GANs for Limited Labeled Data

Ian Goodfellow, Staff Research Scientist, Google Brain

NIPS 2017 Workshop on Limited Labeled Data: Weak Supervision and Beyond

Long Beach, 2017-12-09

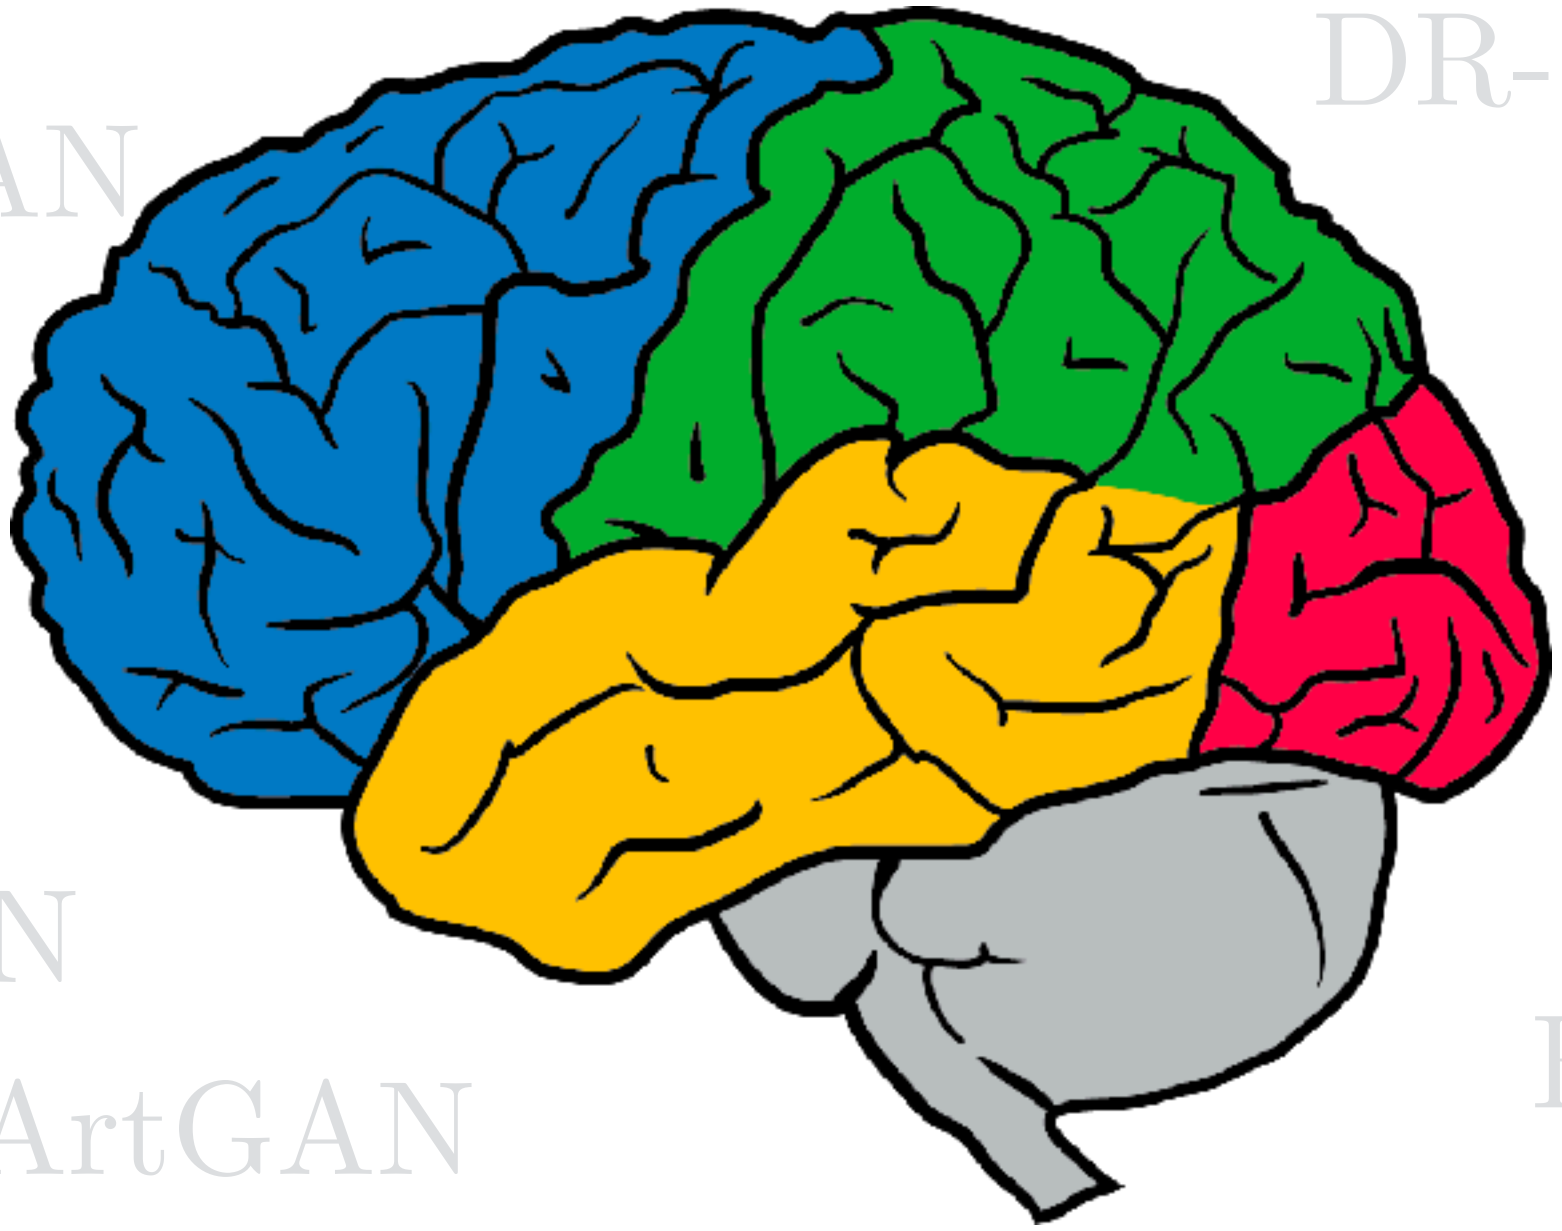

MedGAN Progressive GAN CoGAN LR-GAN CGAN IcGAN
b-GAN LS-GAN AffGAN LAPGAN DiscoGAN MPM-GAN AdaGAN
LSGAN InfoGAN CatGAN AMGAN iGAN IAN
McGAN MIX+GAN
MGAN FF-GAN GoGAN BS-GAN
C-VAE-GAN C-RNN-GAN DR-GAN DCGAN
MAGAN 3D-GAN CCGAN AC-GAN
GAWWN DualGAN CycleGAN BiGAN
Bayesian GAN GP-GAN
EBGAN SN-GAN AnoGAN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

Adversarial Nets Framework

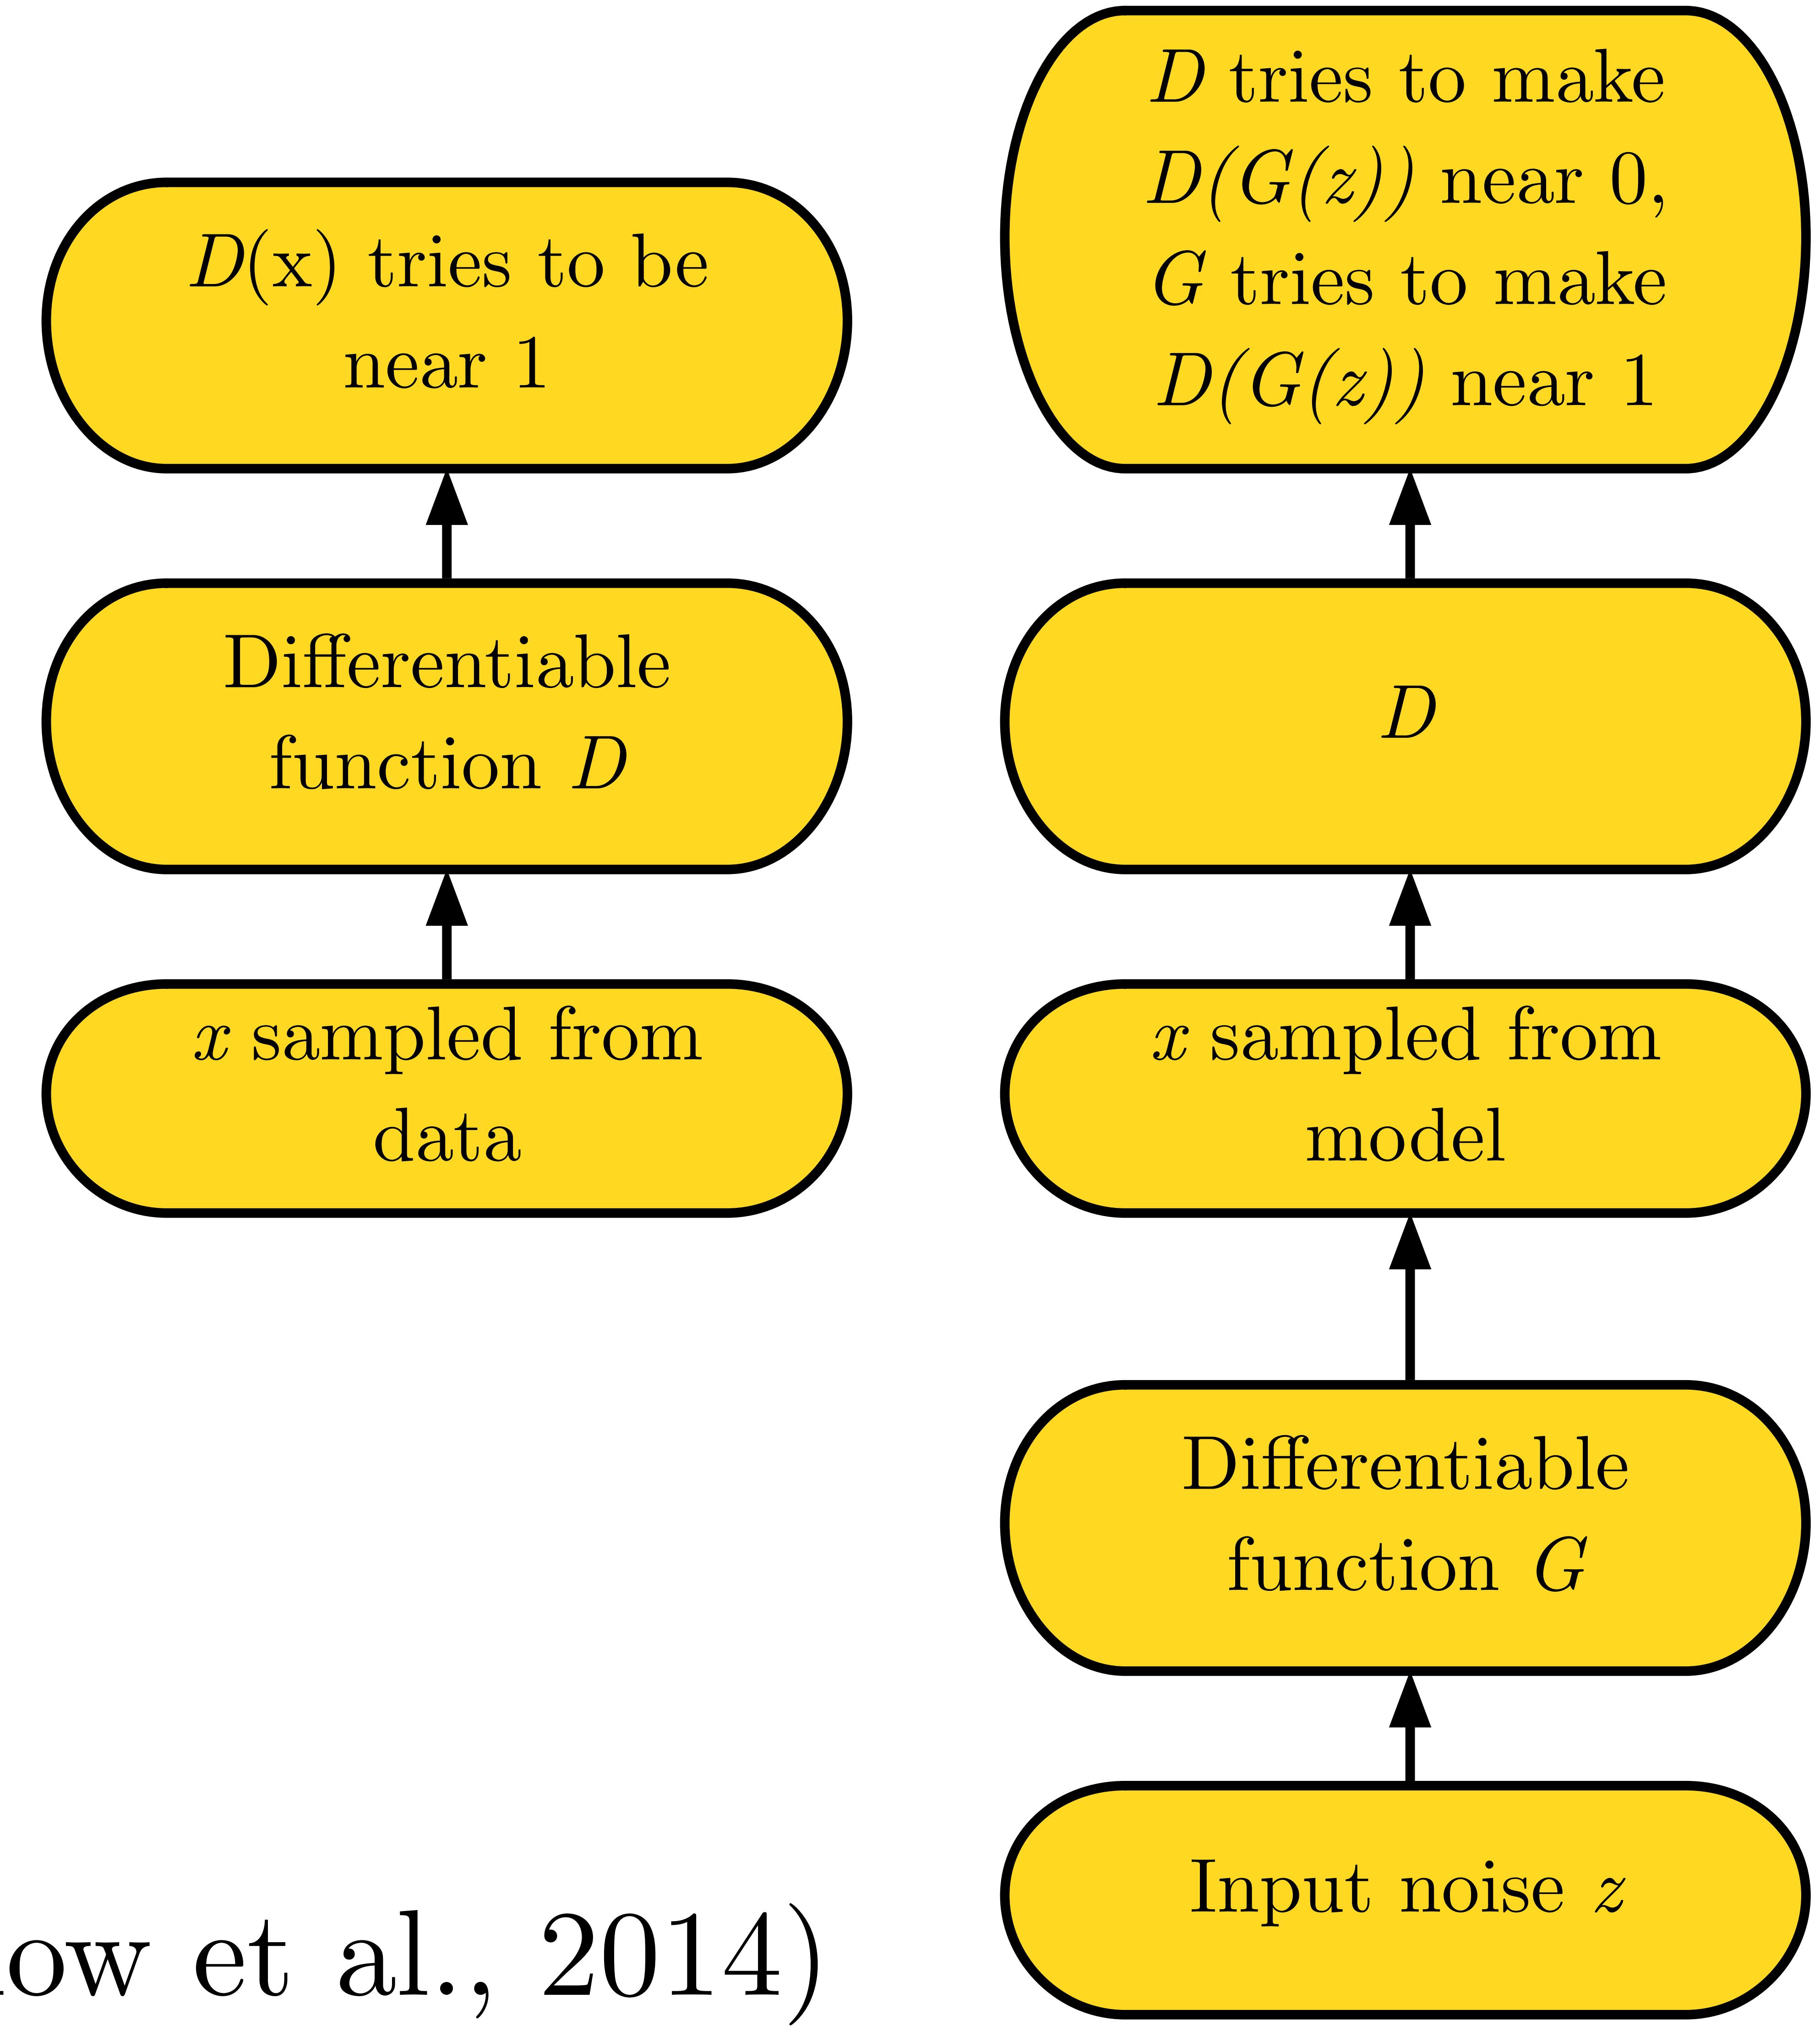

(Goodfellow et al., 2014)

(Goodfellow 2017)

Overcoming limited data with GANs

- Missing data
  - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

What is in this image?

(Yeh et al., 2016)

Generative modeling reveals a face

(Yeh et al., 2016)

Overcoming limited data with GANs

- Missing data
  - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

Supervised Discriminator

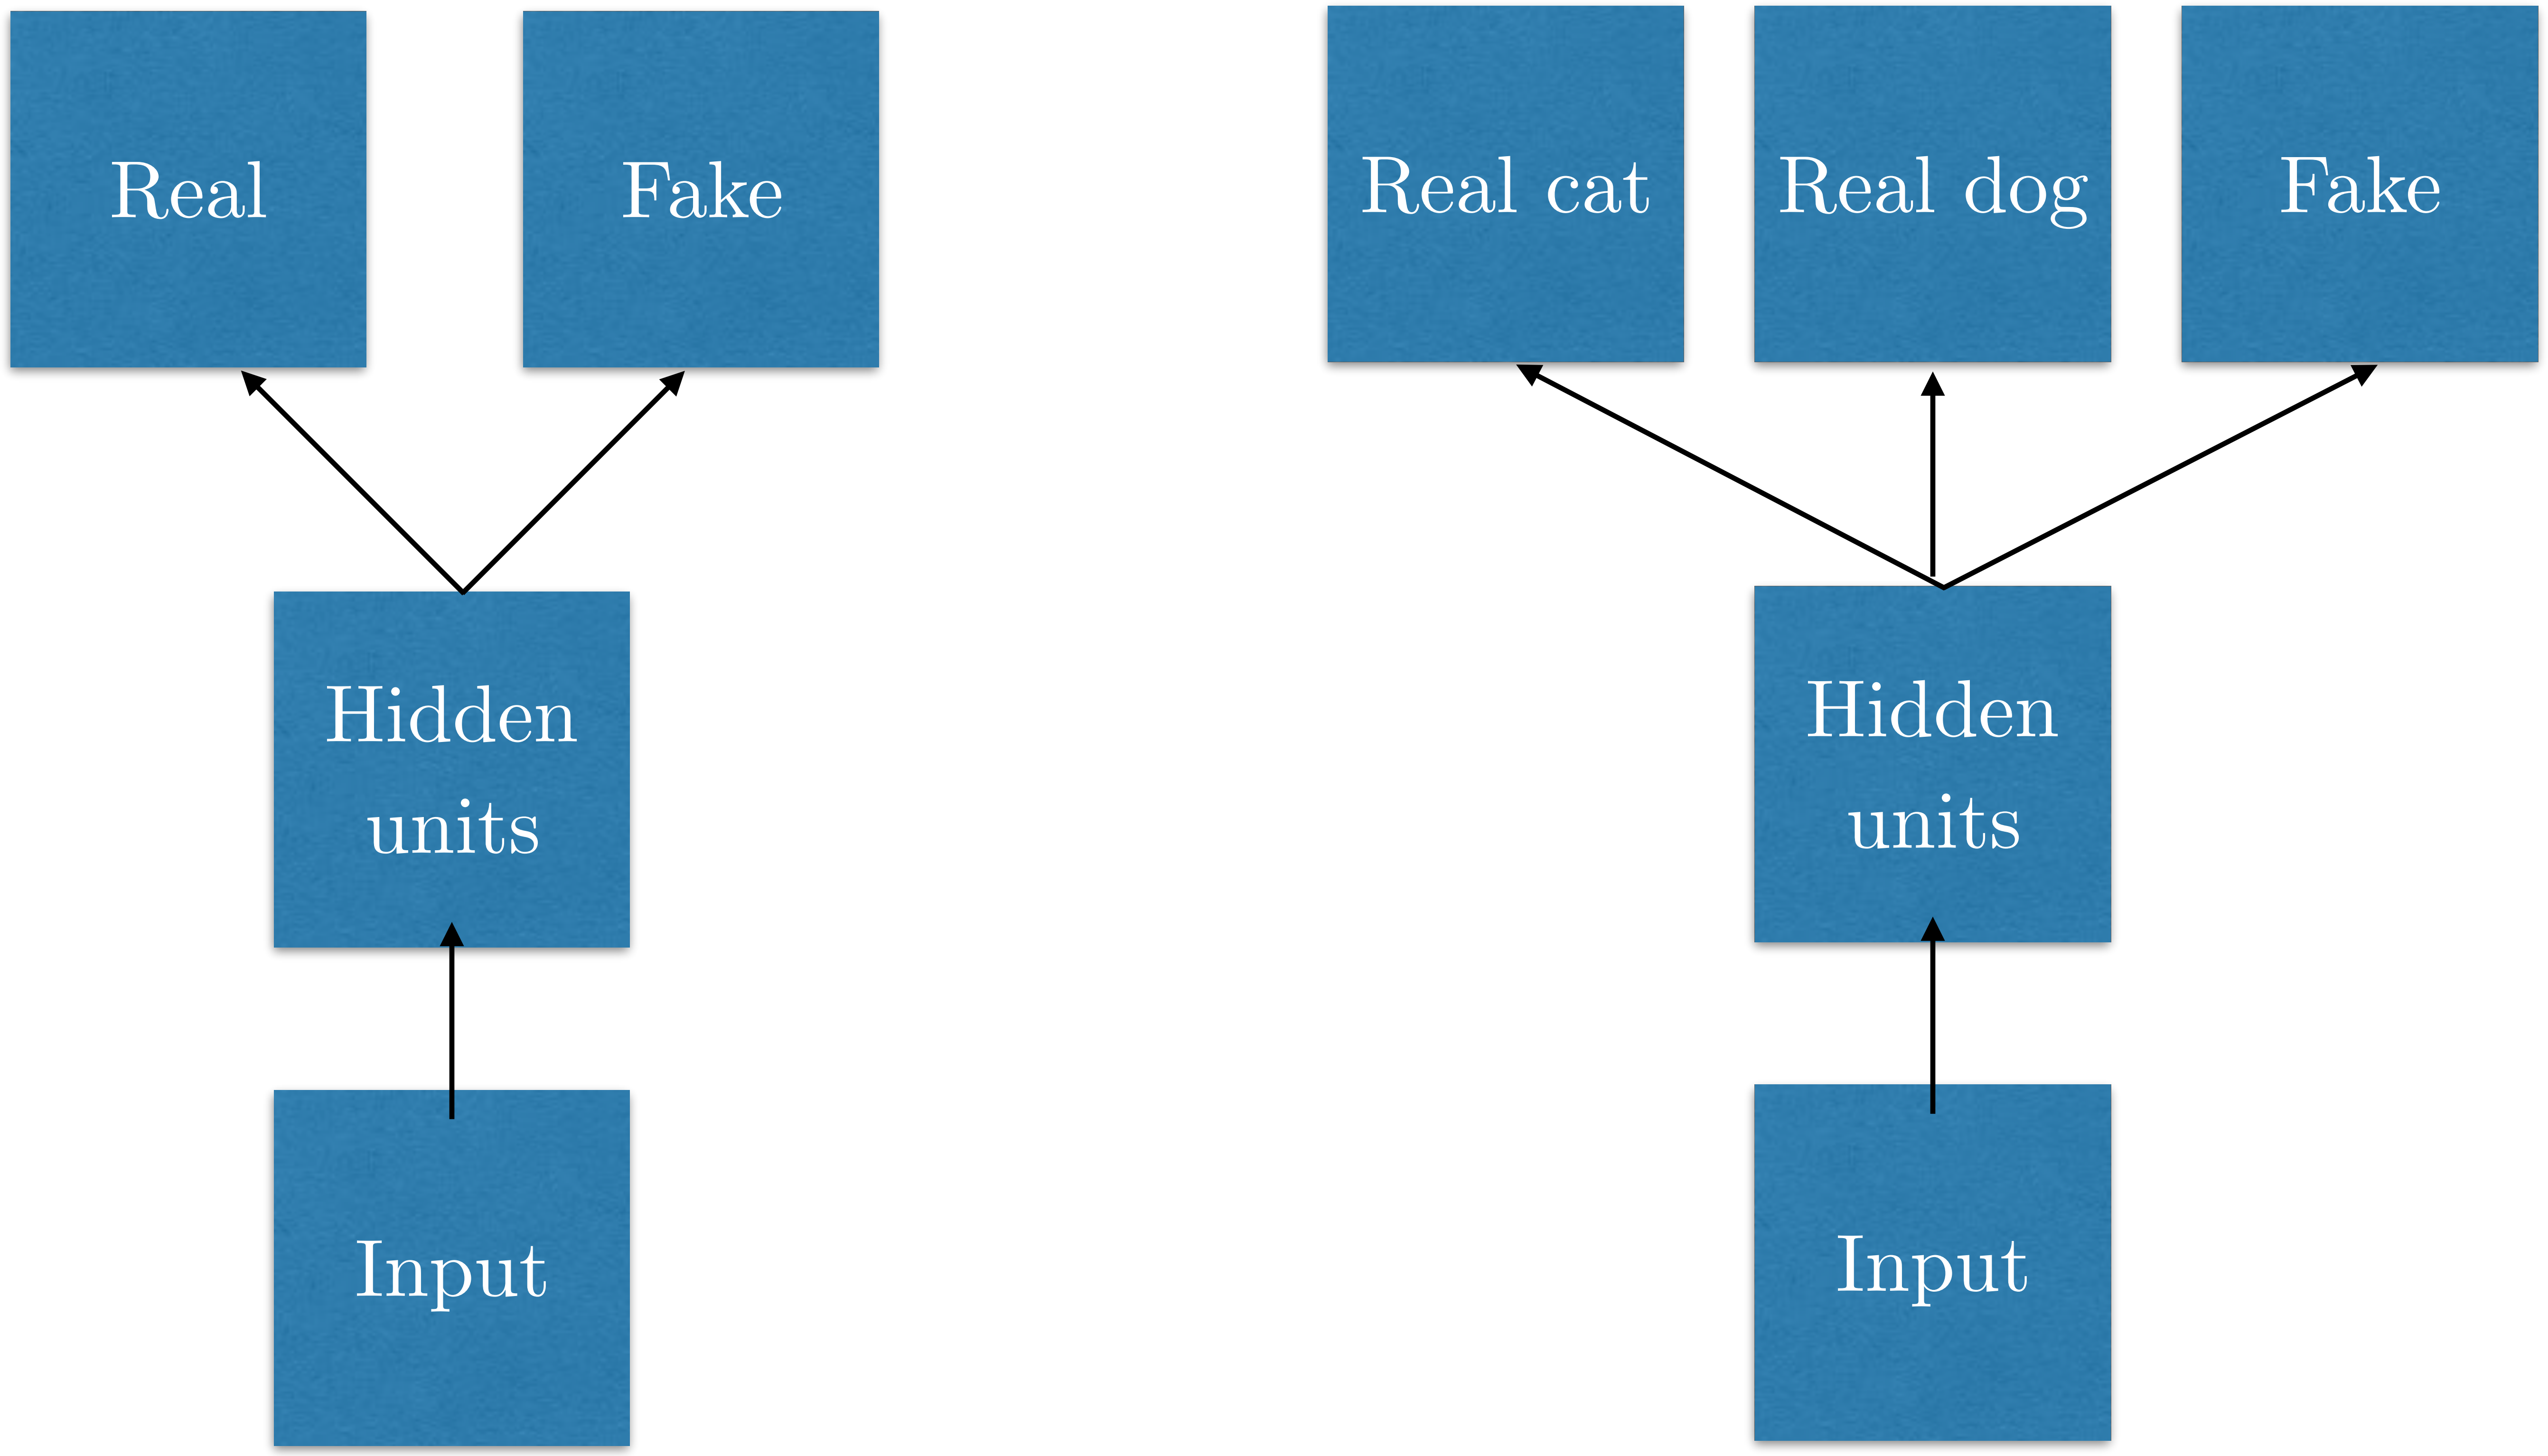

(Odena 2016, Salimans et al 2016)

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)

Useful for differential privacy: Papernot et al, 2016

Overcoming limited data with GANs

- Missing data
    - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

(Goodfellow 2017)

Next Video Frame Prediction

Ground Truth

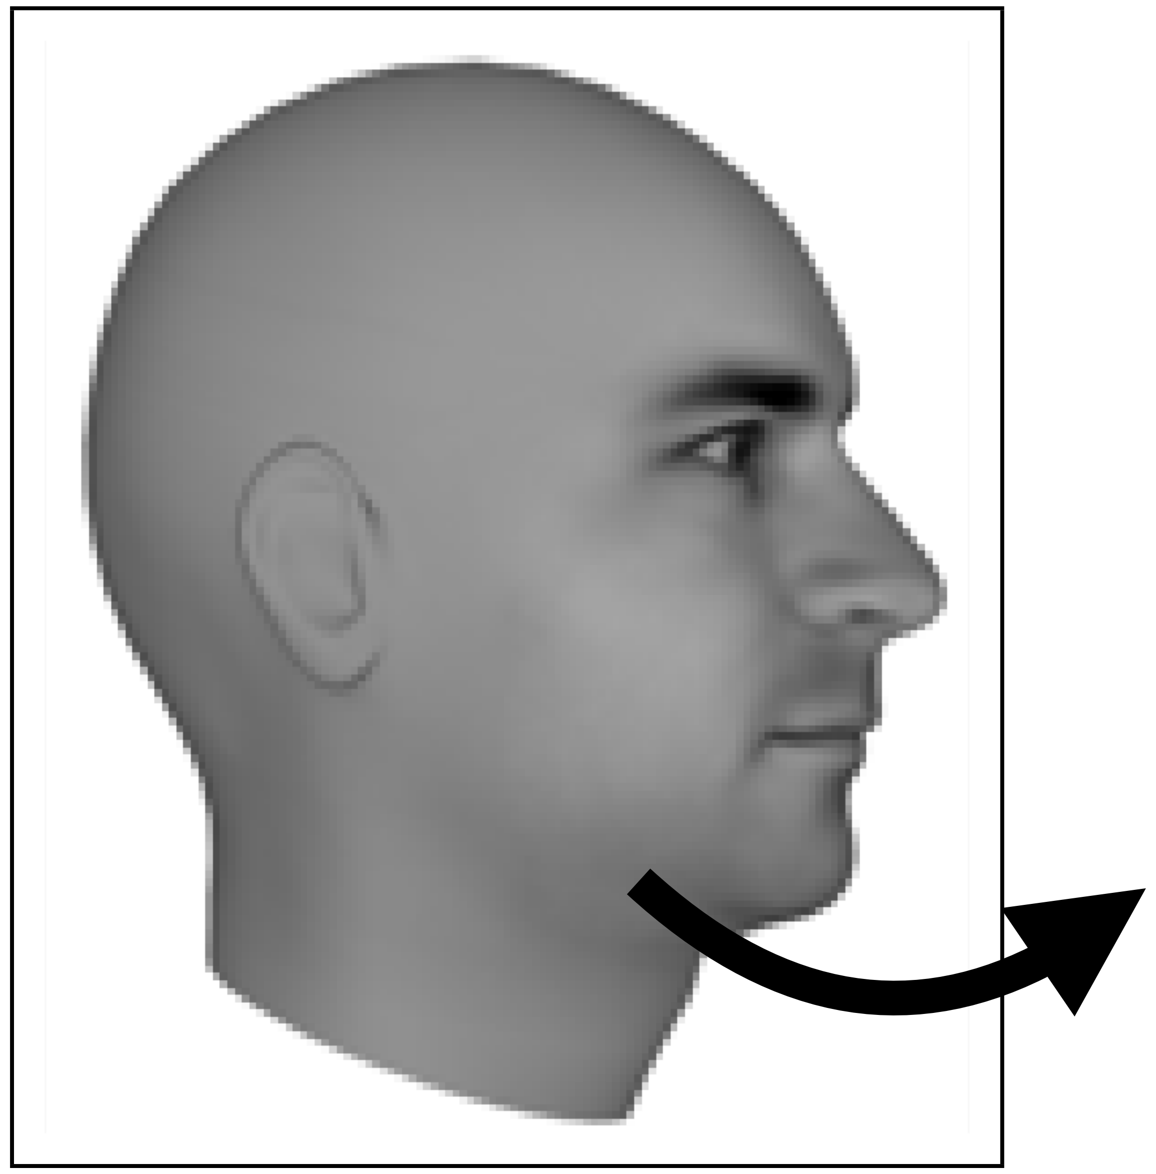

What happens next?

(Lotter et al 2016)

Next Video Frame Prediction

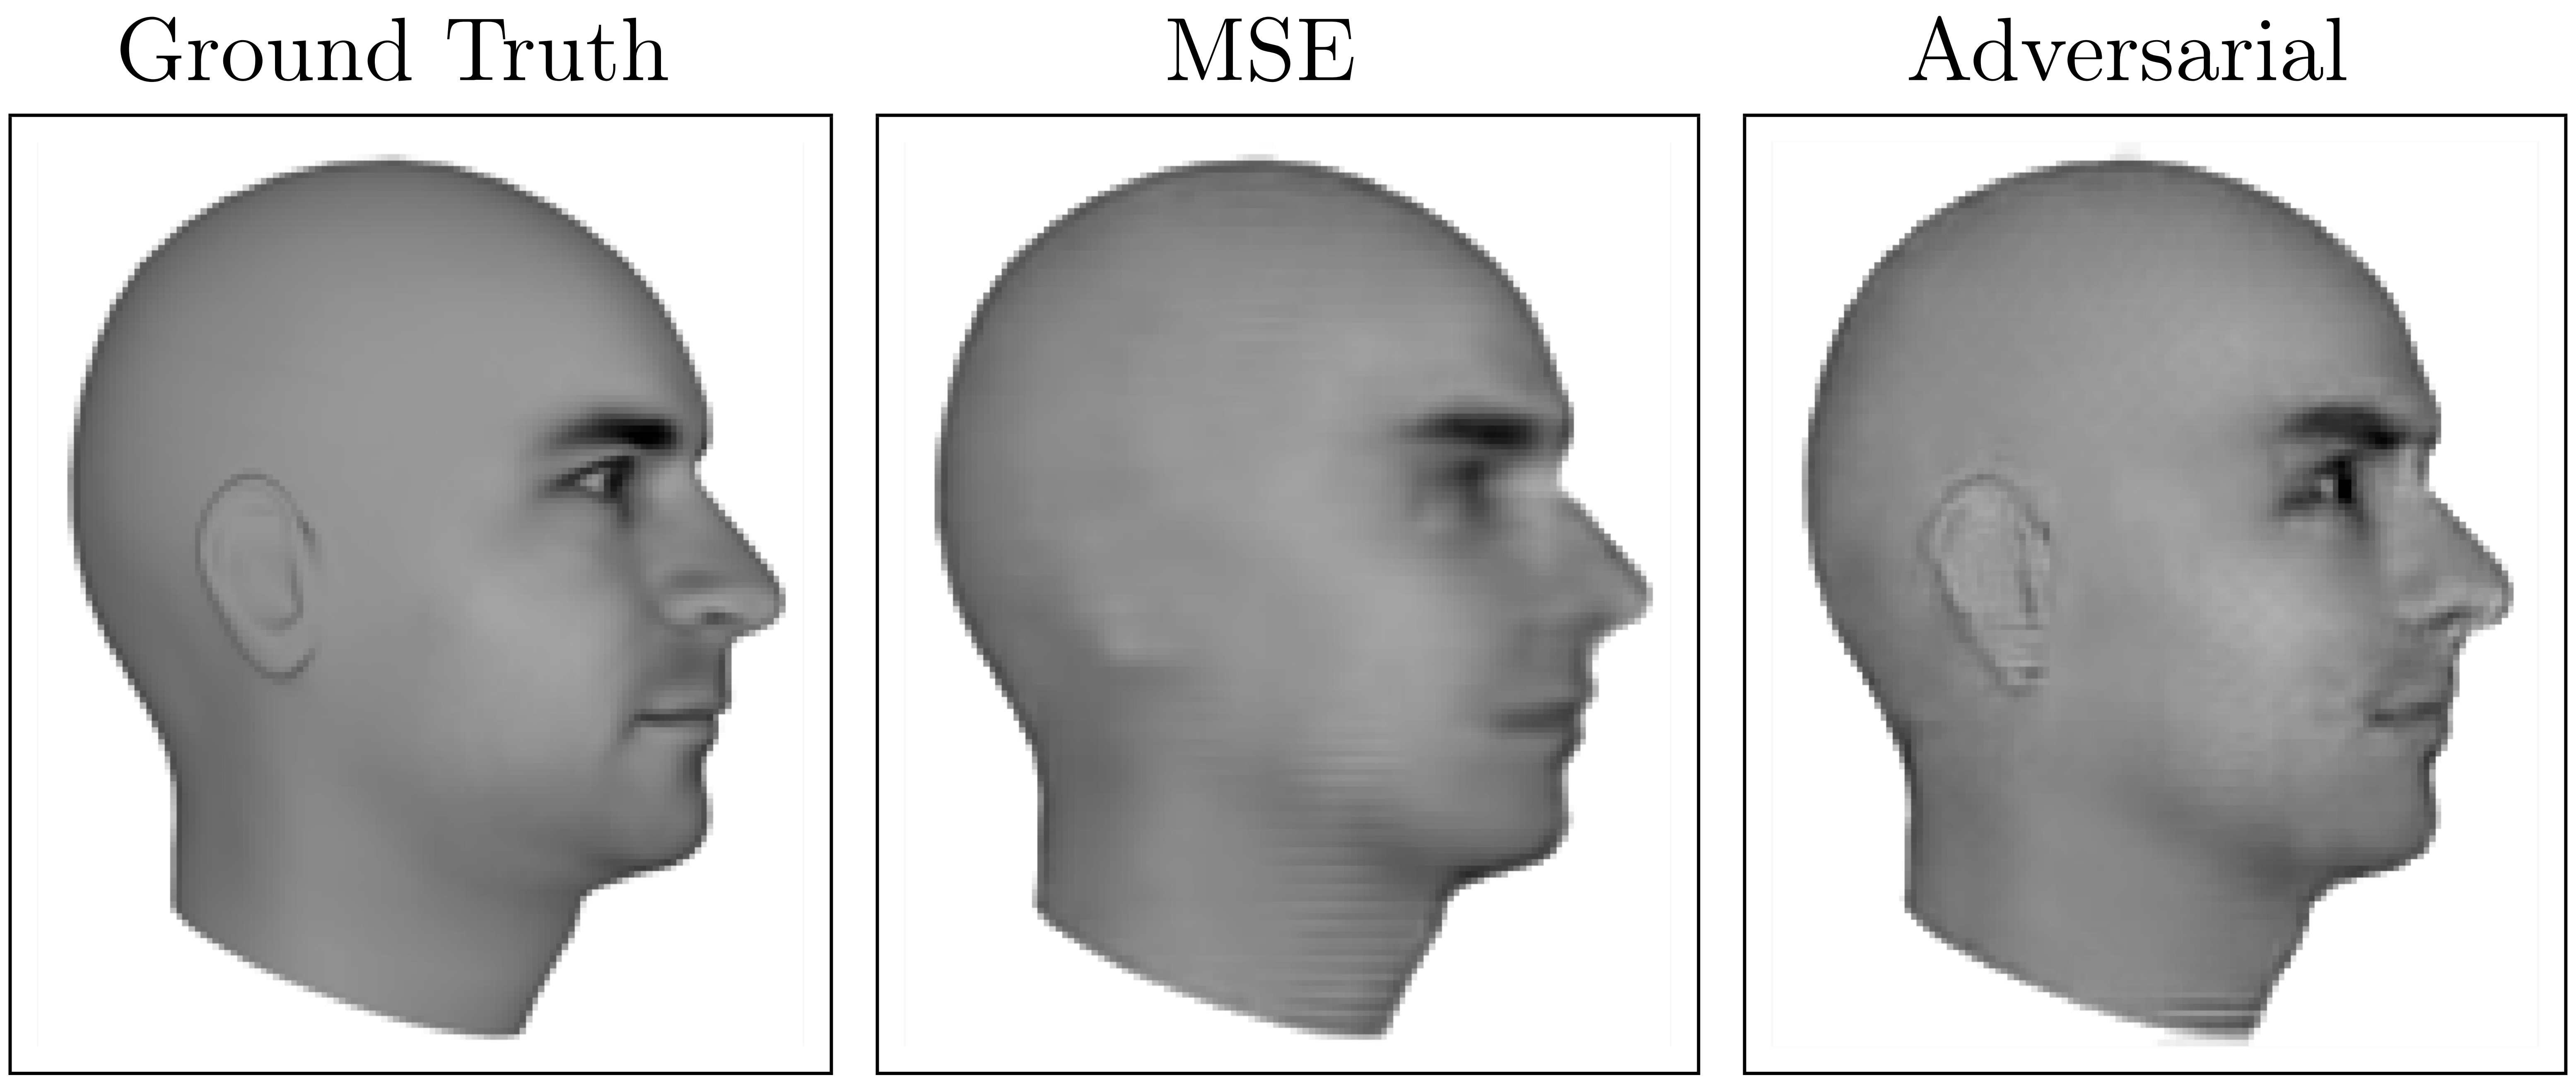

(Lotter et al 2016)

Next Video Frame(s) Prediction

Mean Squared Error

Mean Absolute Error

Adversarial

(Mathieu et al. 2015)

Overcoming limited data with GANs

- Missing data
  - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

Unsupervised Image-to-Image Translation

Day to night

(Liu et al., 2017)

CycleGAN

(Zhu et al., 2017)

Translation without parallel corpora

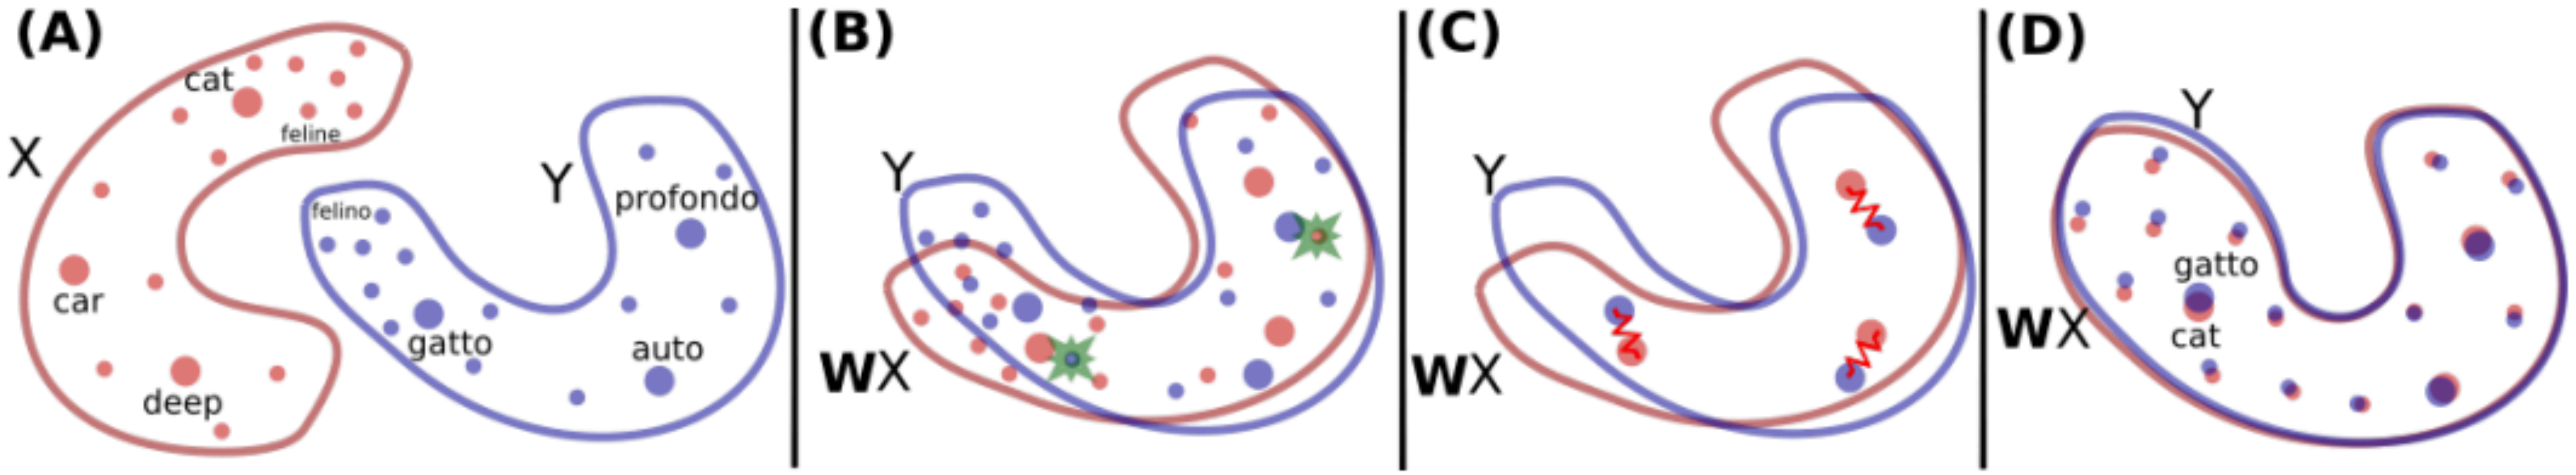

(Conneau et al., 2017)

Overcoming limited data with GANs

- Missing data
  - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

Simulating particle physics

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

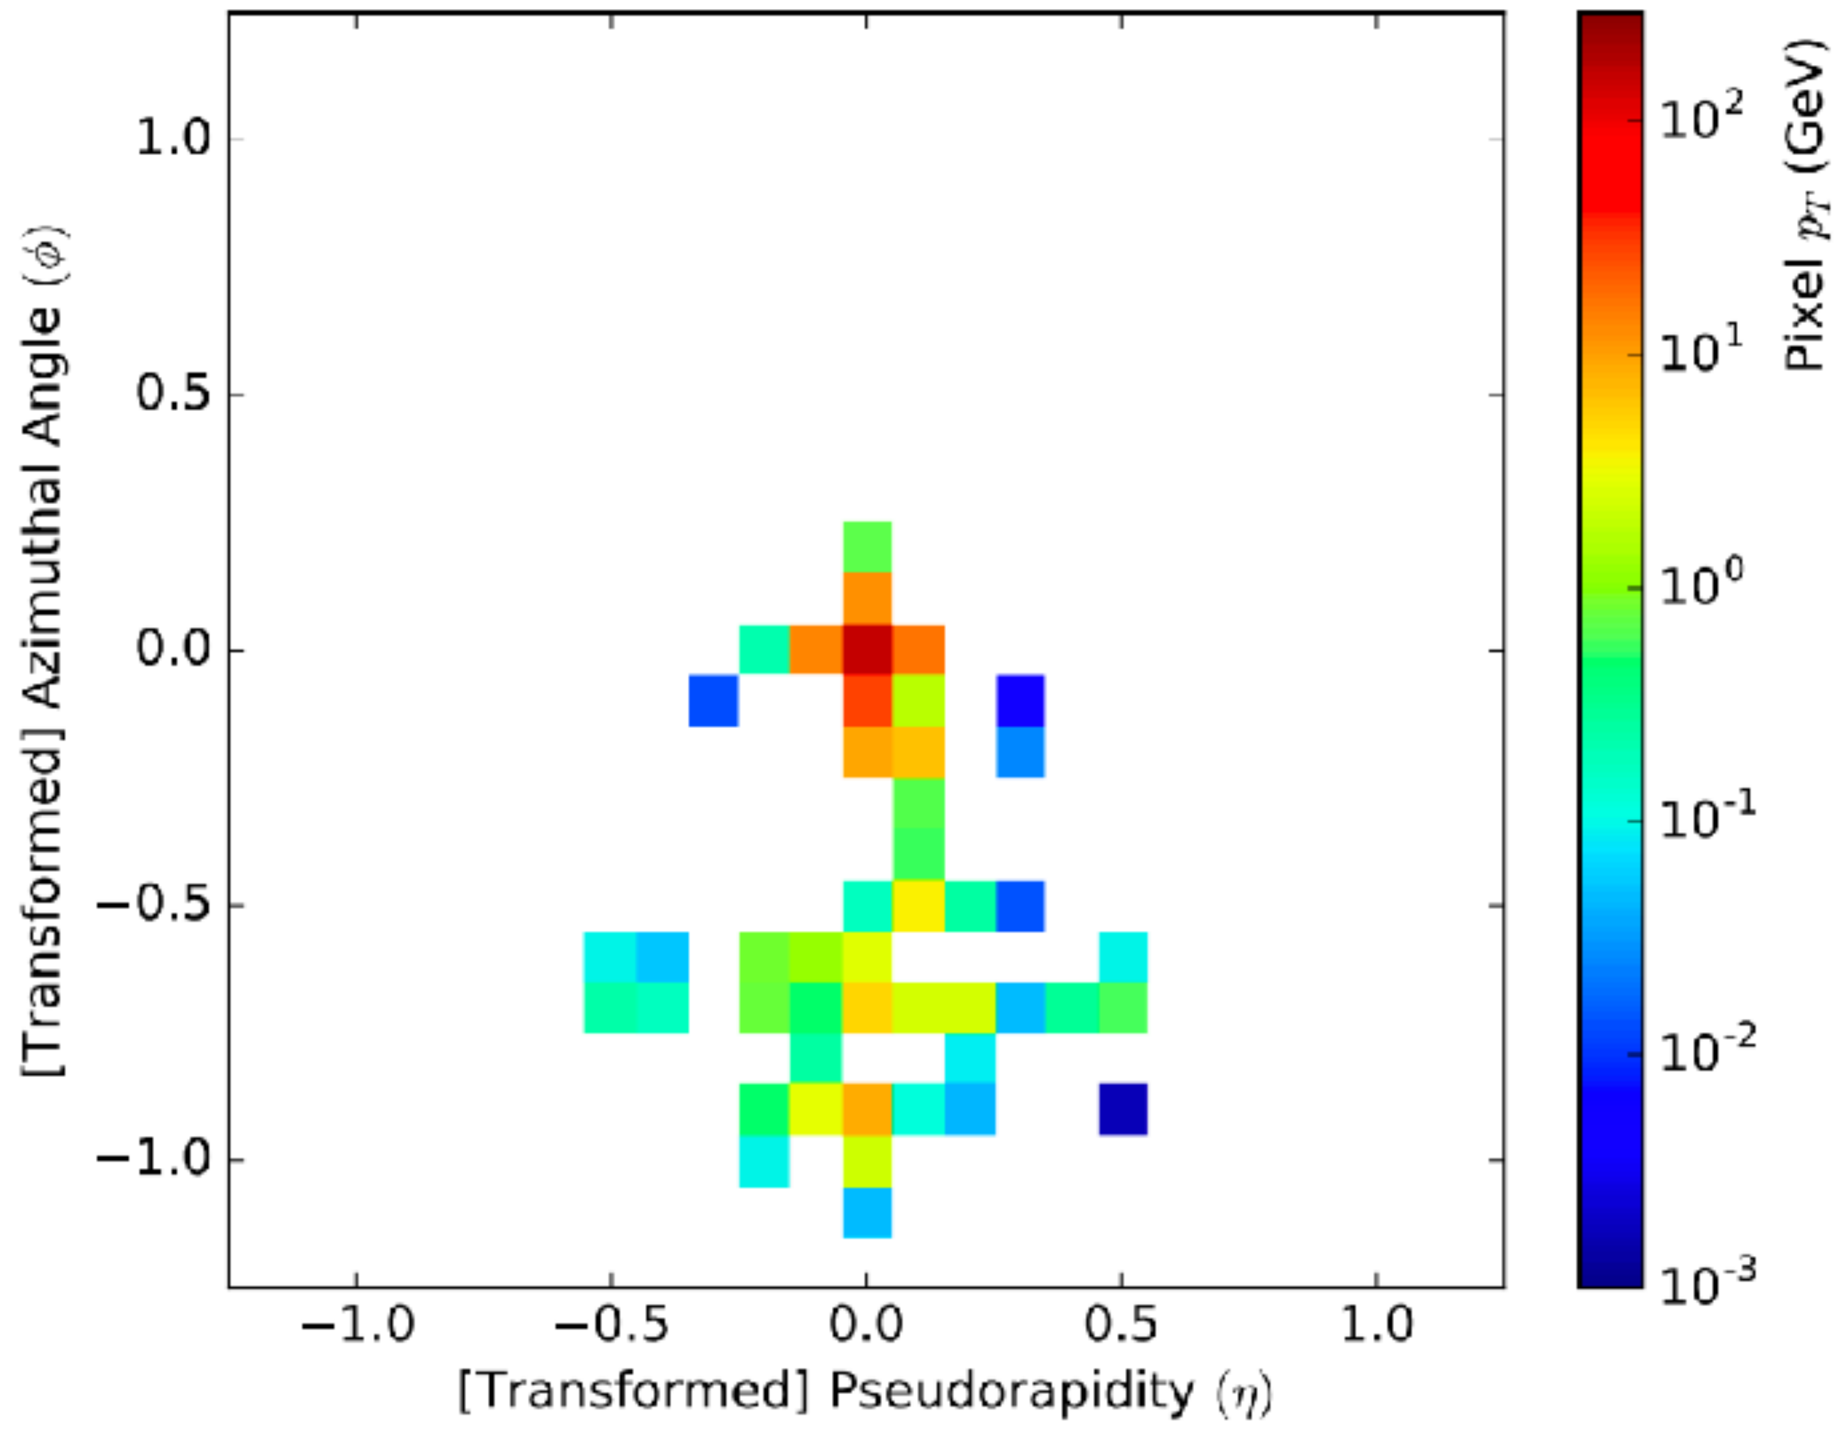

(de Oliveira et al., 2017)

Overcoming limited data with GANs

- Missing data
    - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

(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

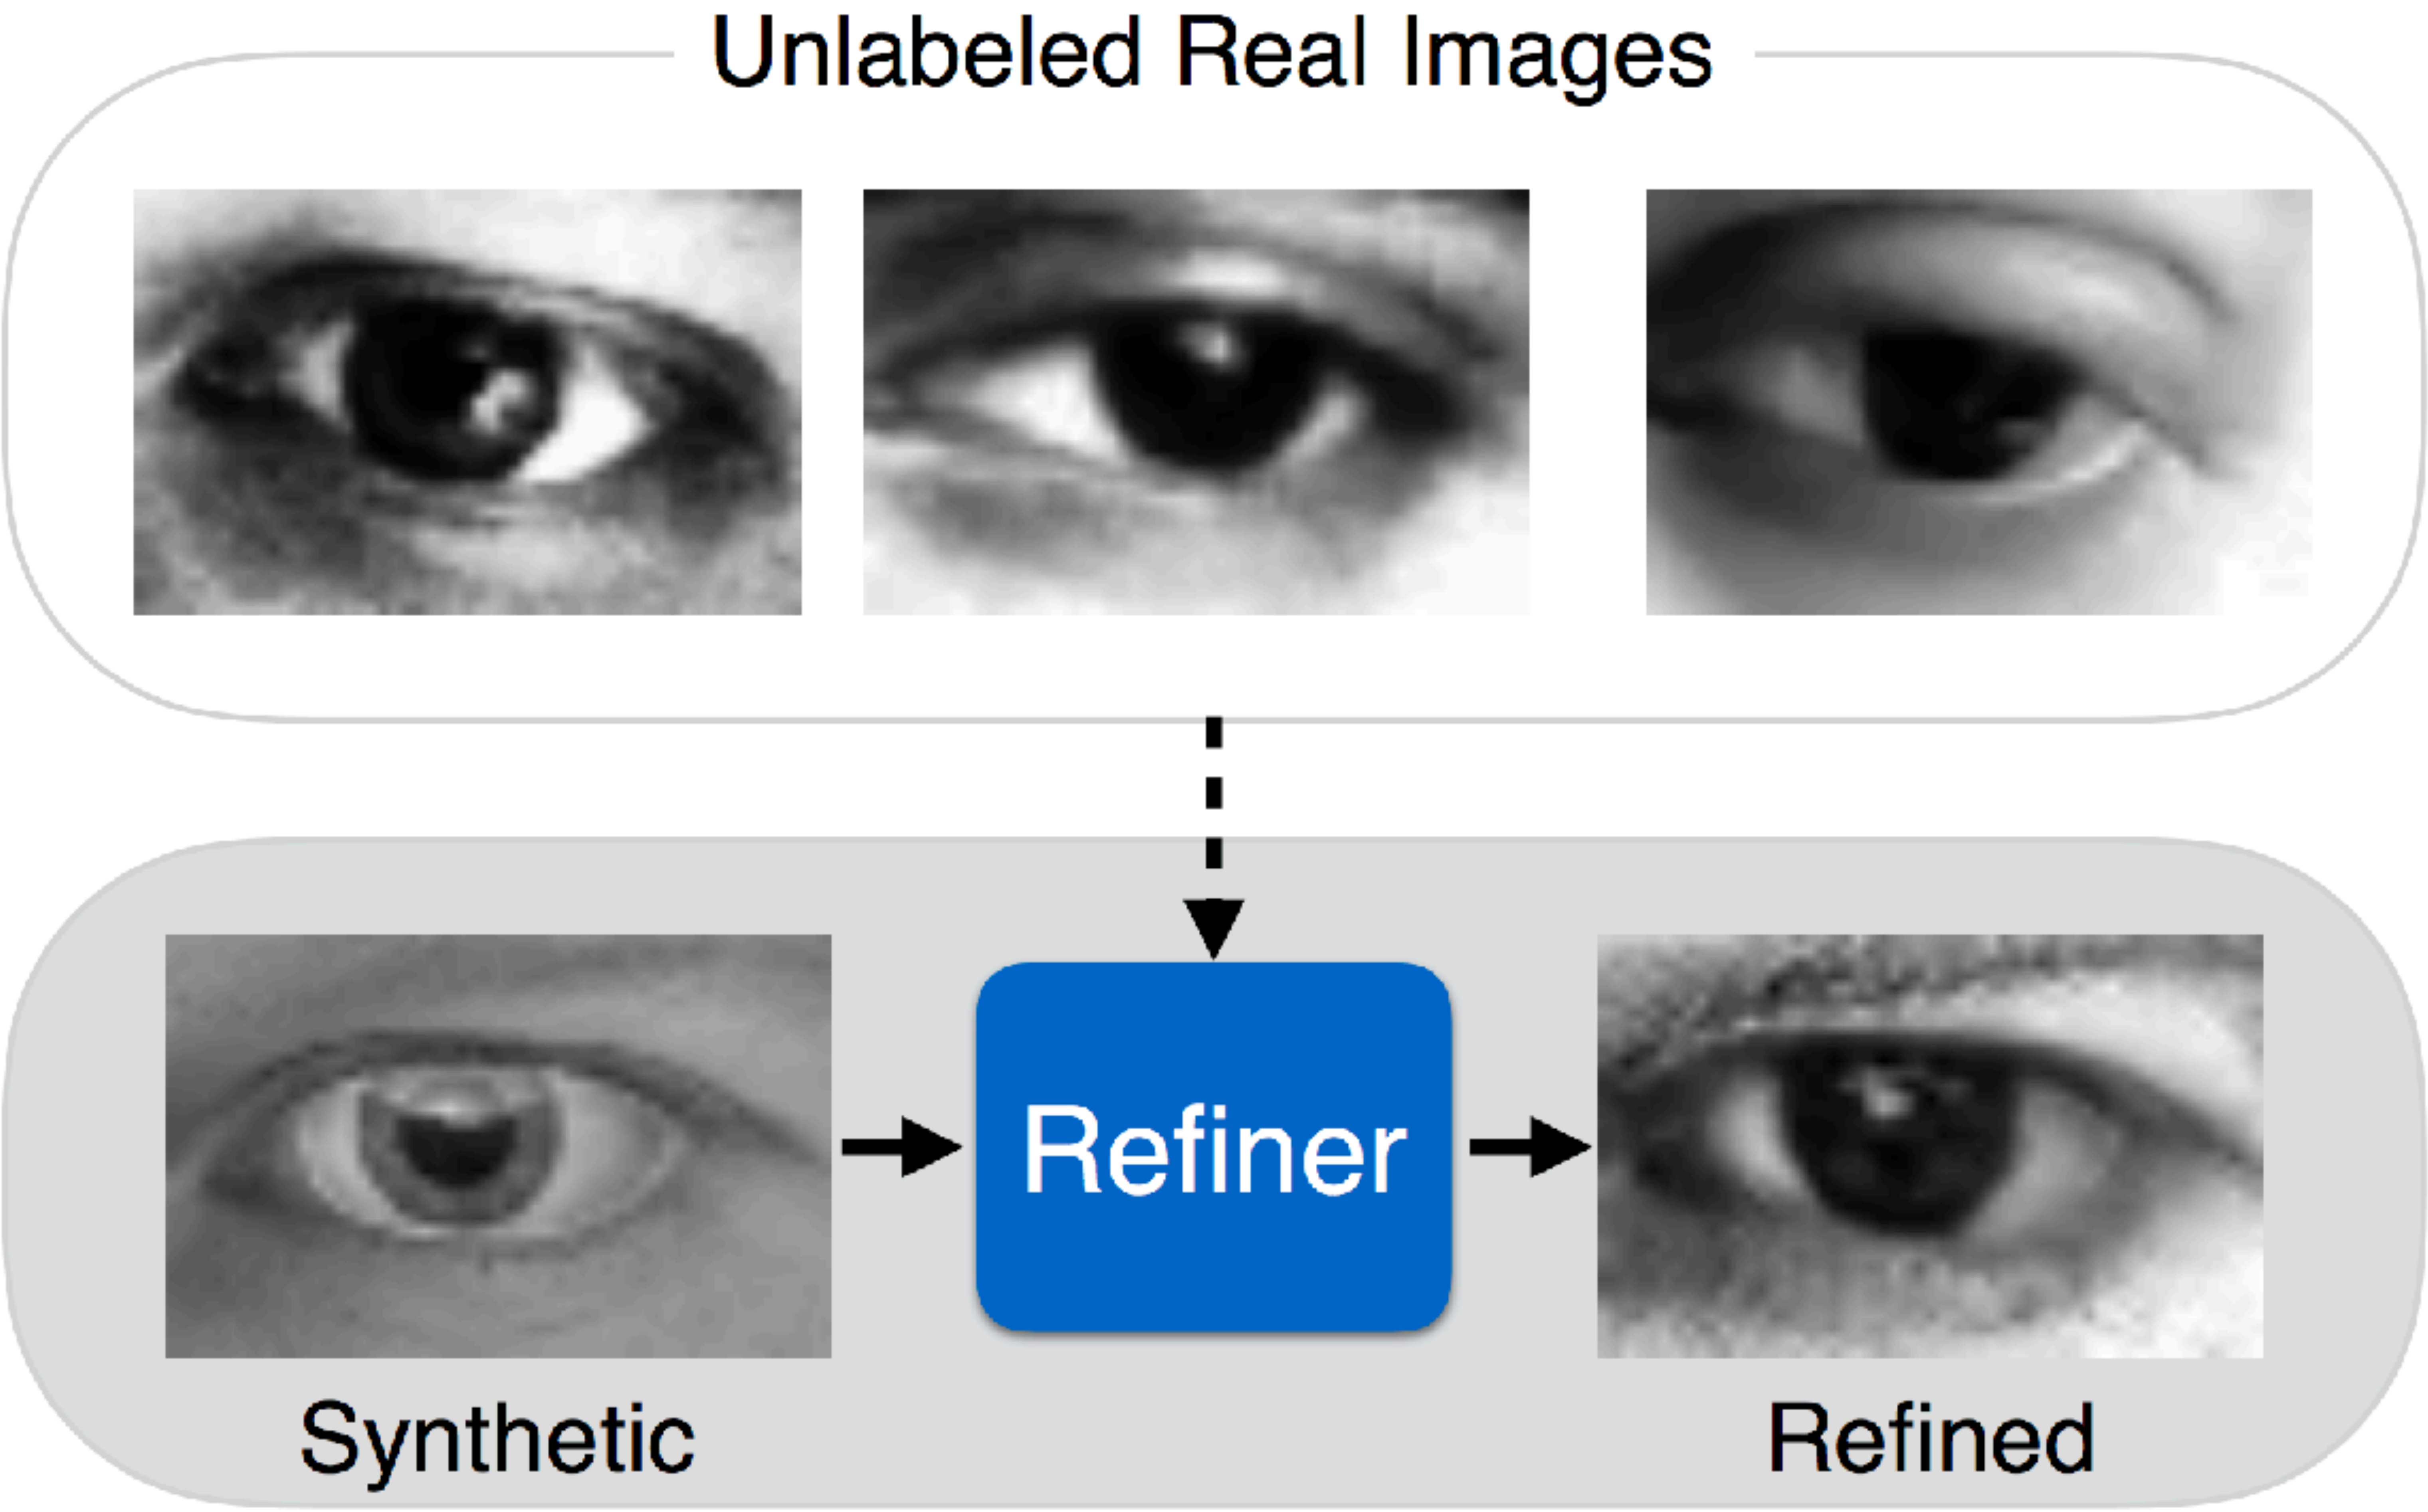

(Shrivastava et al., 2016)

GANs for domain adaptation

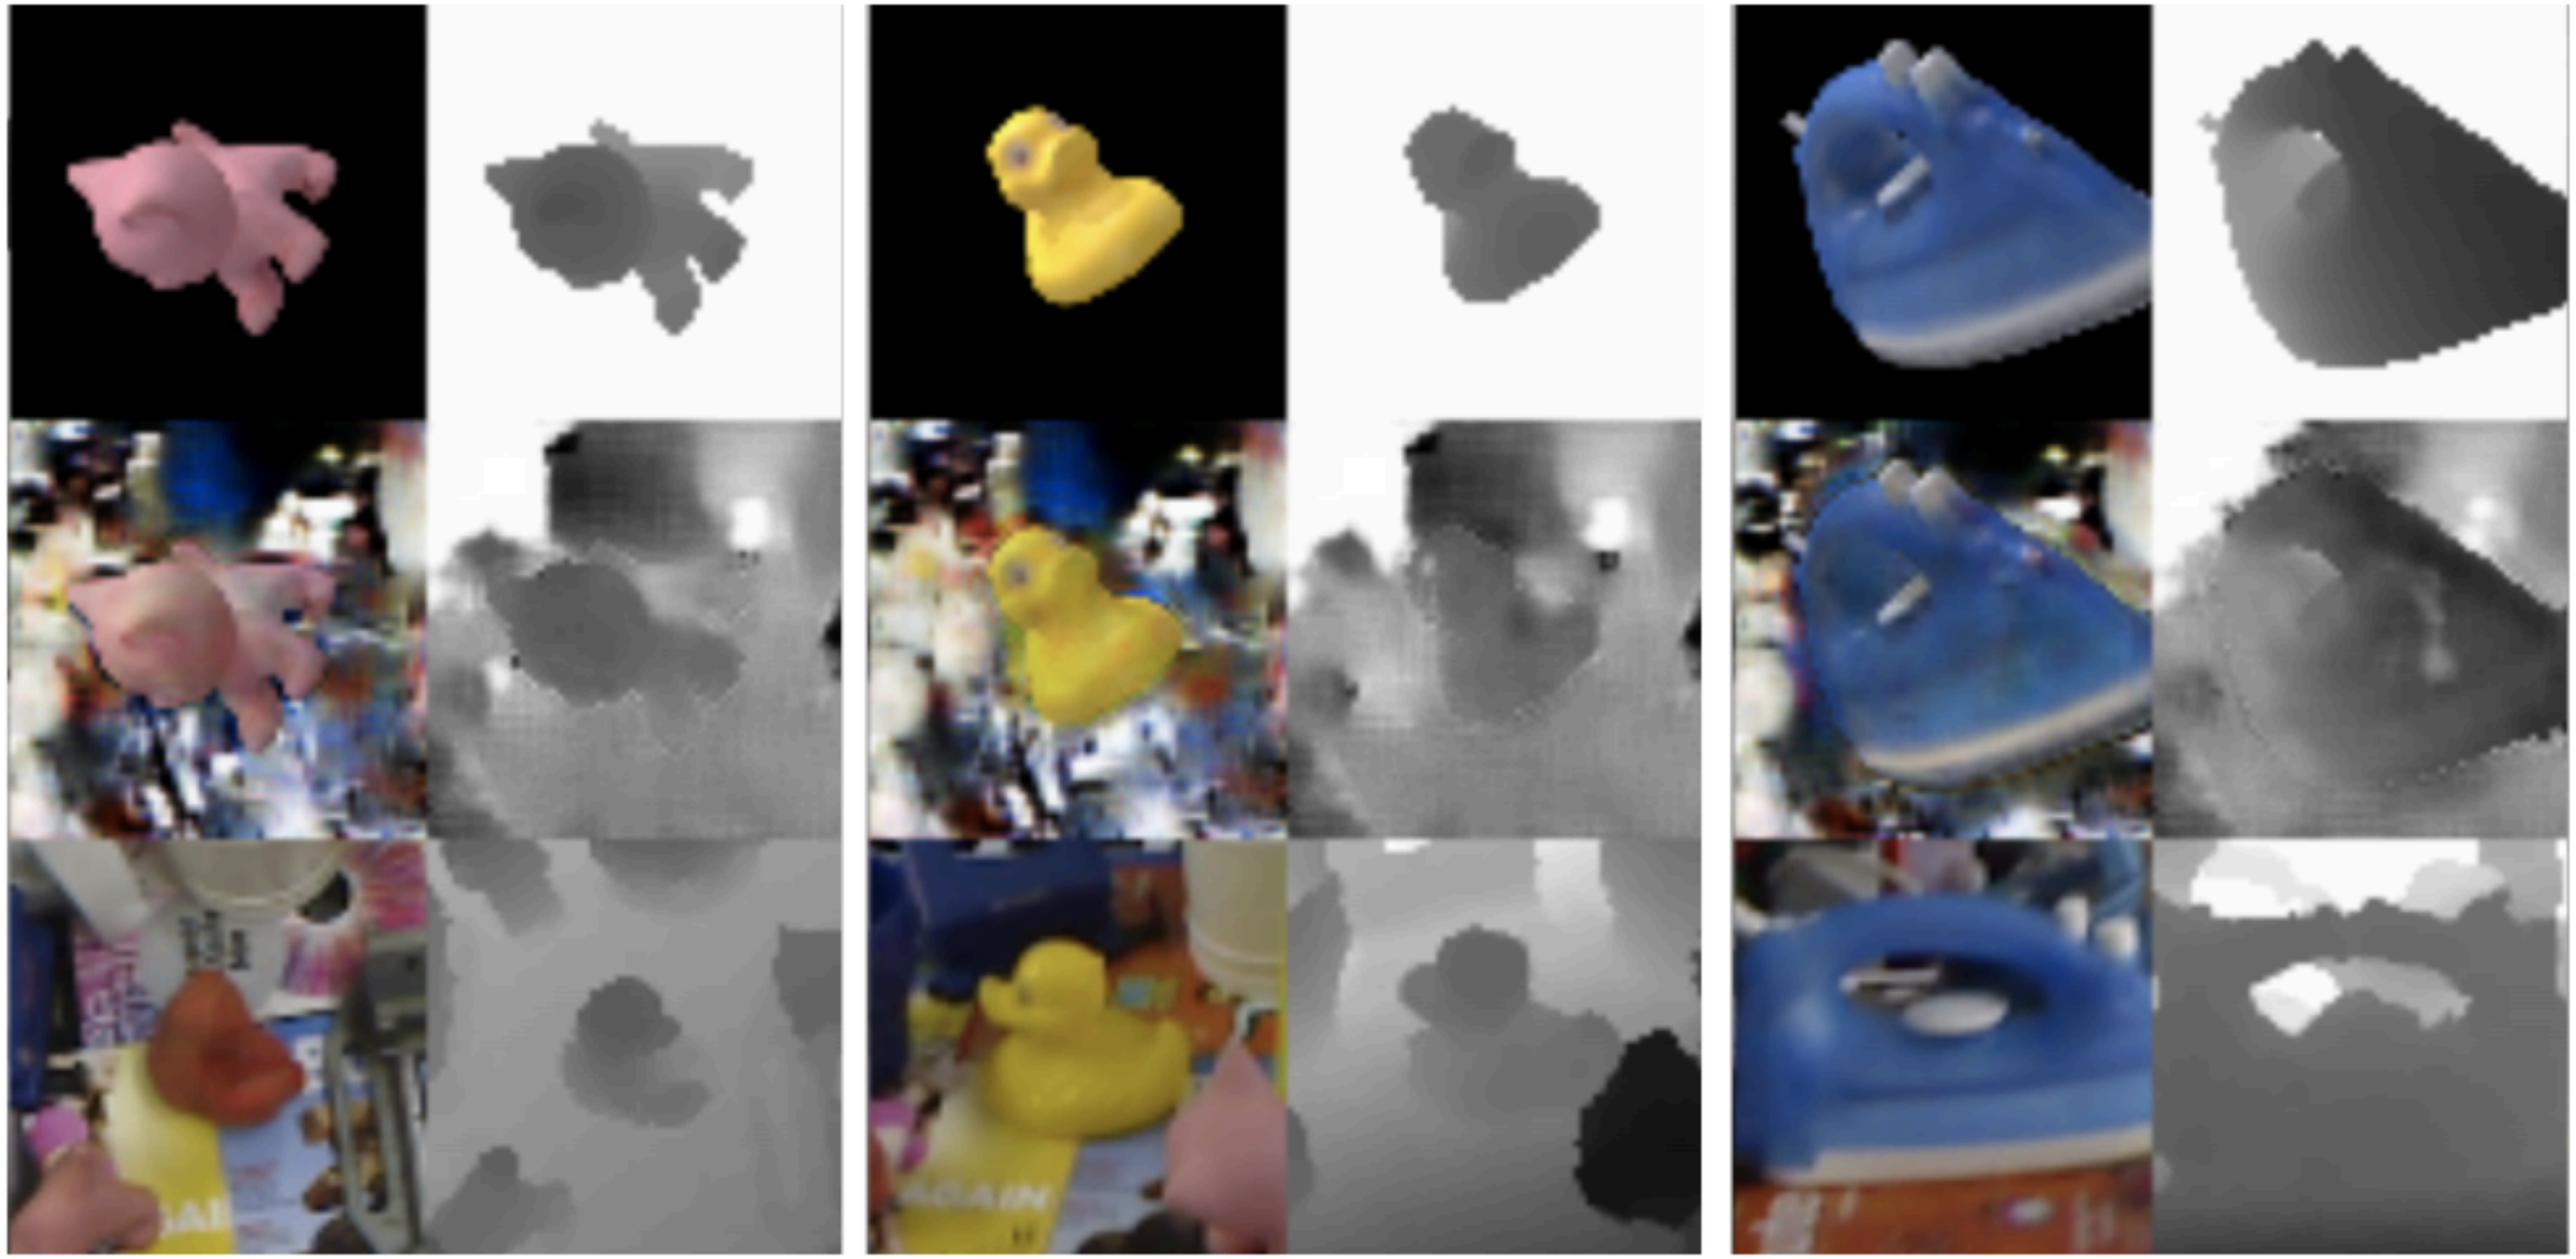

(Bousmalis et al., 2016)

(Raffel, 2017)

Autonomous Driving Data

Input labels

Synthesized image

(Wang et al., 2017)

Overcoming limited data with GANs

- Missing data
    - Semi-supervised learning
- Set-member supervision
- Unsupervised correspondence learning
- Replace data collection with simulation
- Simulated environments and training data
- Domain adaptation

(Goodfellow 2017)

Domain Adaptation

- Domain Adversarial Networks (Ganin et al, 2015)

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

Questions