Generative Adversarial Networks (GANs)

Ian Goodfellow, OpenAI Research Scientist

Re-Work Deep Learning Summit

San Francisco, 2017-01-26

Generative Modeling

- Density estimation

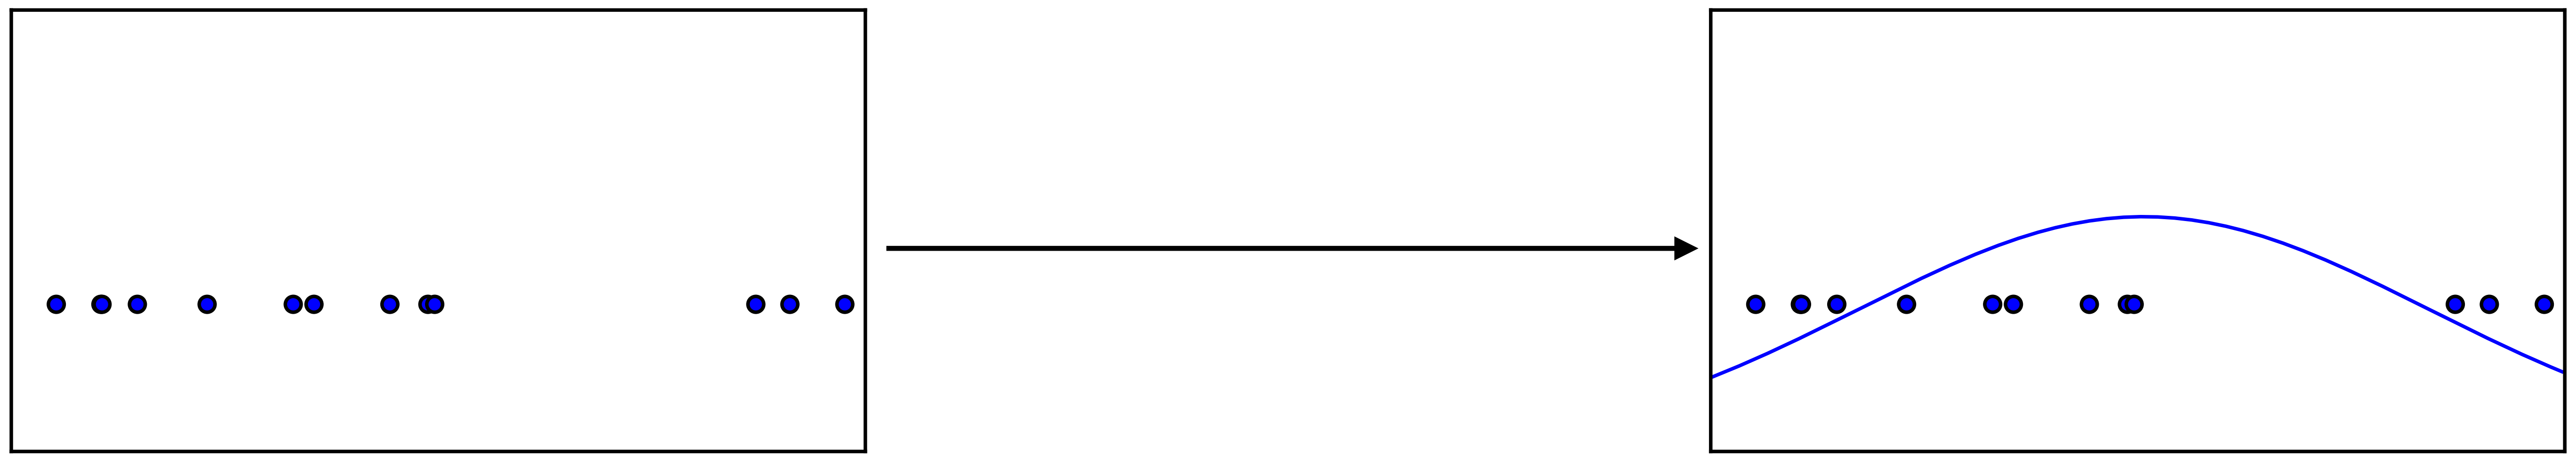

- Sample generation

Training examples Model samples

(Goodfellow 2016)

Next Video Frame Prediction

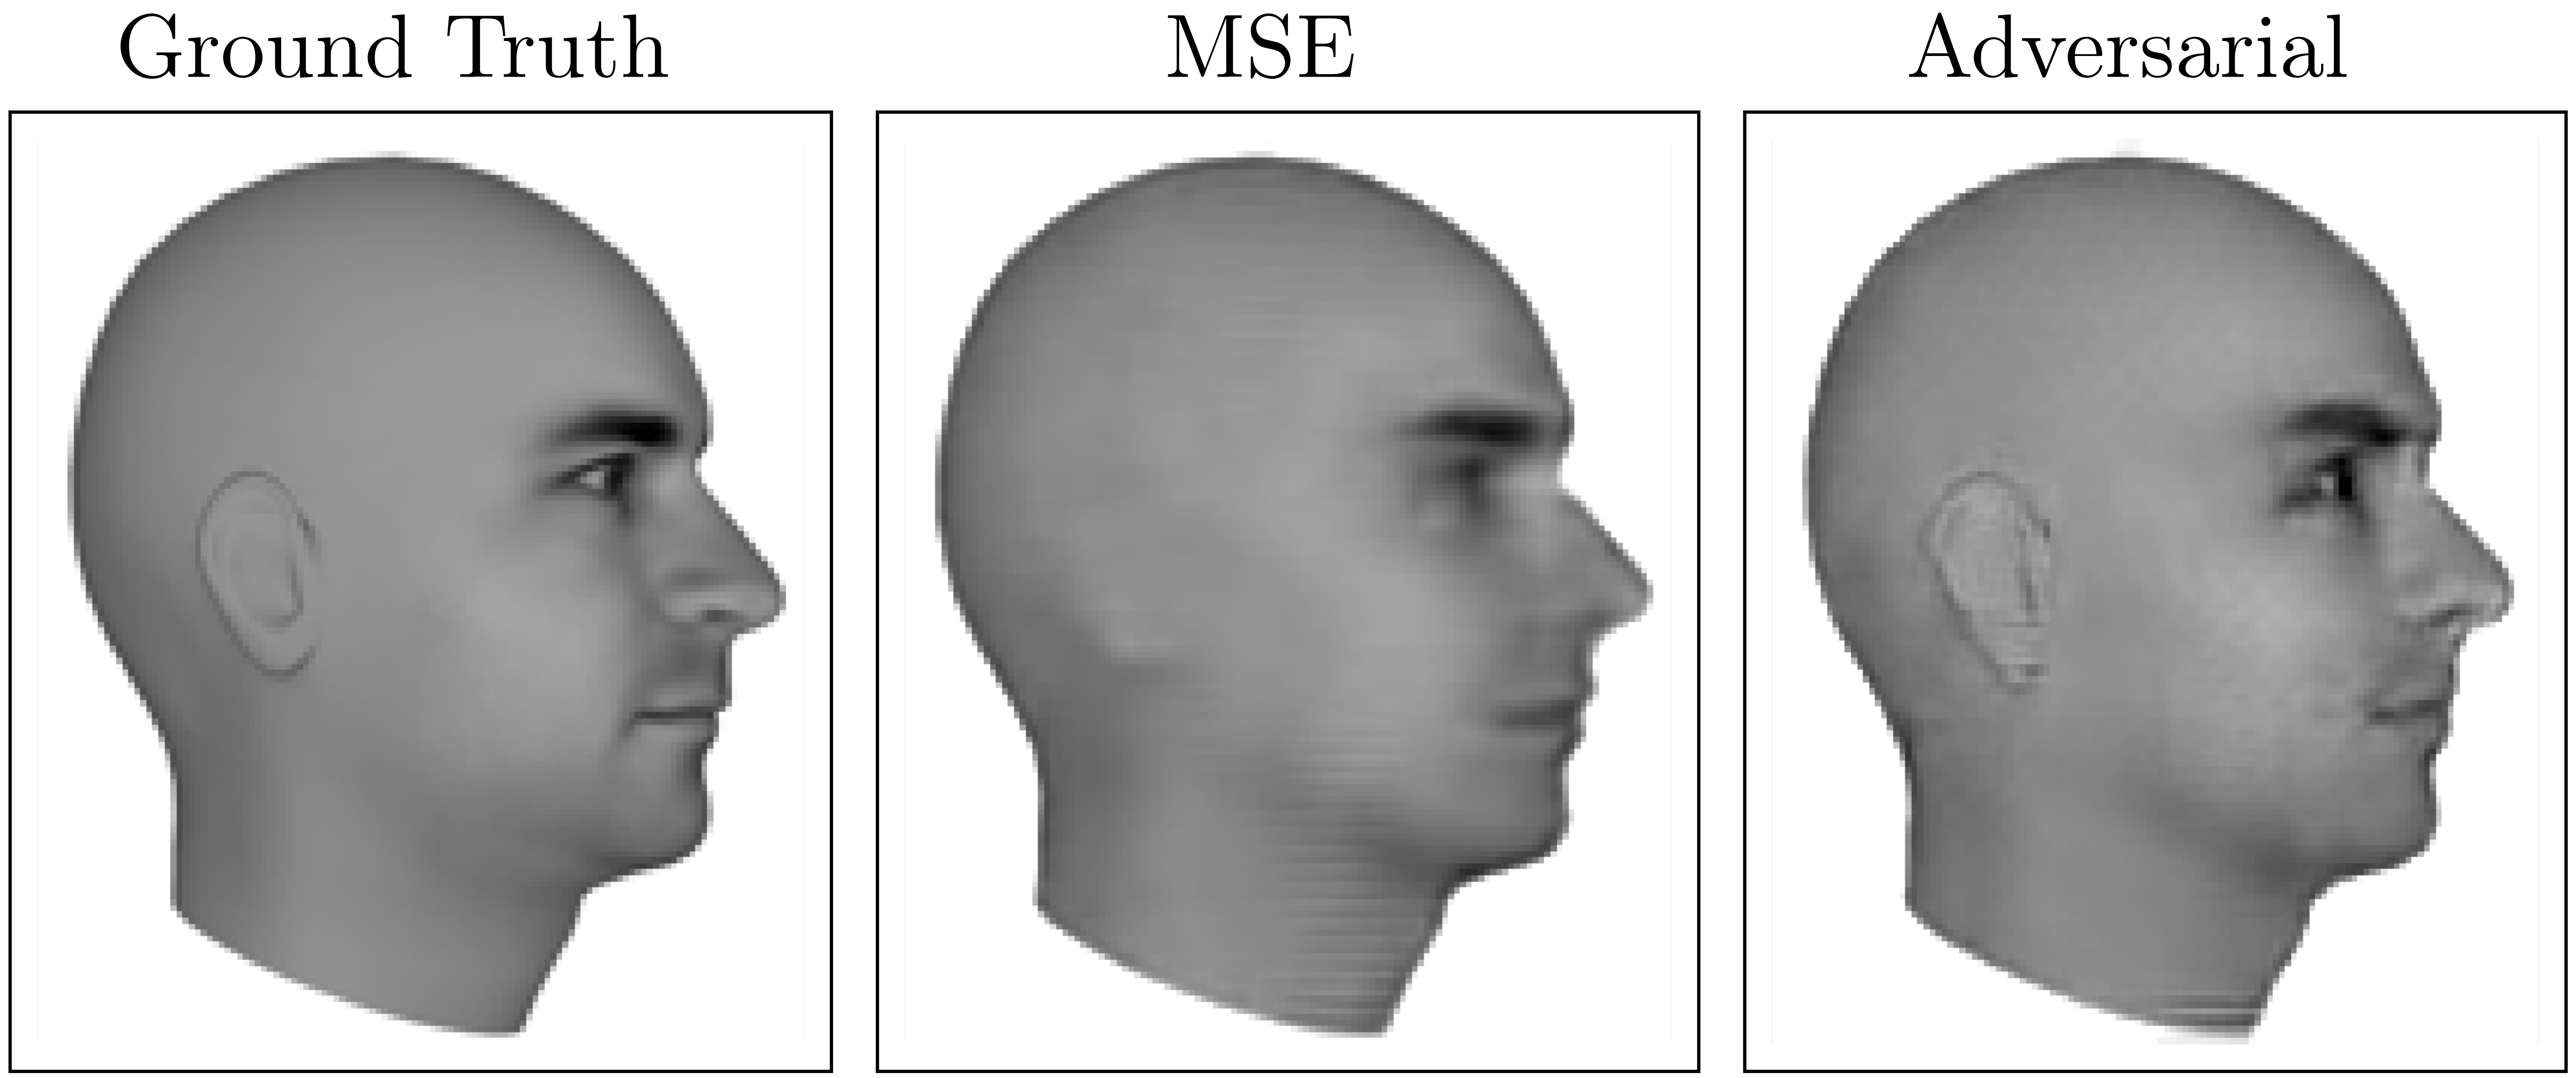

(Lotter et al 2016)

iGAN

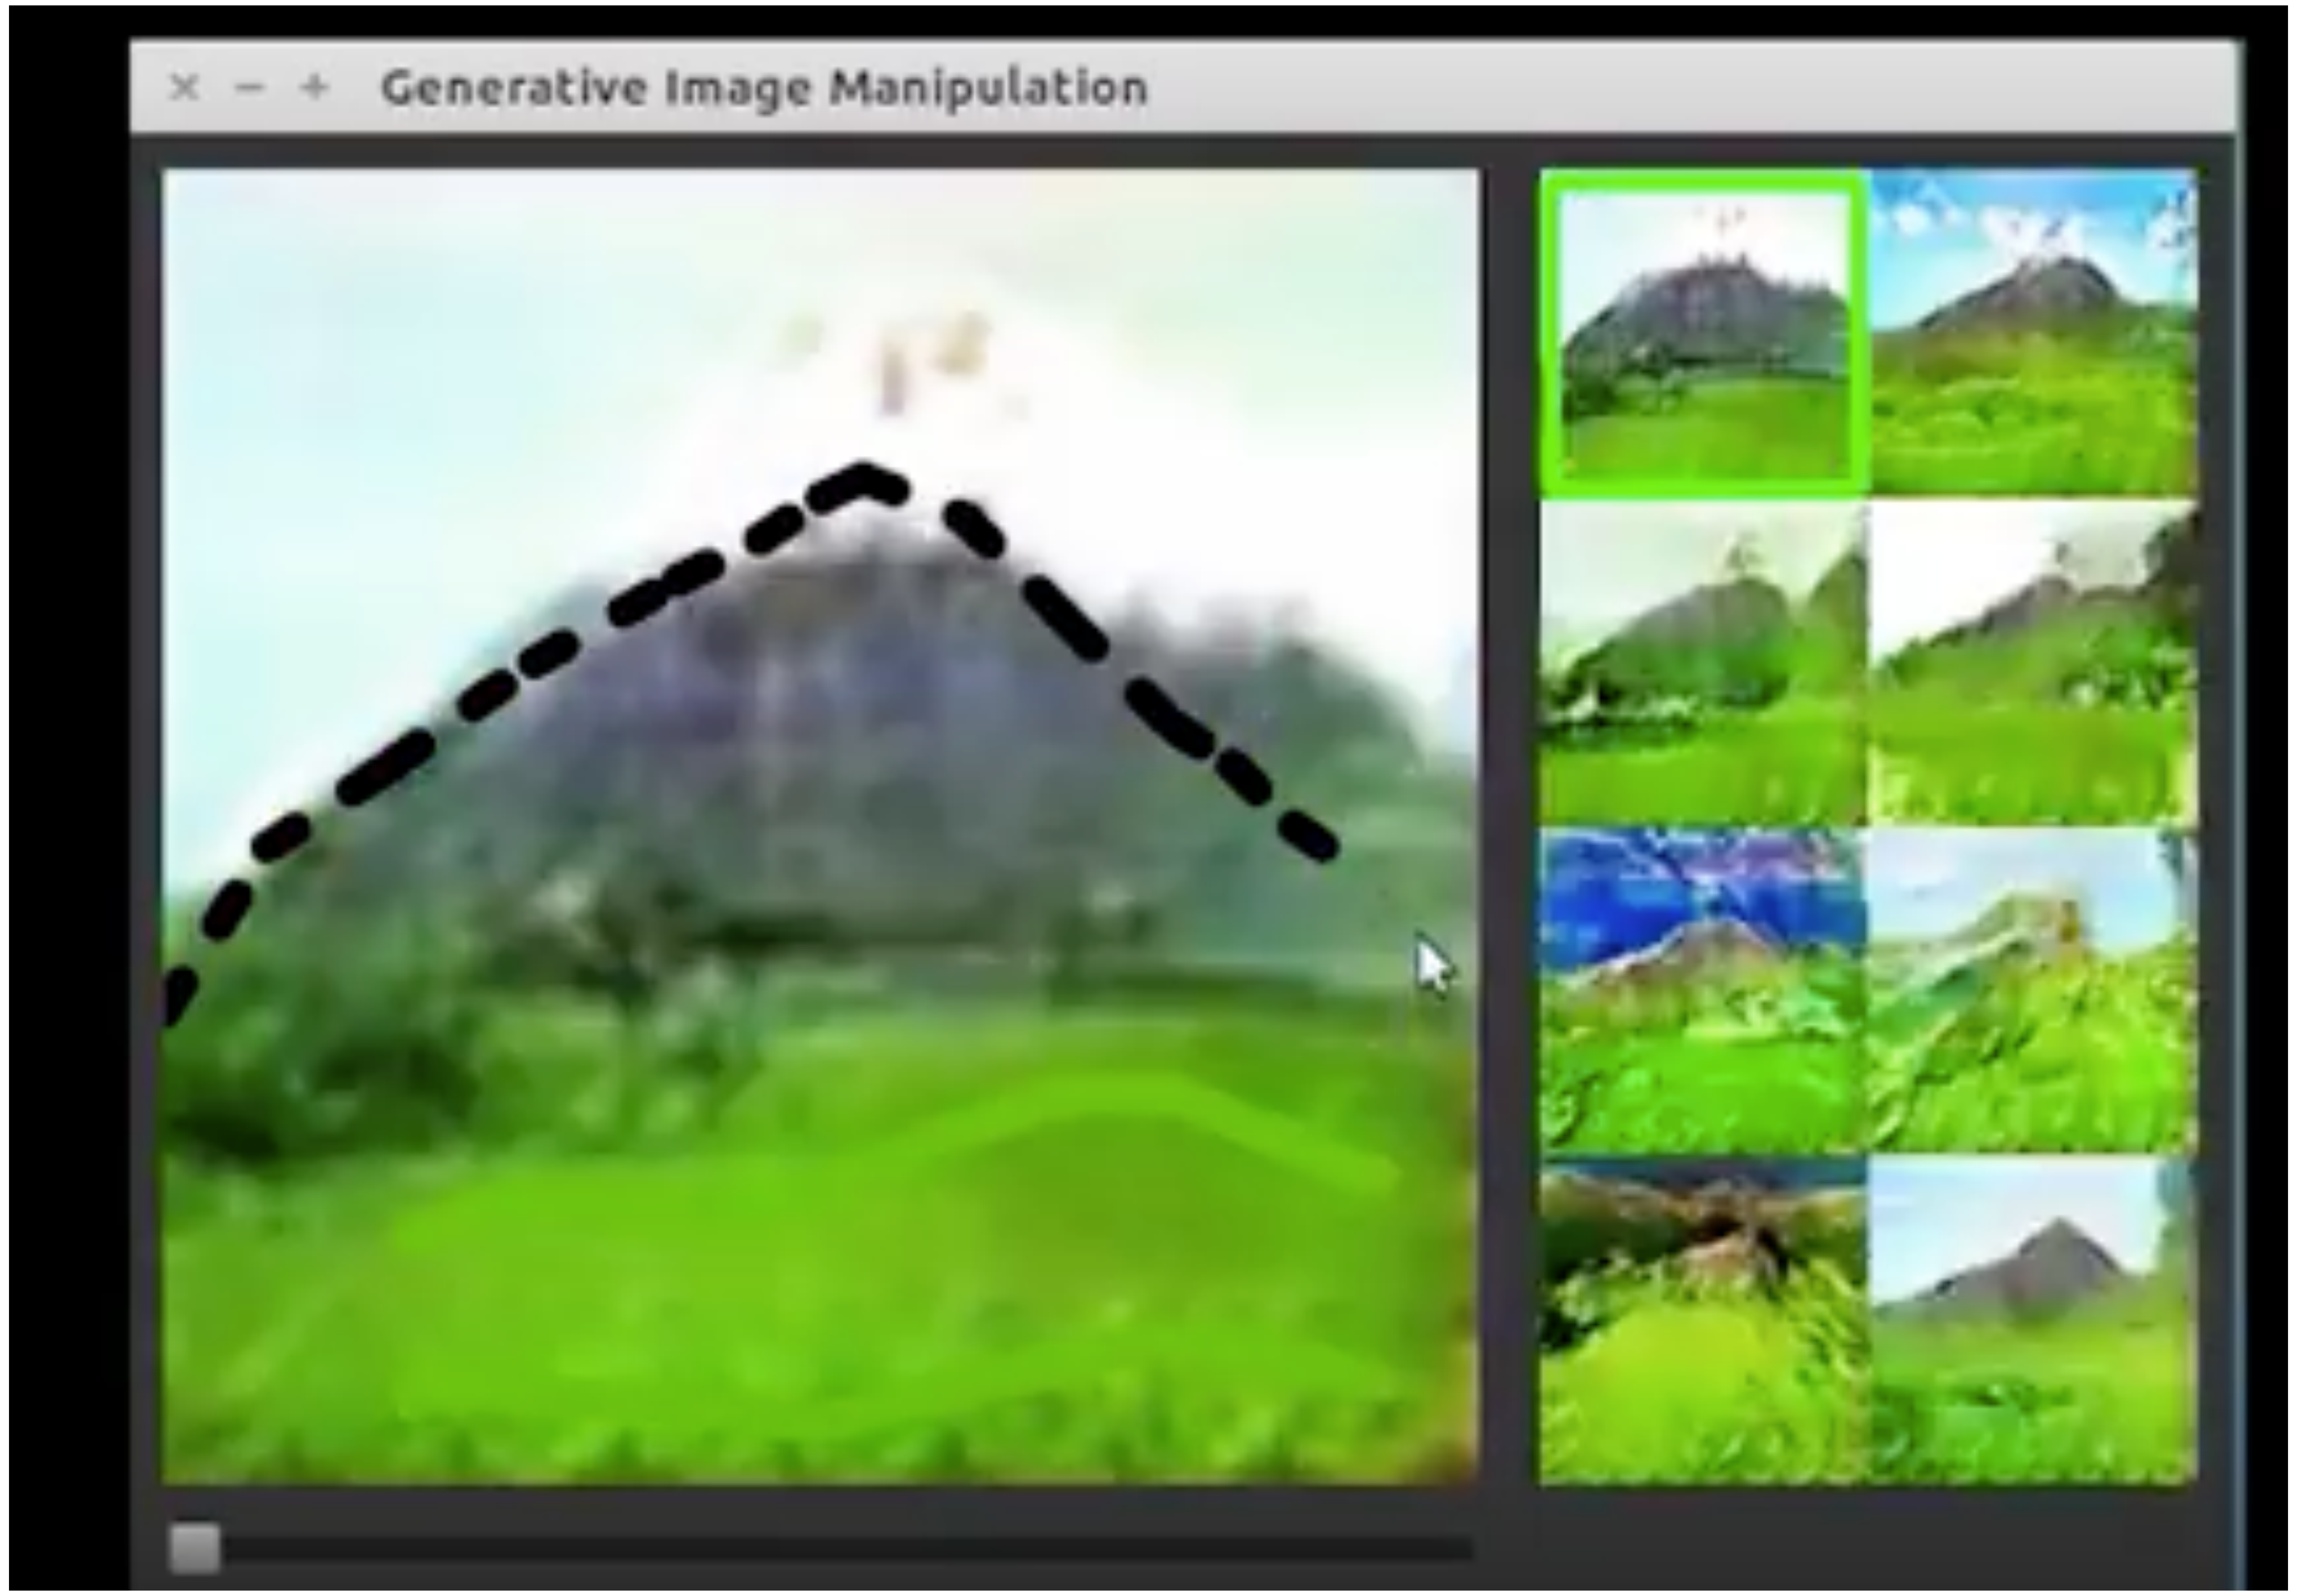

youtube

(Zhu et al 2016)

IAN

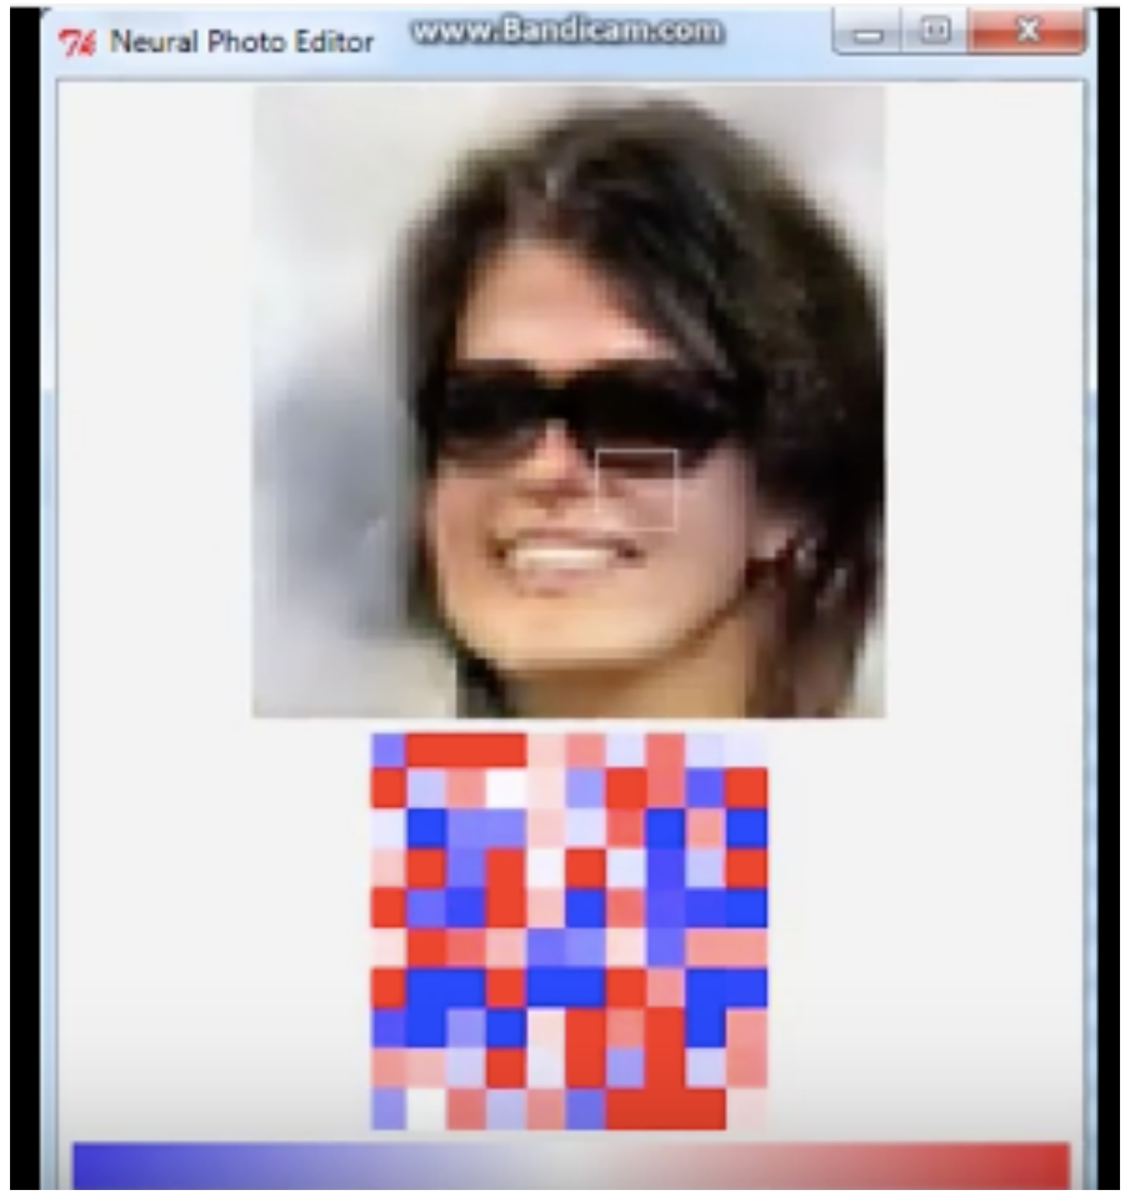

youtube

(Brock et al 2016)

Image to Image Translation

(Isola et al 2016)

Fully Visible Belief Nets

- Explicit formula based on chain rule: (Frey et al, 1996)

$$p_{\text{model}}(\boldsymbol{x}) = p_{\text{model}}(x_1) \prod_{i=2}^{n} p_{\text{model}}(x_i \mid x_1, \ldots, x_{i-1})$$

- Disadvantages:
  - $\mathrm{O}(n)$ sample generation cost
  - Generation not controlled by a latent code

PixelCNN elephants (van den Ord et al 2016)

WaveNet

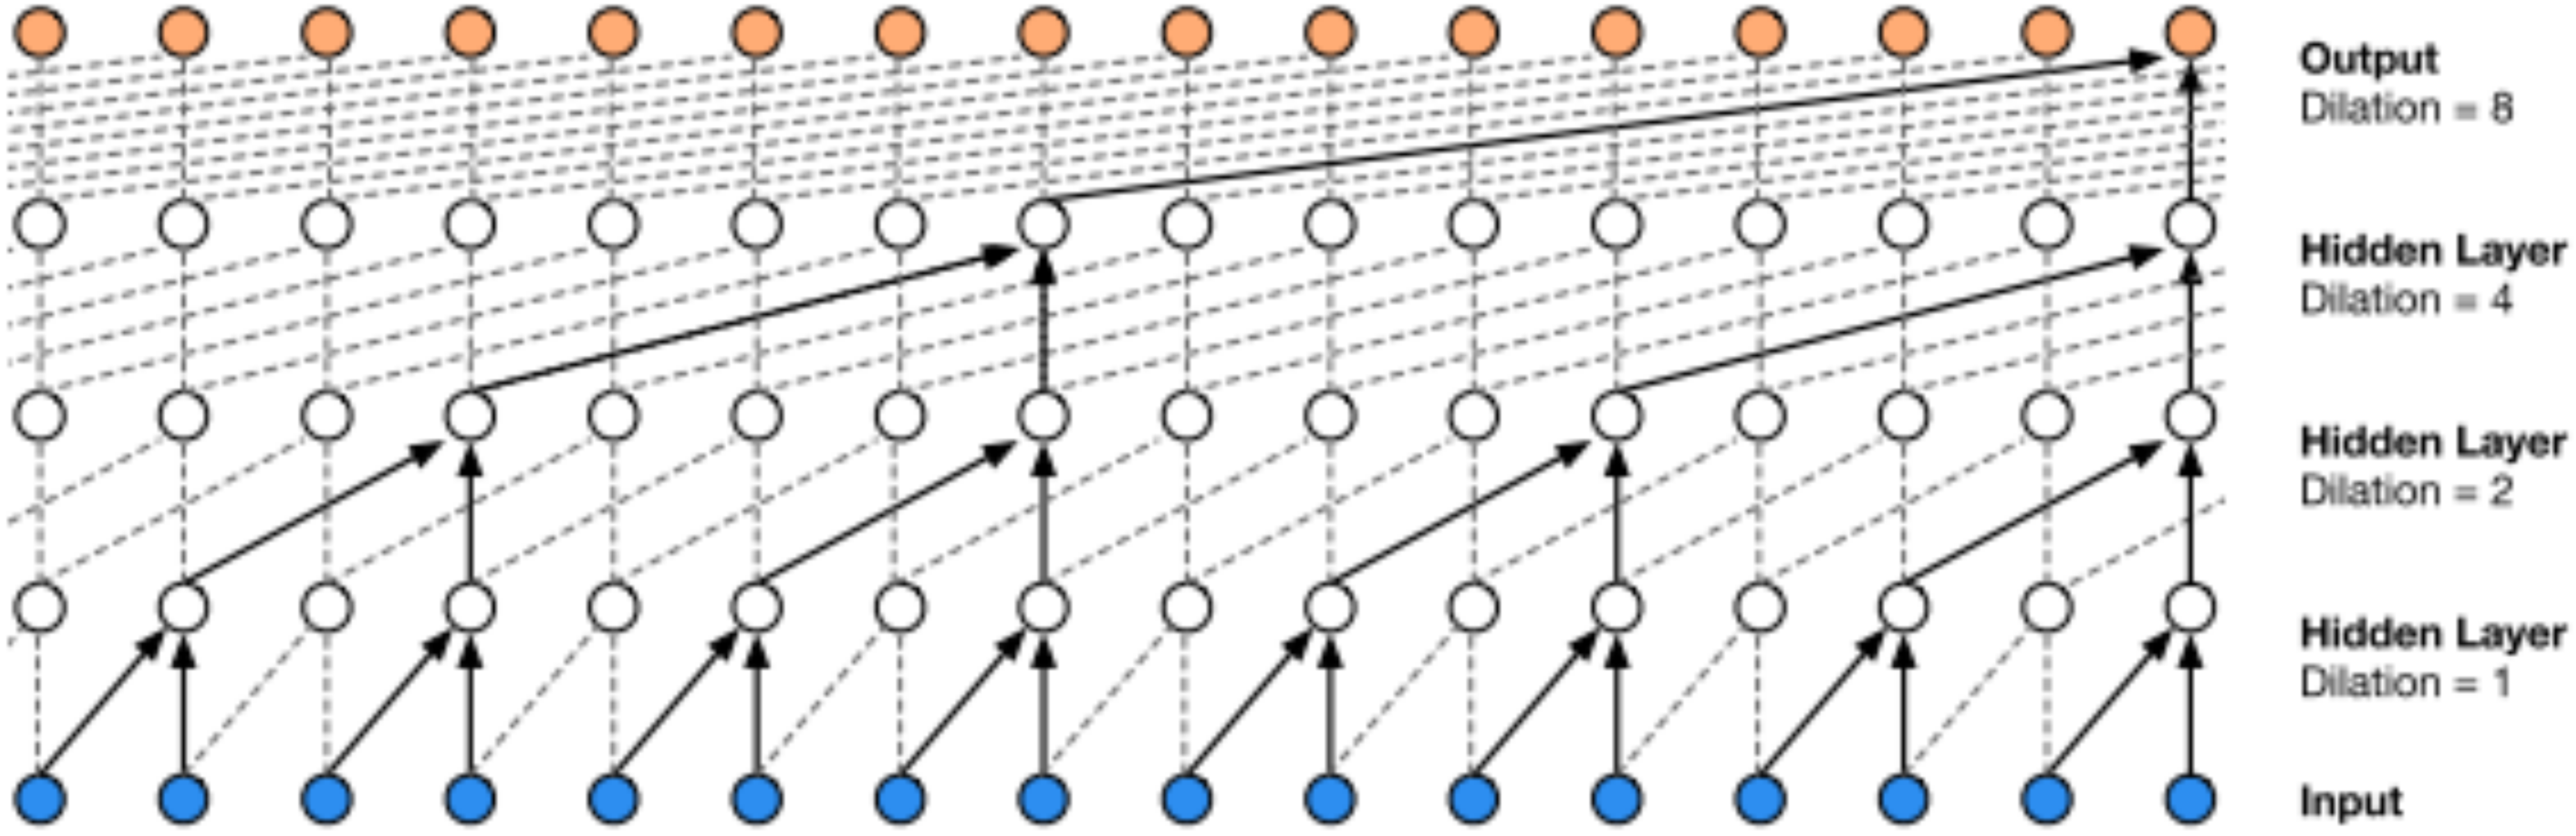

Amazing quality
Sample generation slow

Two minutes to synthesize
one second of audio

(Goodfellow 2016)

Adversarial Nets Framework

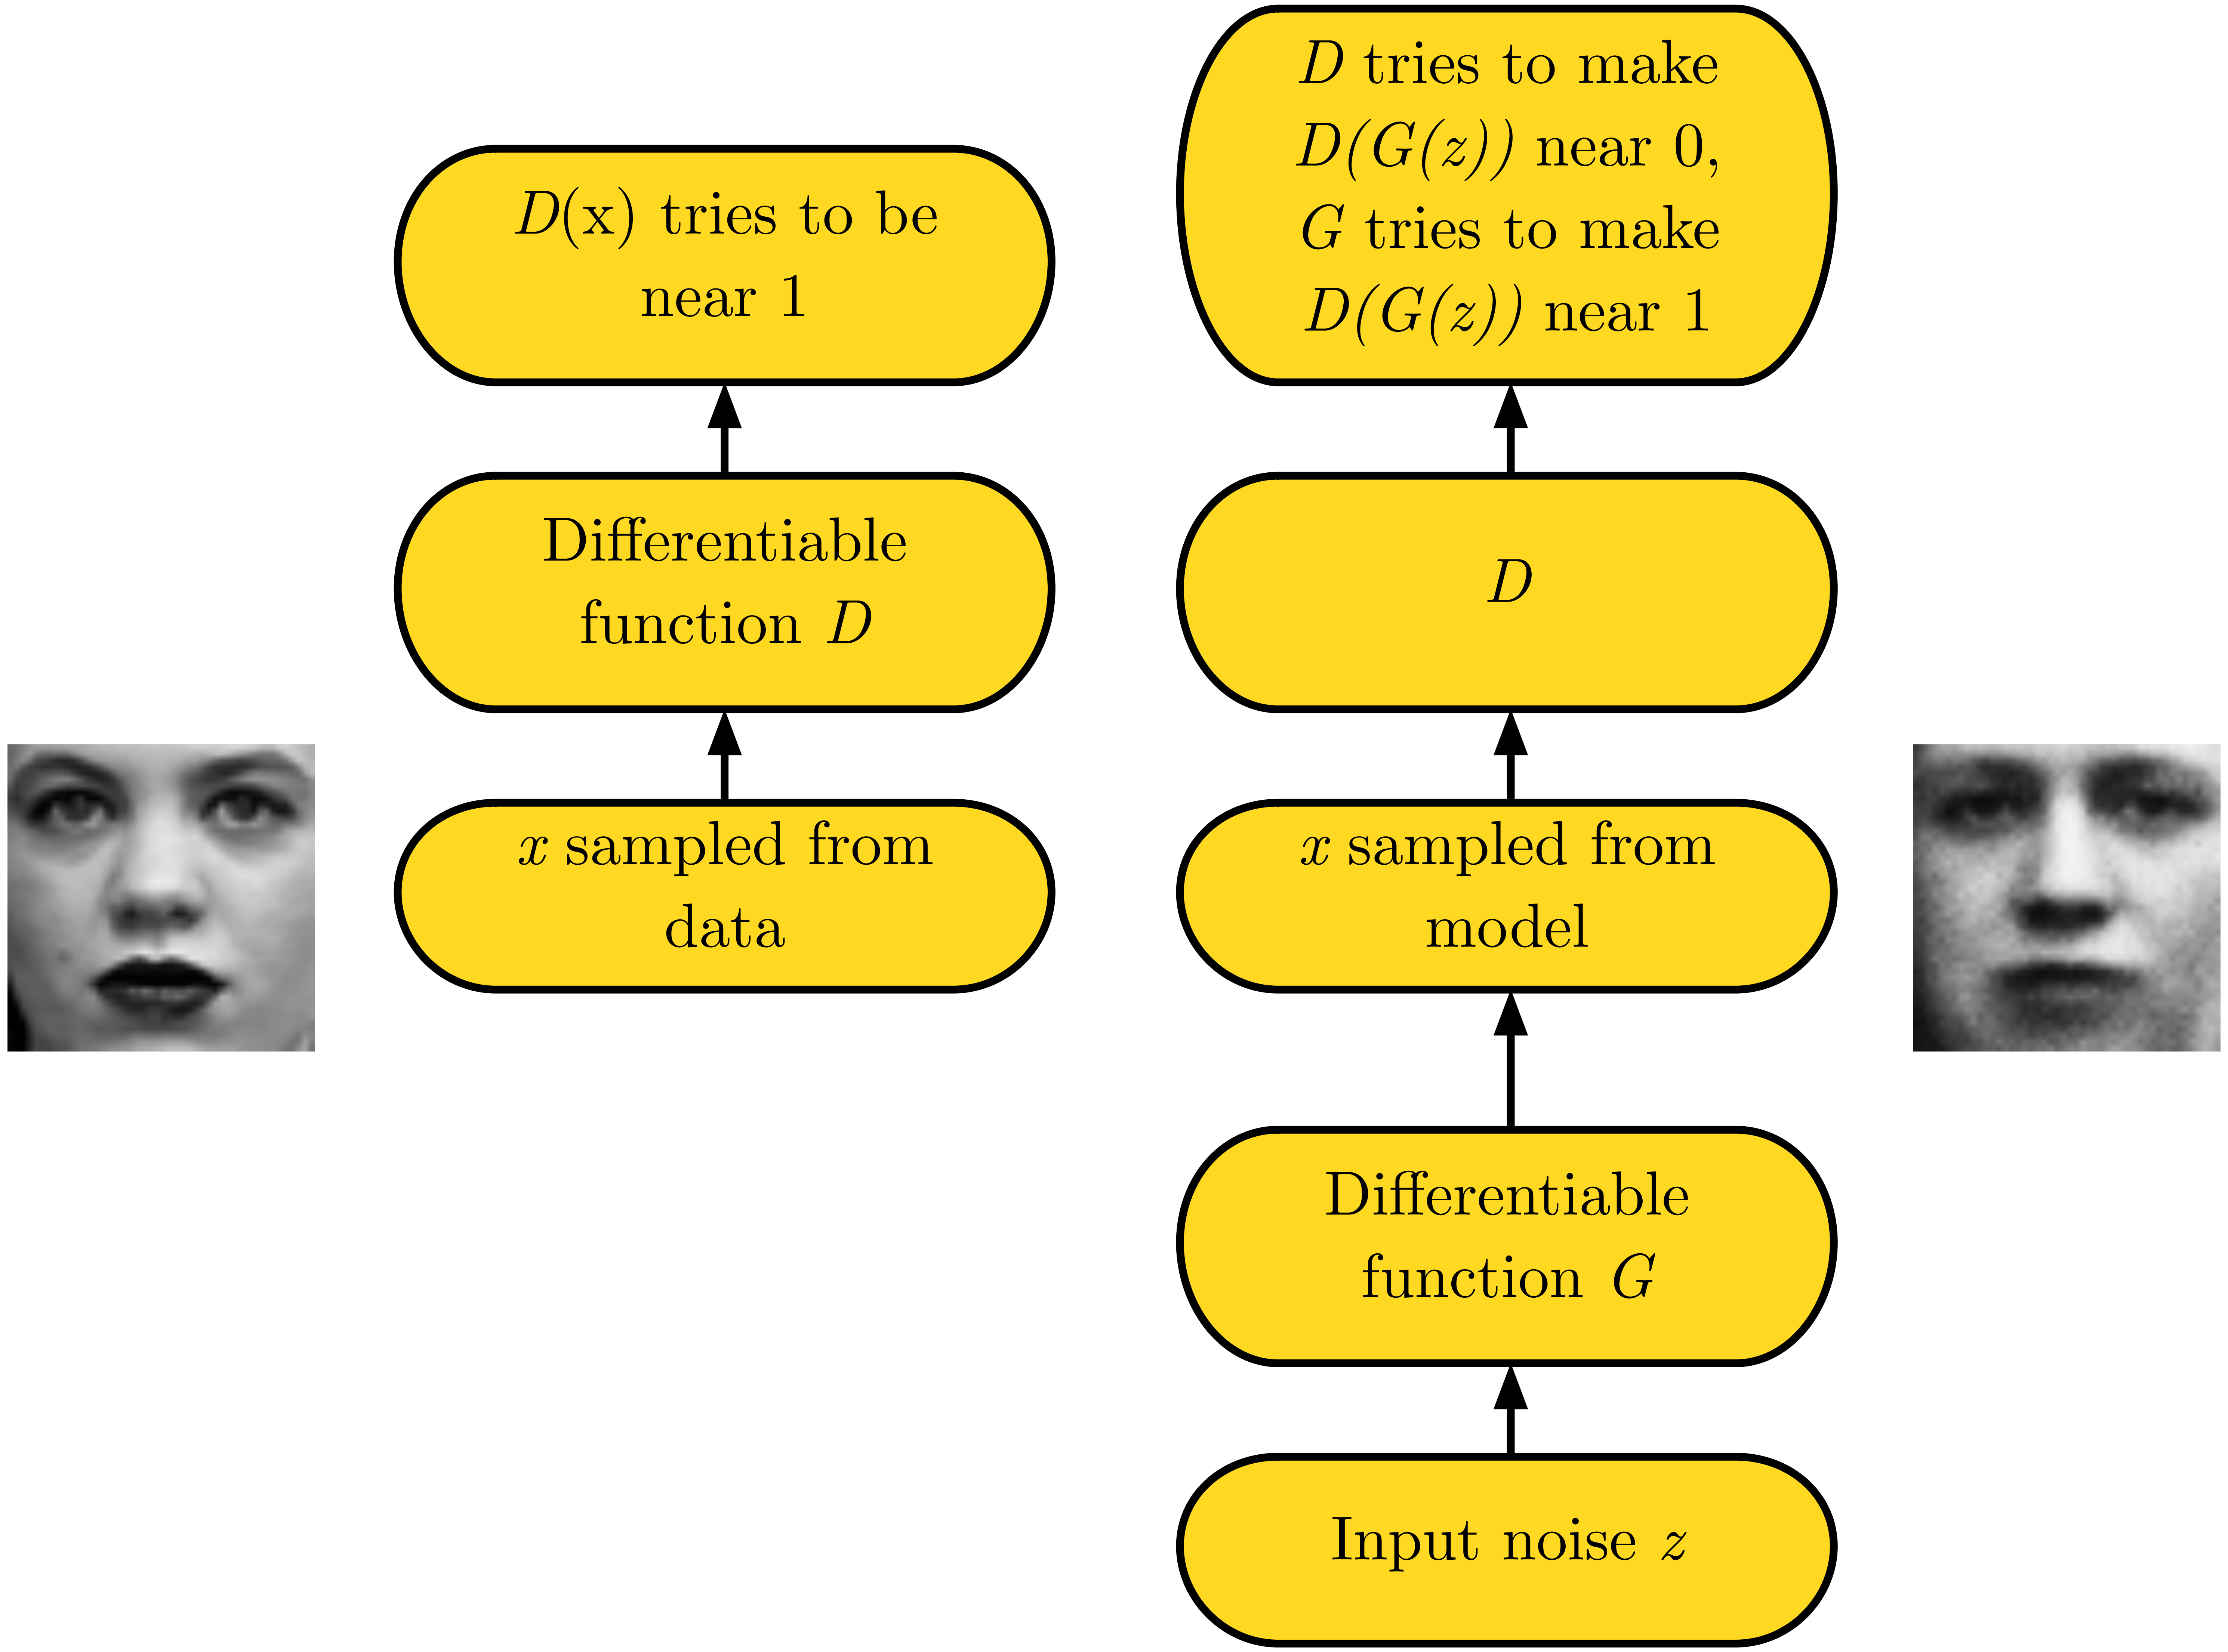

(Goodfellow 2016)

Vector Space Arithmetic


Man with glasses

-

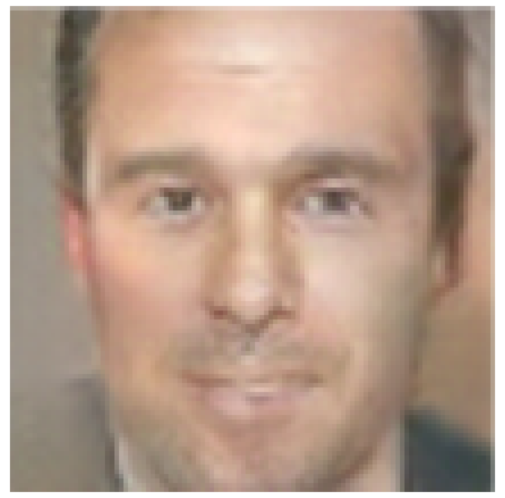
Man

+

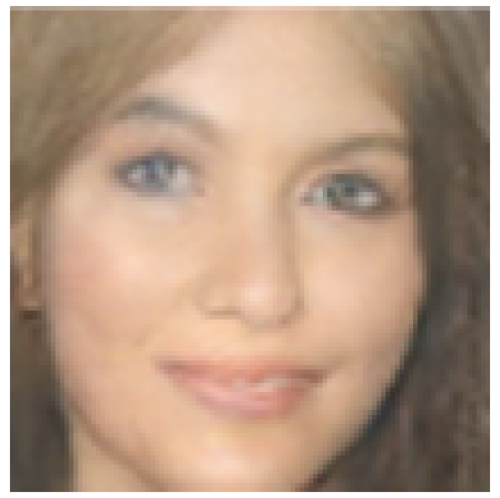
Woman

=


Woman with Glasses

(Radford et al, 2015)

3D GAN

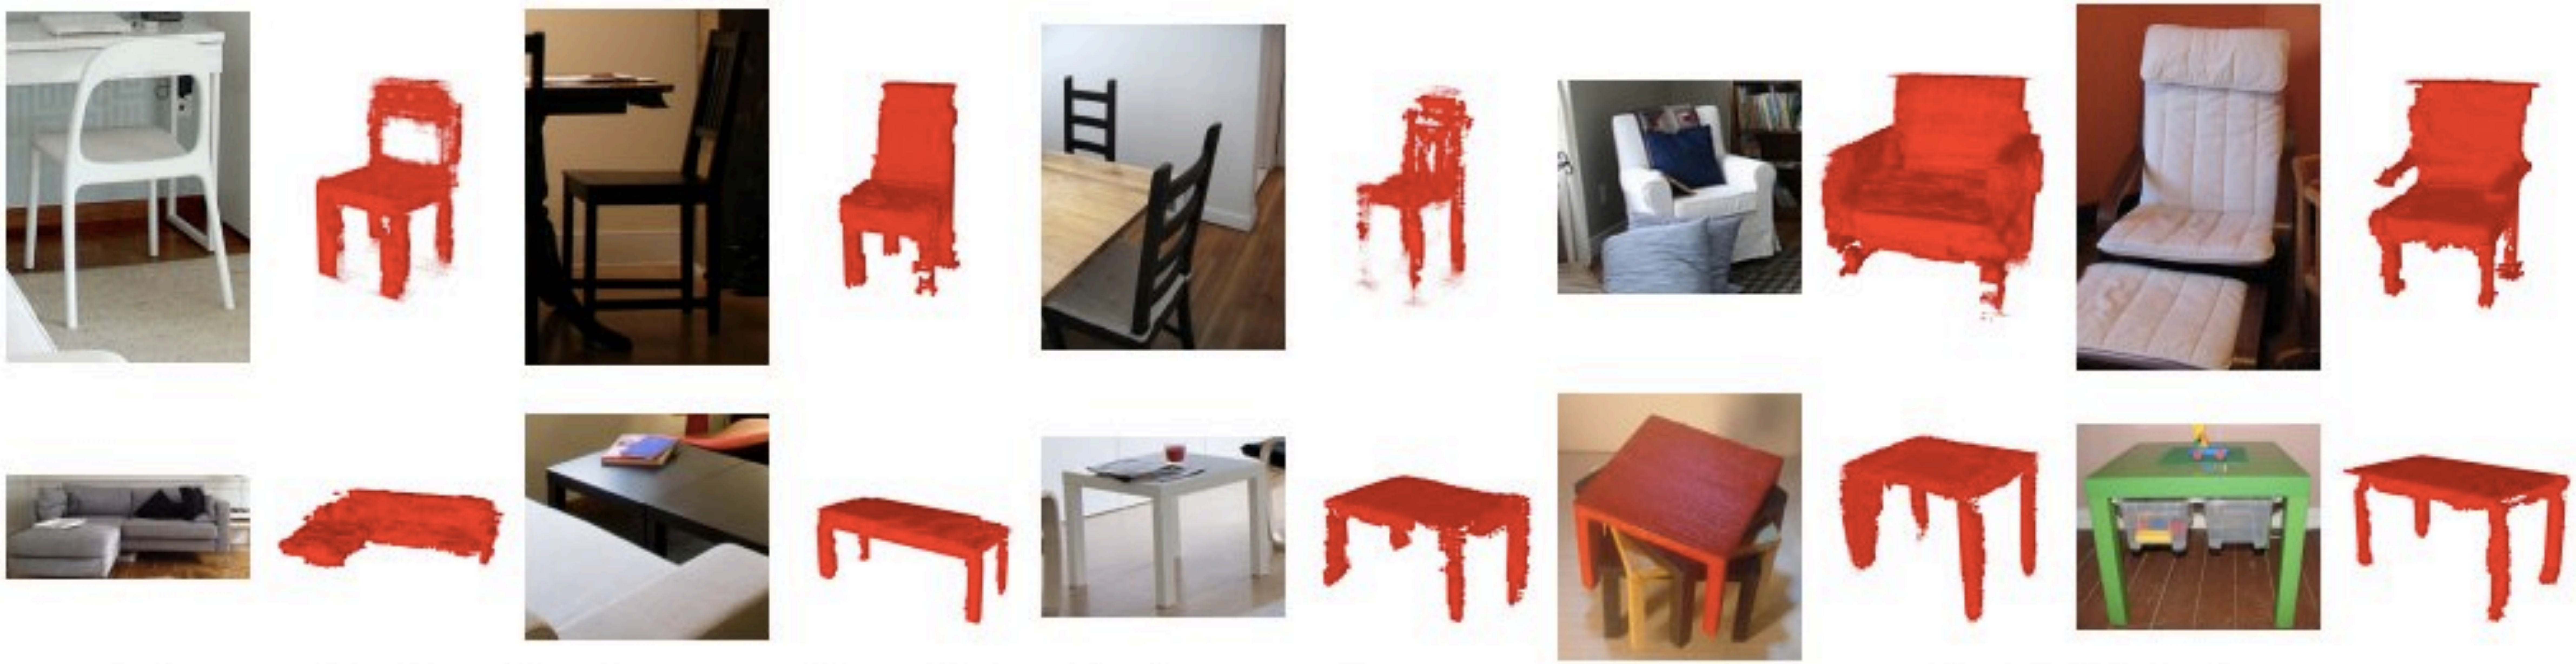

Figure 7: Qualitative results of single image 3D reconstruction on the IKEA dataset

(Wu et al, 2016)

OpenAI GAN-created images

(Goodfellow 2016)

Problems with Counting

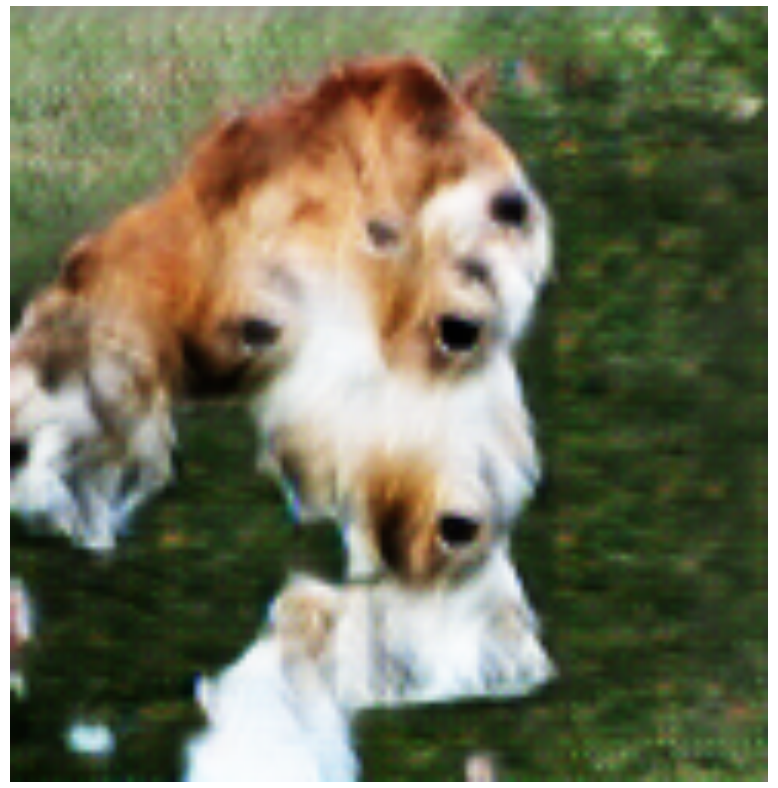

(Goodfellow 2016)

Problems with Perspective

(Goodfellow 2016)

Problems with Global Structure





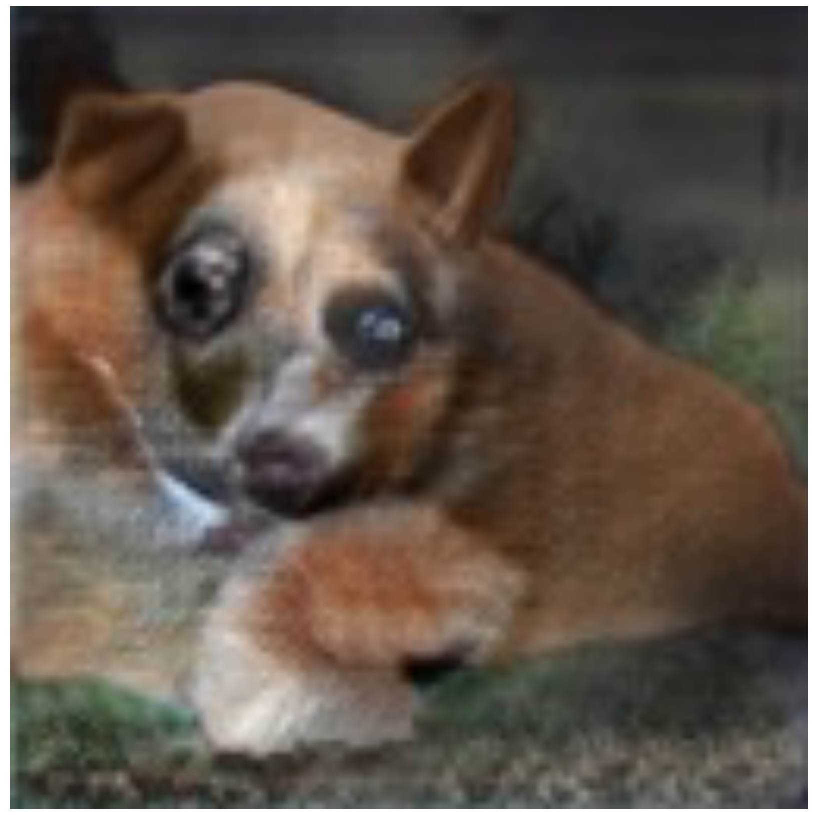
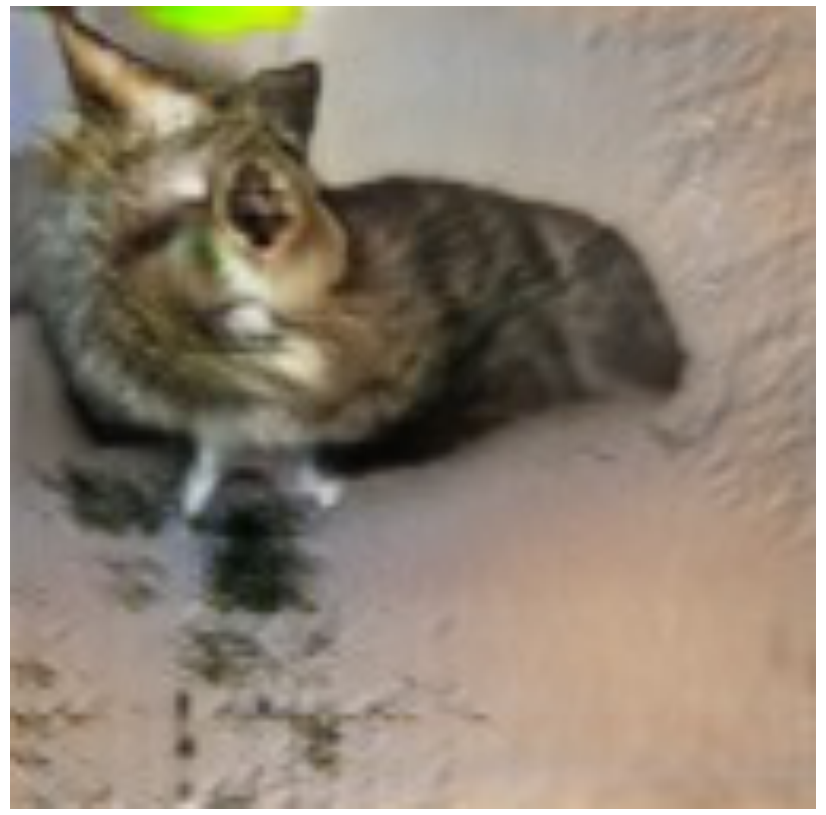

(Goodfellow 2016)

This one is real

(Goodfellow 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)

Learning interpretable latent codes / controlling the generation process

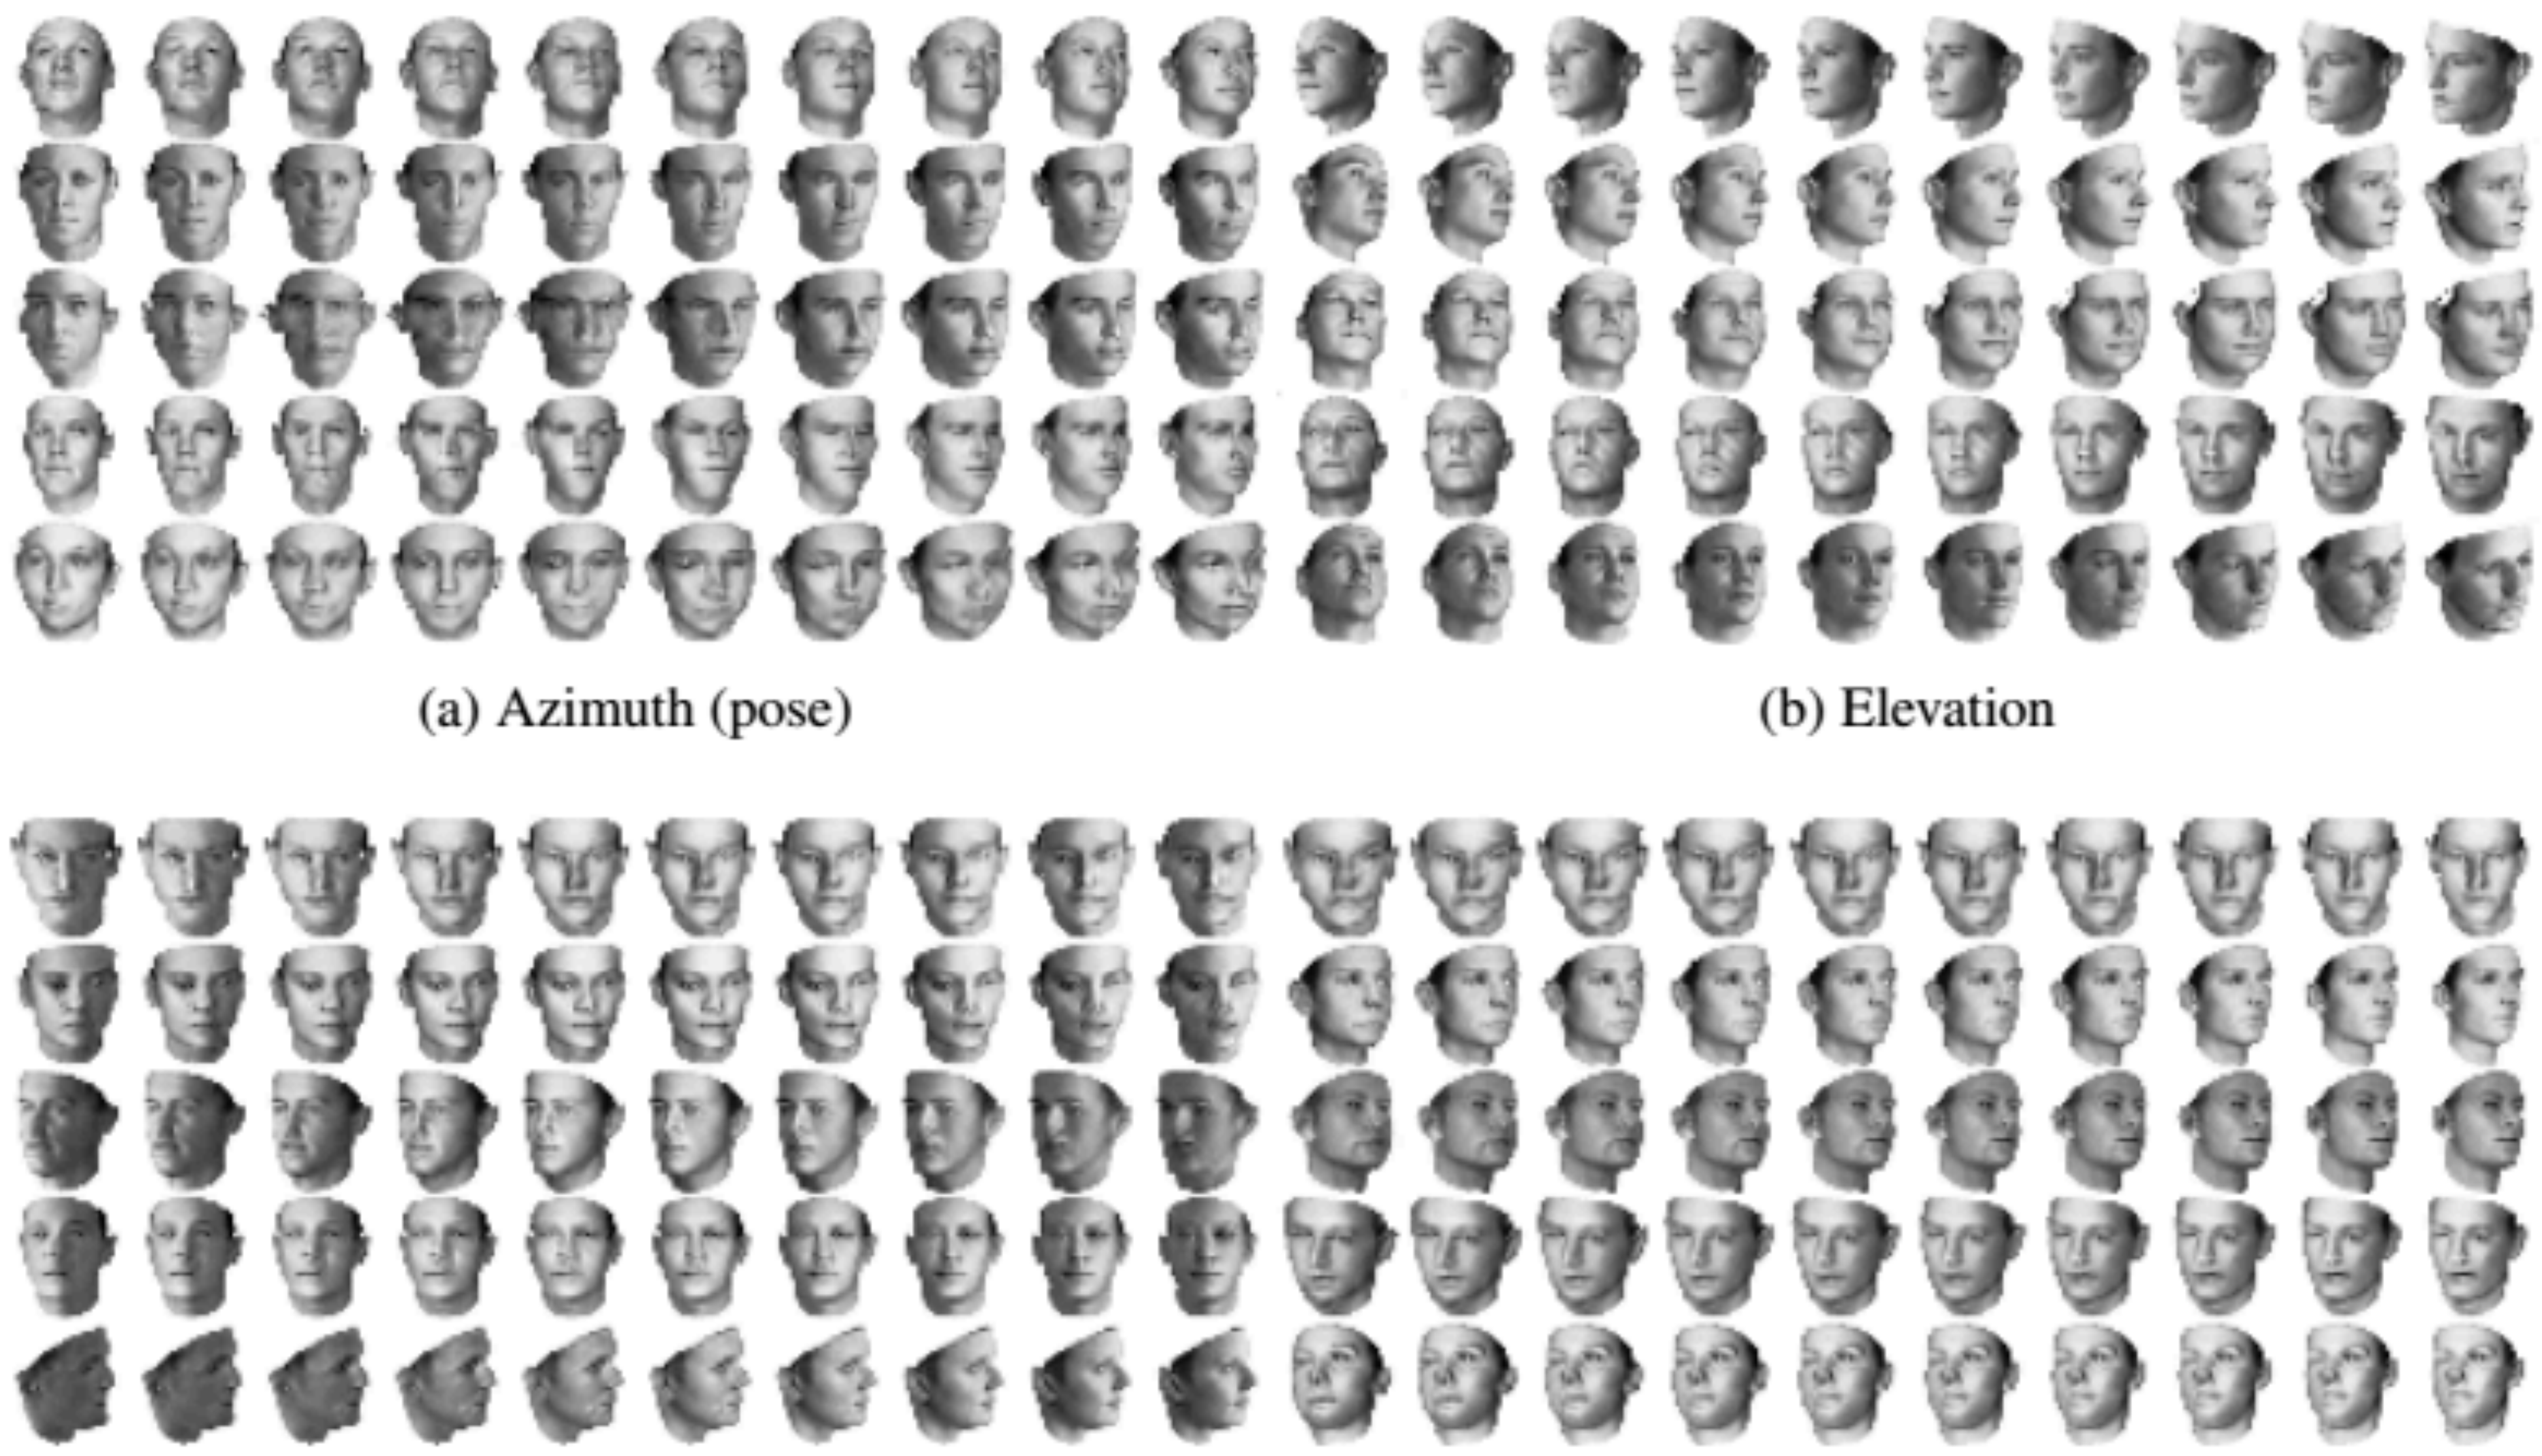

(a) Azimuth (pose) (b) Elevation

(c) Lighting (d) Wide or Narrow

InfoGAN (Chen et al 2016)

(Goodfellow 2016)

Plug and Play Generative Networks

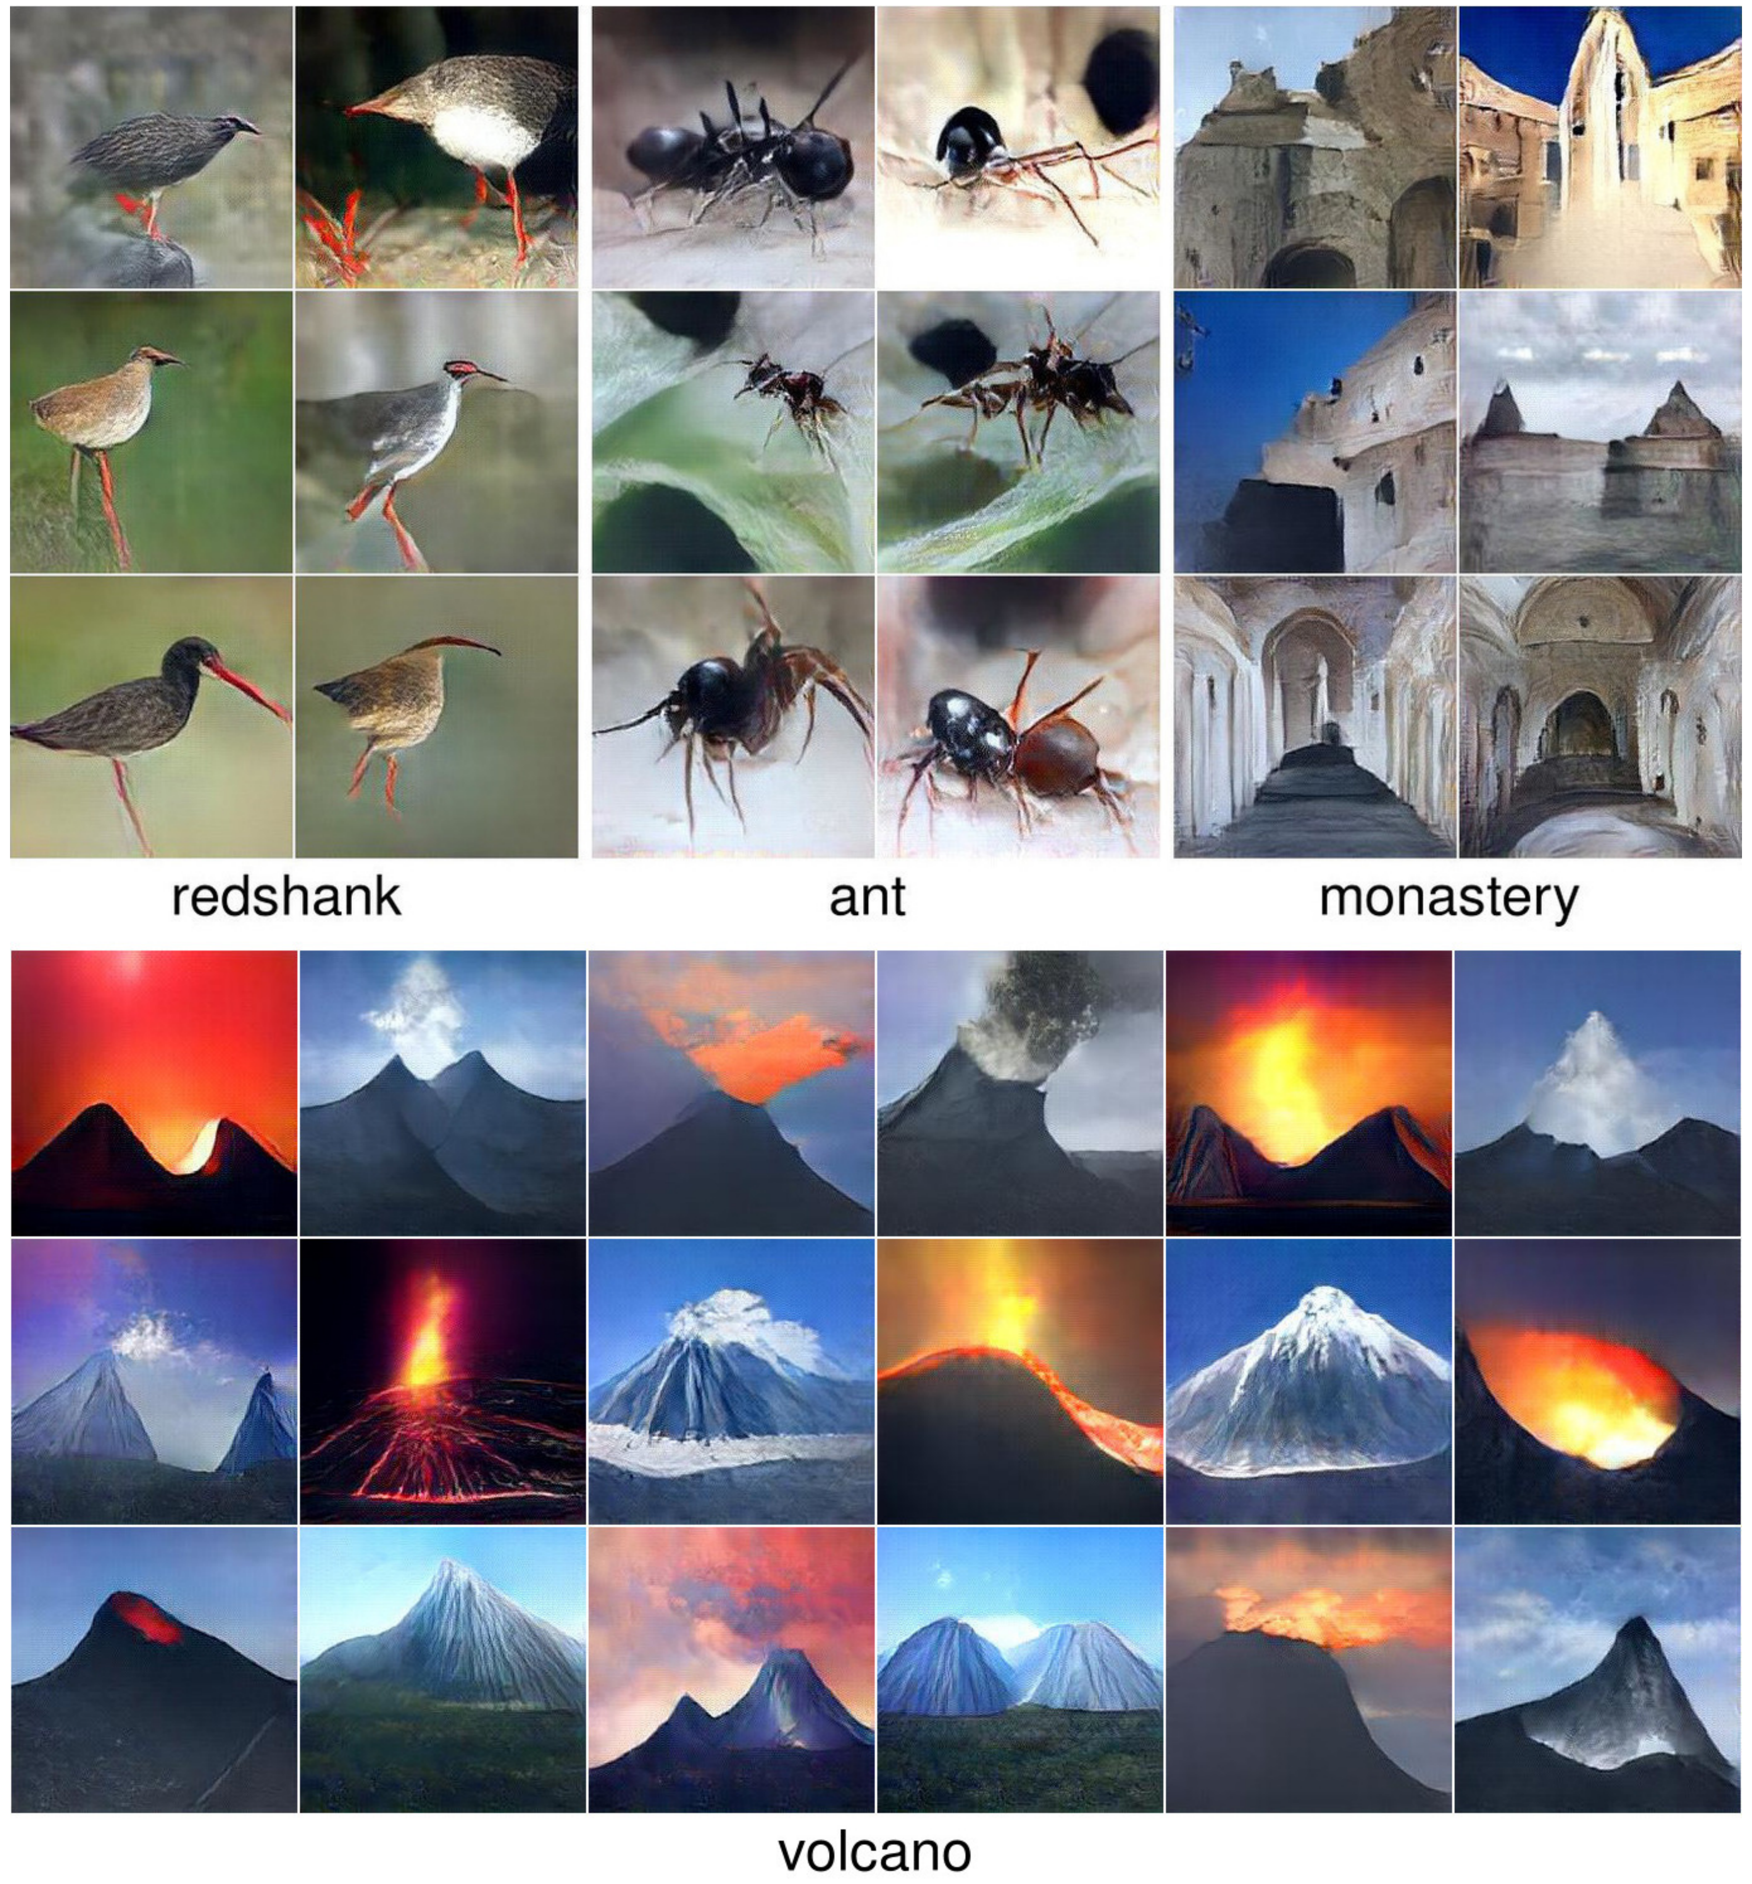

(Nguyen et al 2016)

PPGN for caption to image

oranges on a table next to a liquor bottle

(Nguyen et al 2016)

GAN loss is a key ingredient

Raw data

Reconstruction by PPGN

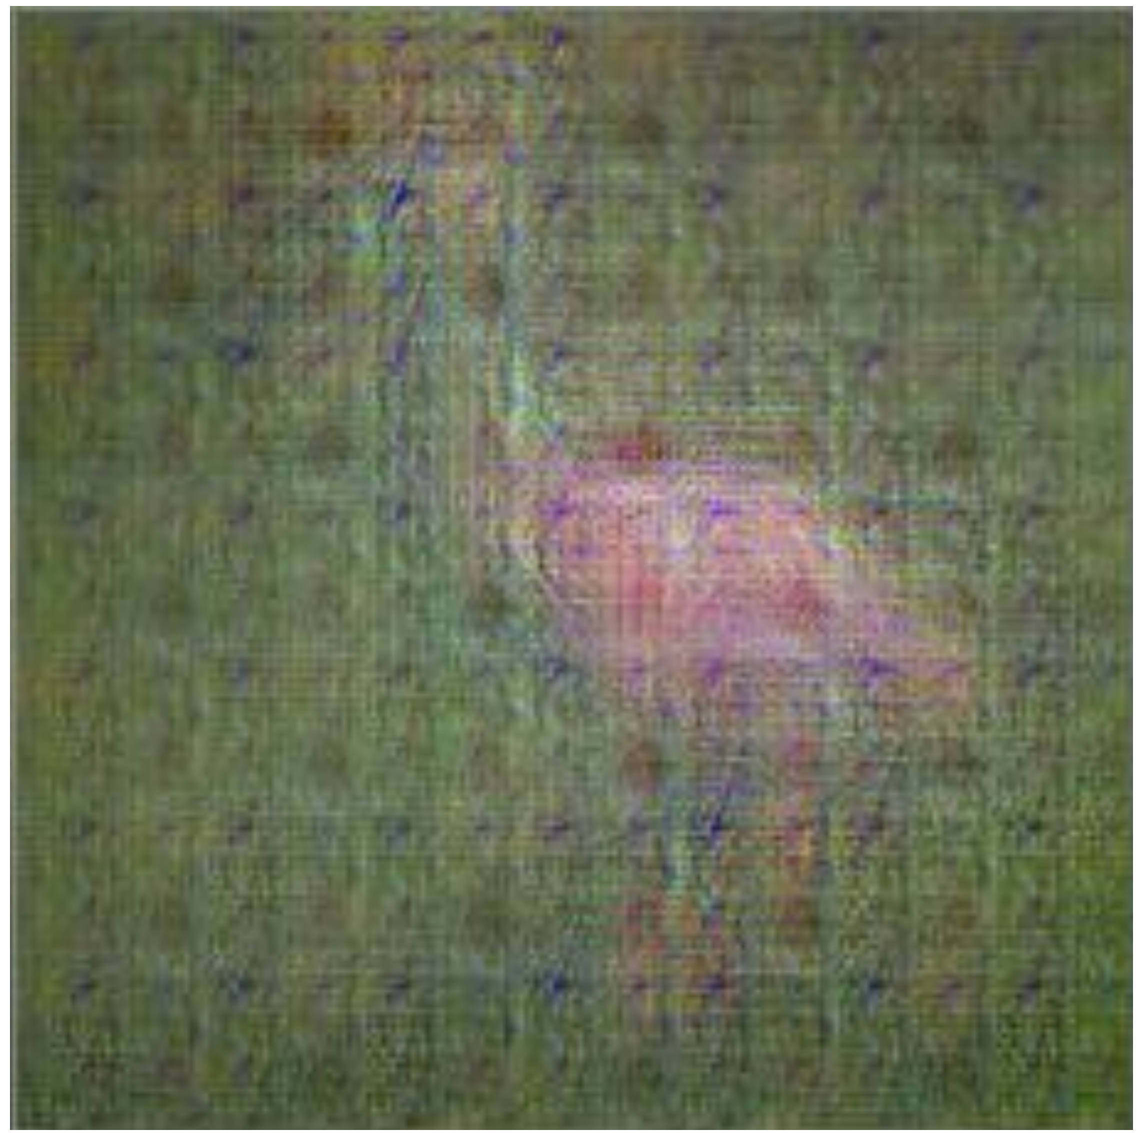

Reconstruction by PPGN without GAN

Images from Nguyen et al 2016

First observed by Dosovitskiy et al 2016

StackGANs

This small blue bird has a short pointy beak and brown on its wings

This bird is completely red with black wings and pointy beak

A small sized bird that has a cream belly and a short pointed bill

A small bird with a black head and wings and features grey wings

(Zhang et al 2016)

Conclusion

- GANs produce rich, realistic imagery
- GANs learn to draw samples from a probability distribution
- Applications include learning from very few labeled examples, interactive artwork generation, and differential privacy

(Goodfellow 2016)