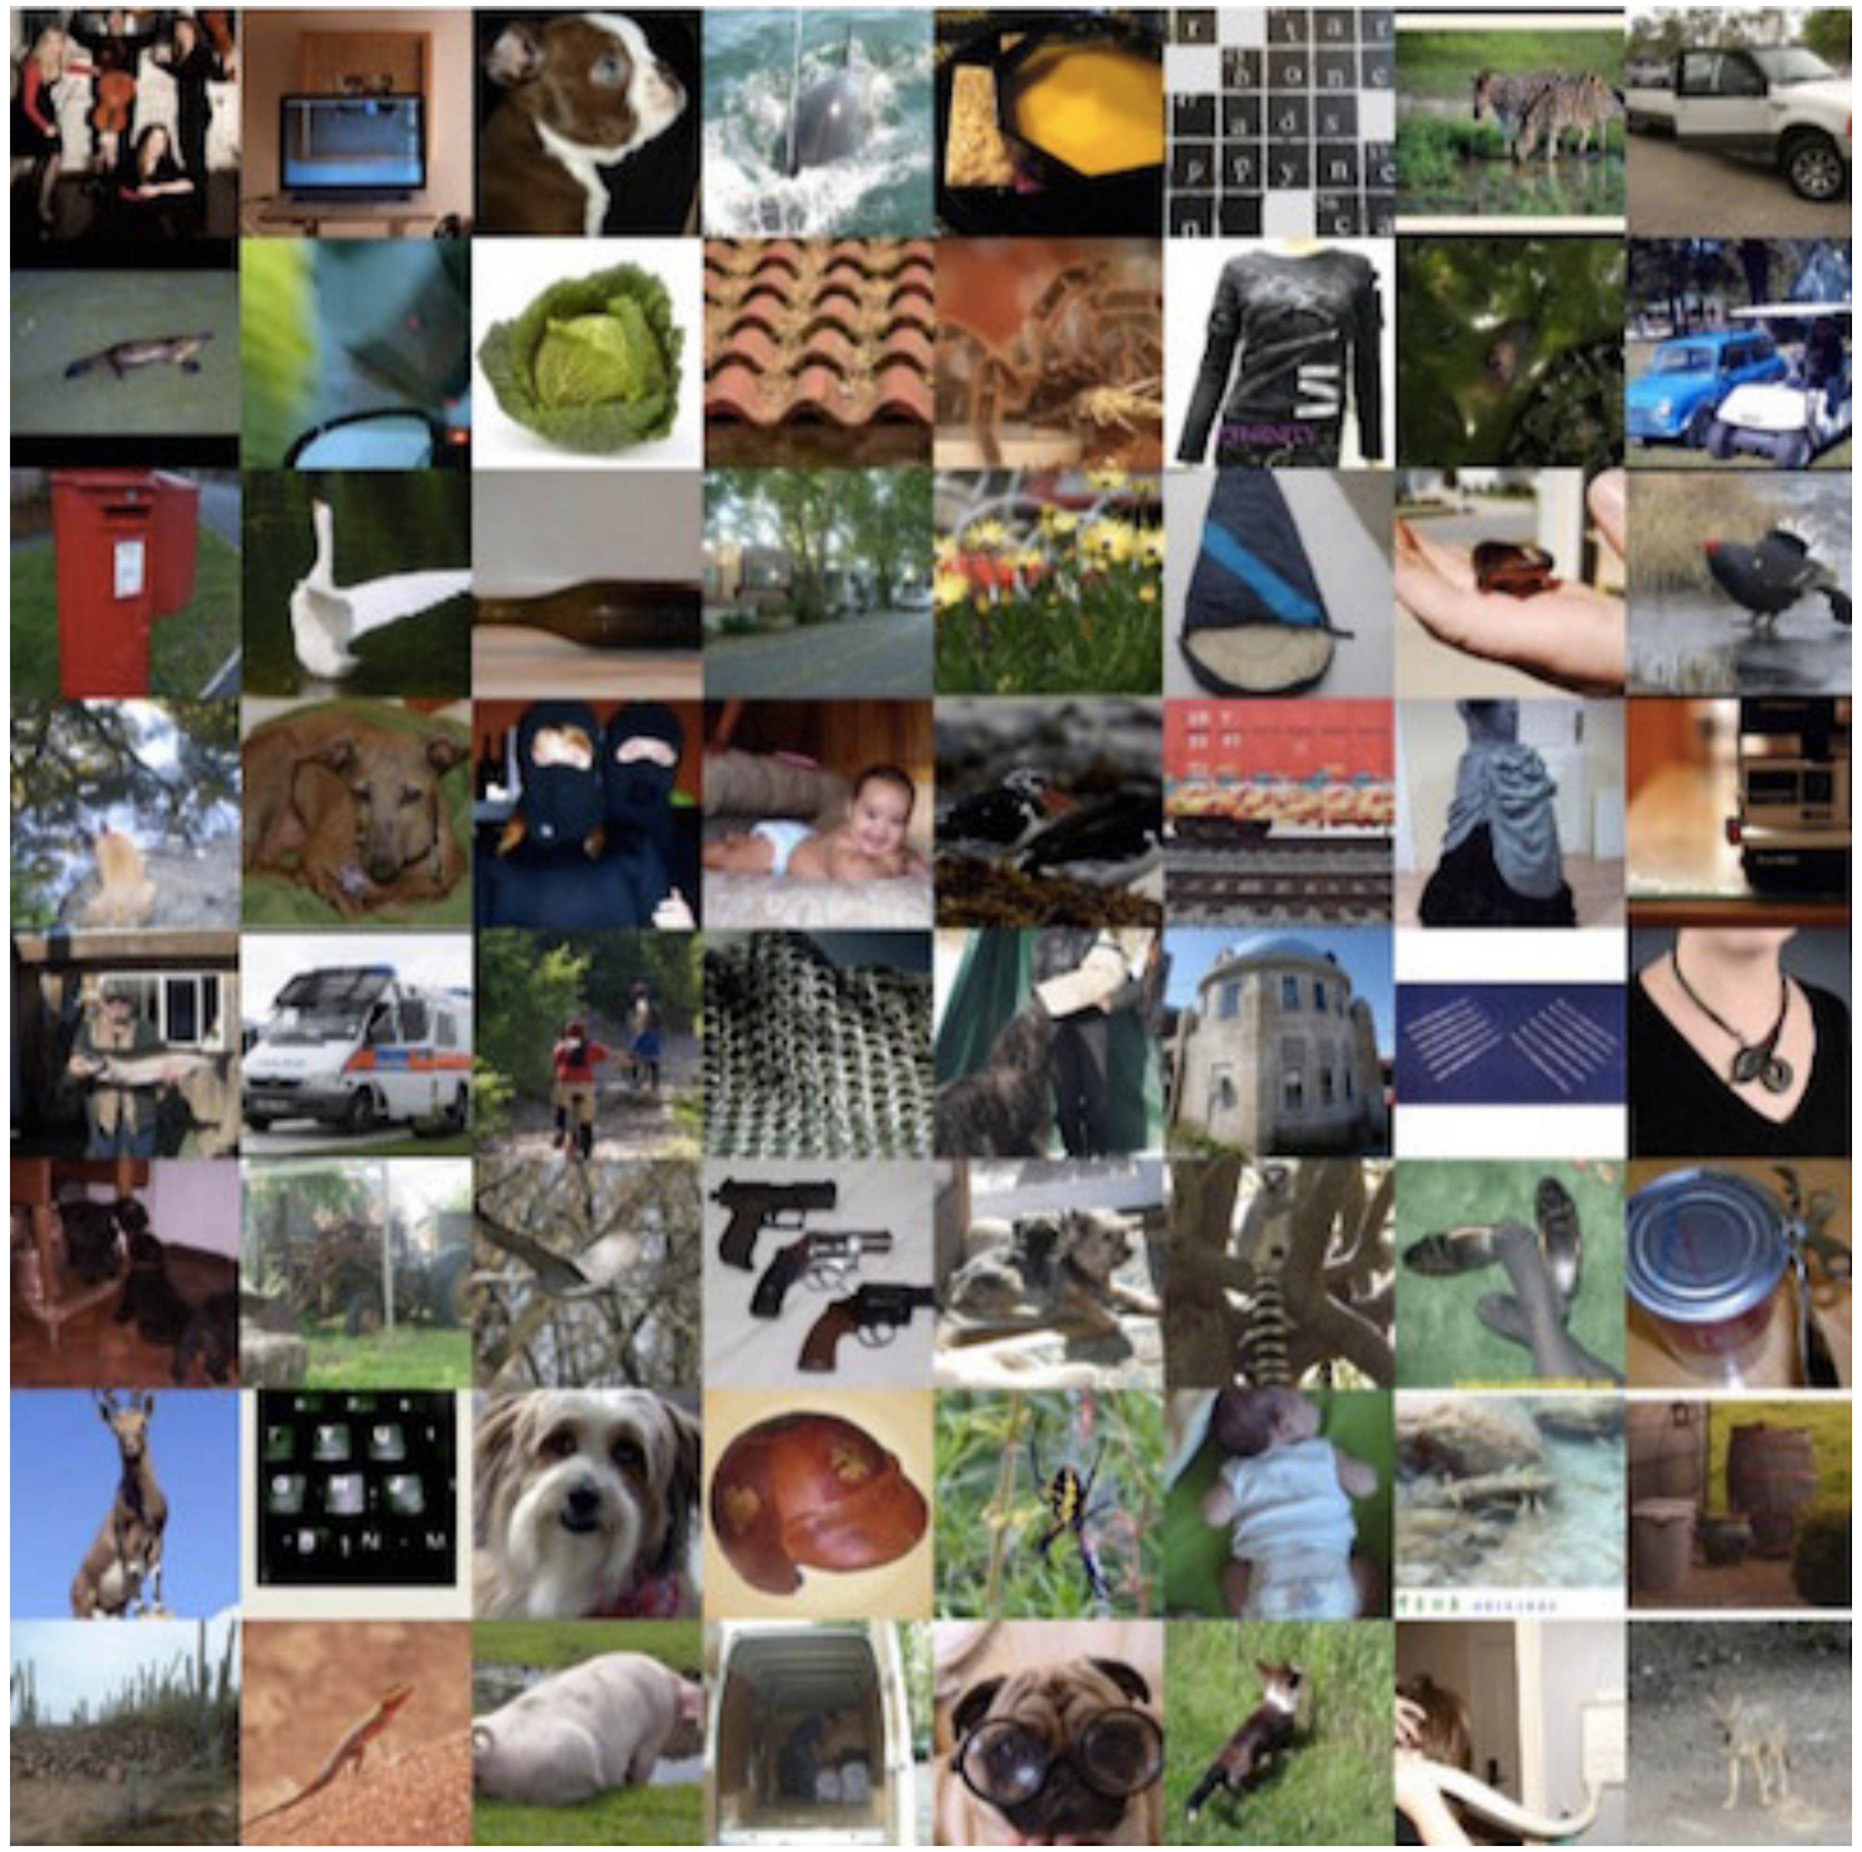

(Salimans et al., 2016)

Cherry-Picked Results





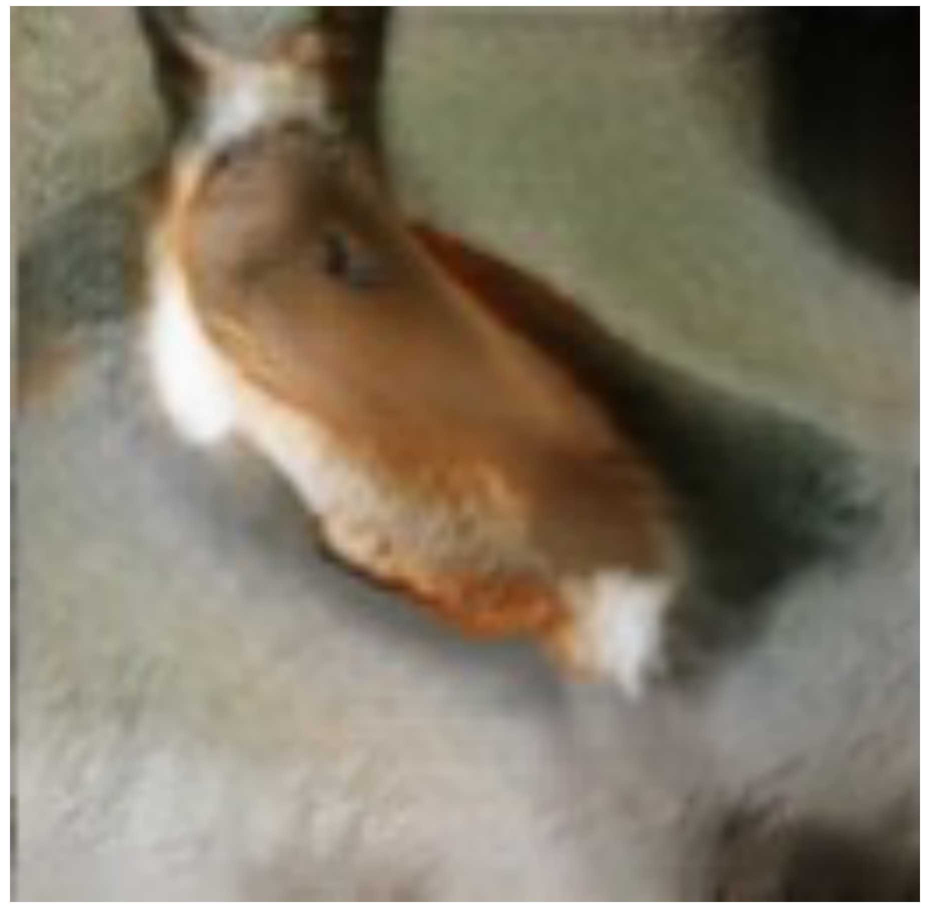

(Goodfellow 2017)

Problems with Counting

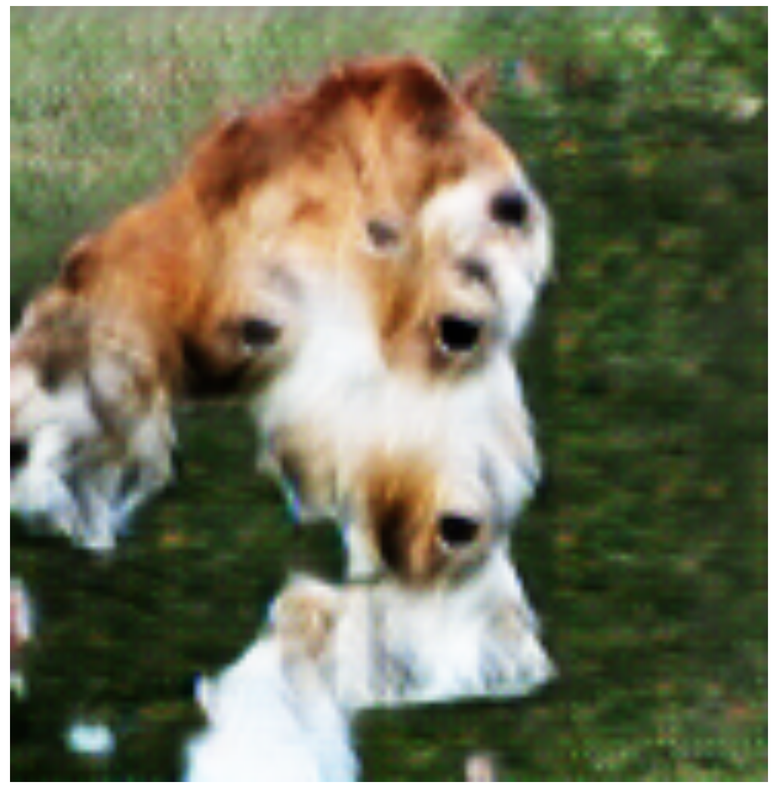




(Goodfellow 2017)

Problems with Perspective

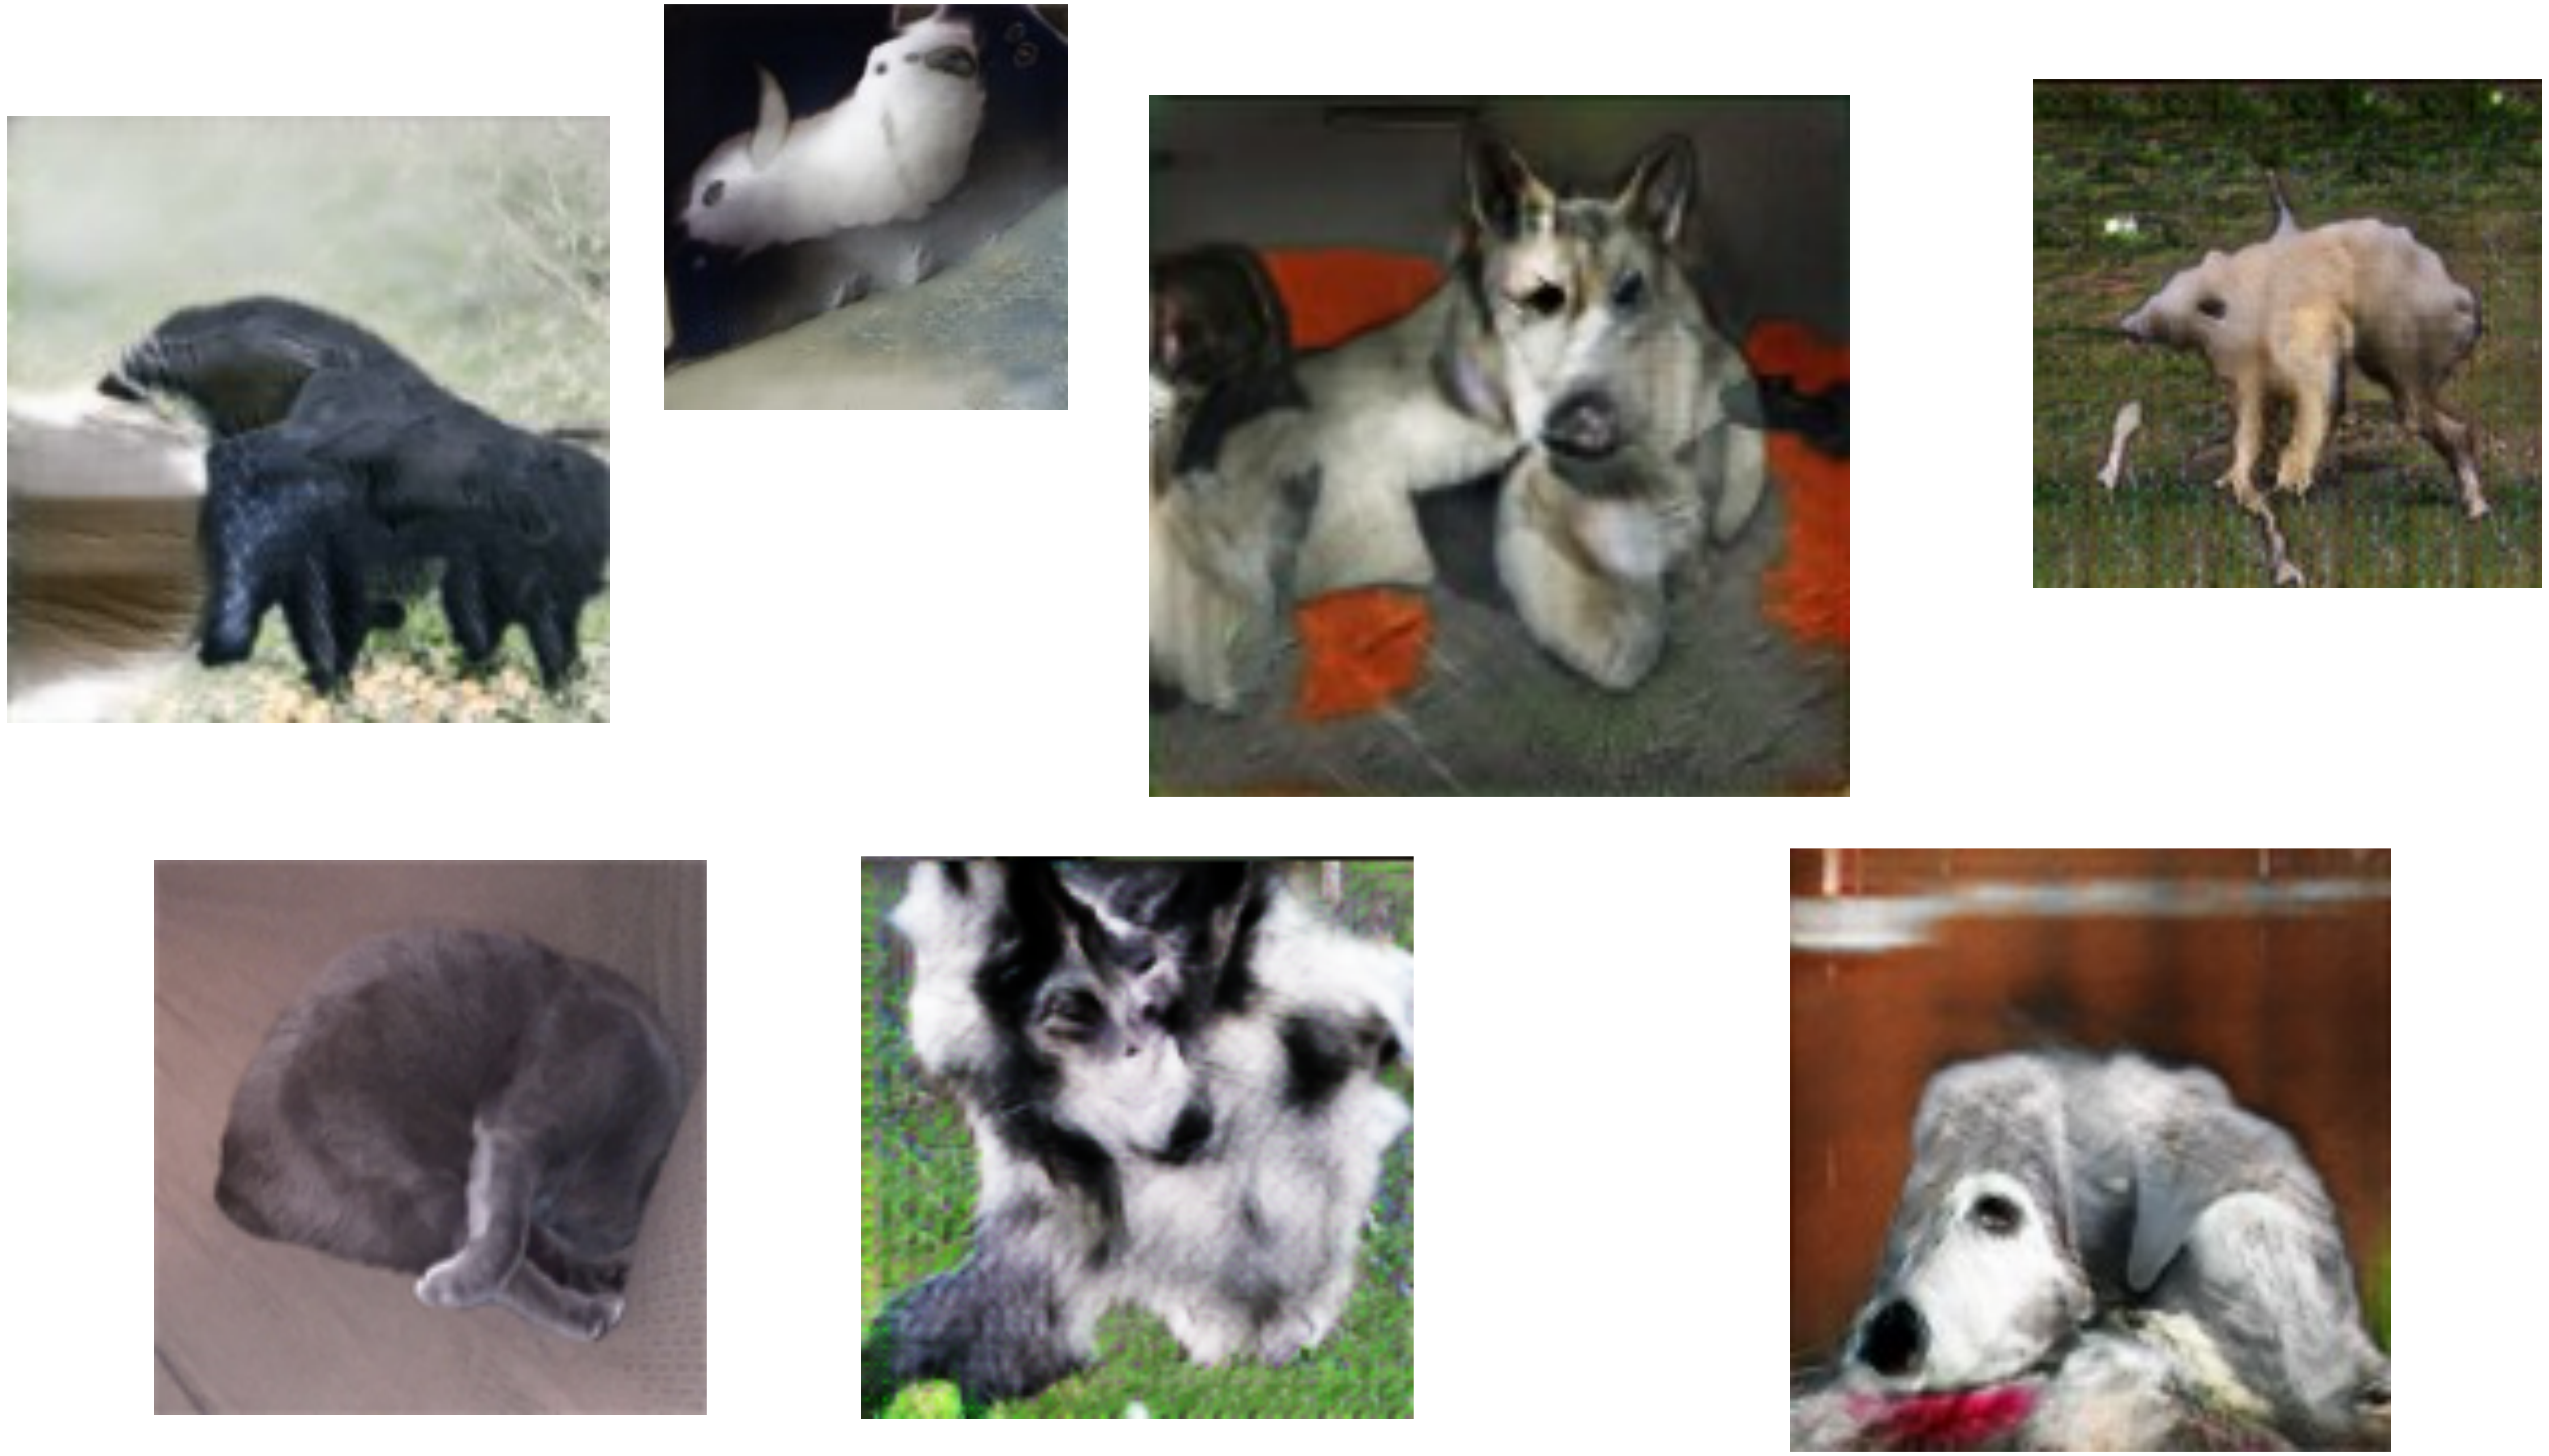

(Goodfellow 2017)

Problems with Global Structure





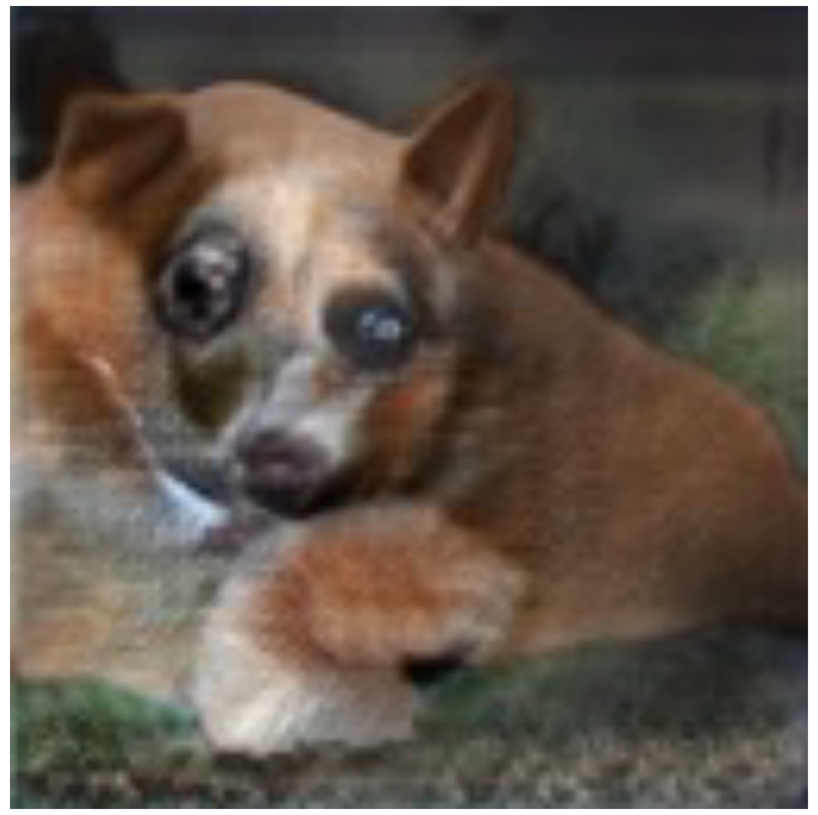
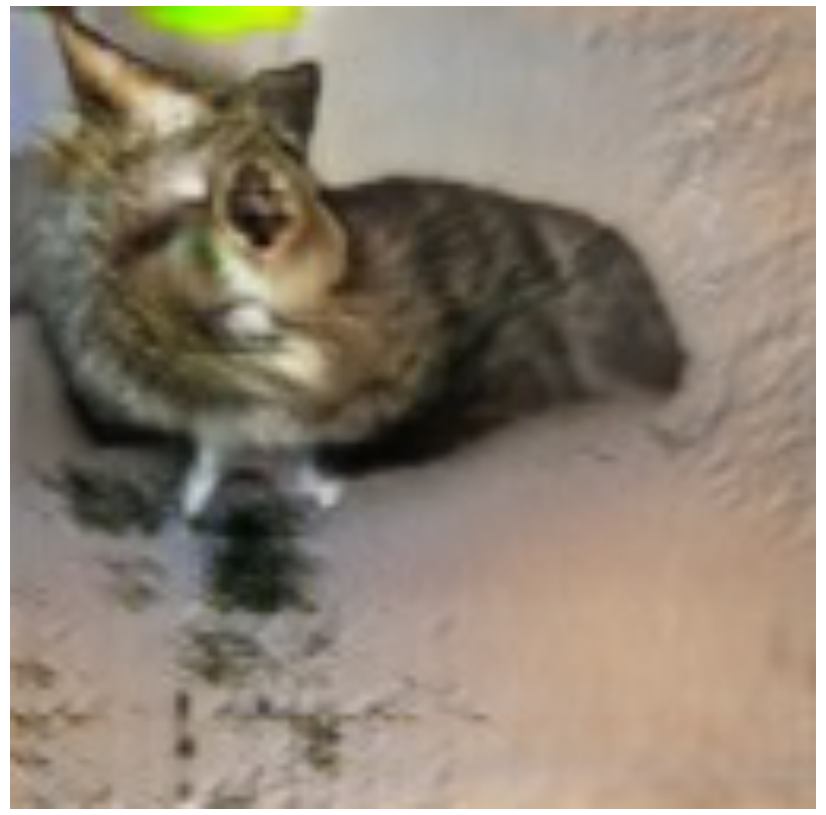

(Goodfellow 2017)

This one is real

(Goodfellow 2017)

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