Generative Adversarial Networks

Philipp Krähenbühl, UT Austin

Generative models

- Two tasks of a generative model $P(X)$
  - Sampling: $x \sim P(X)$
  - Density estimation: $P(X = x)$

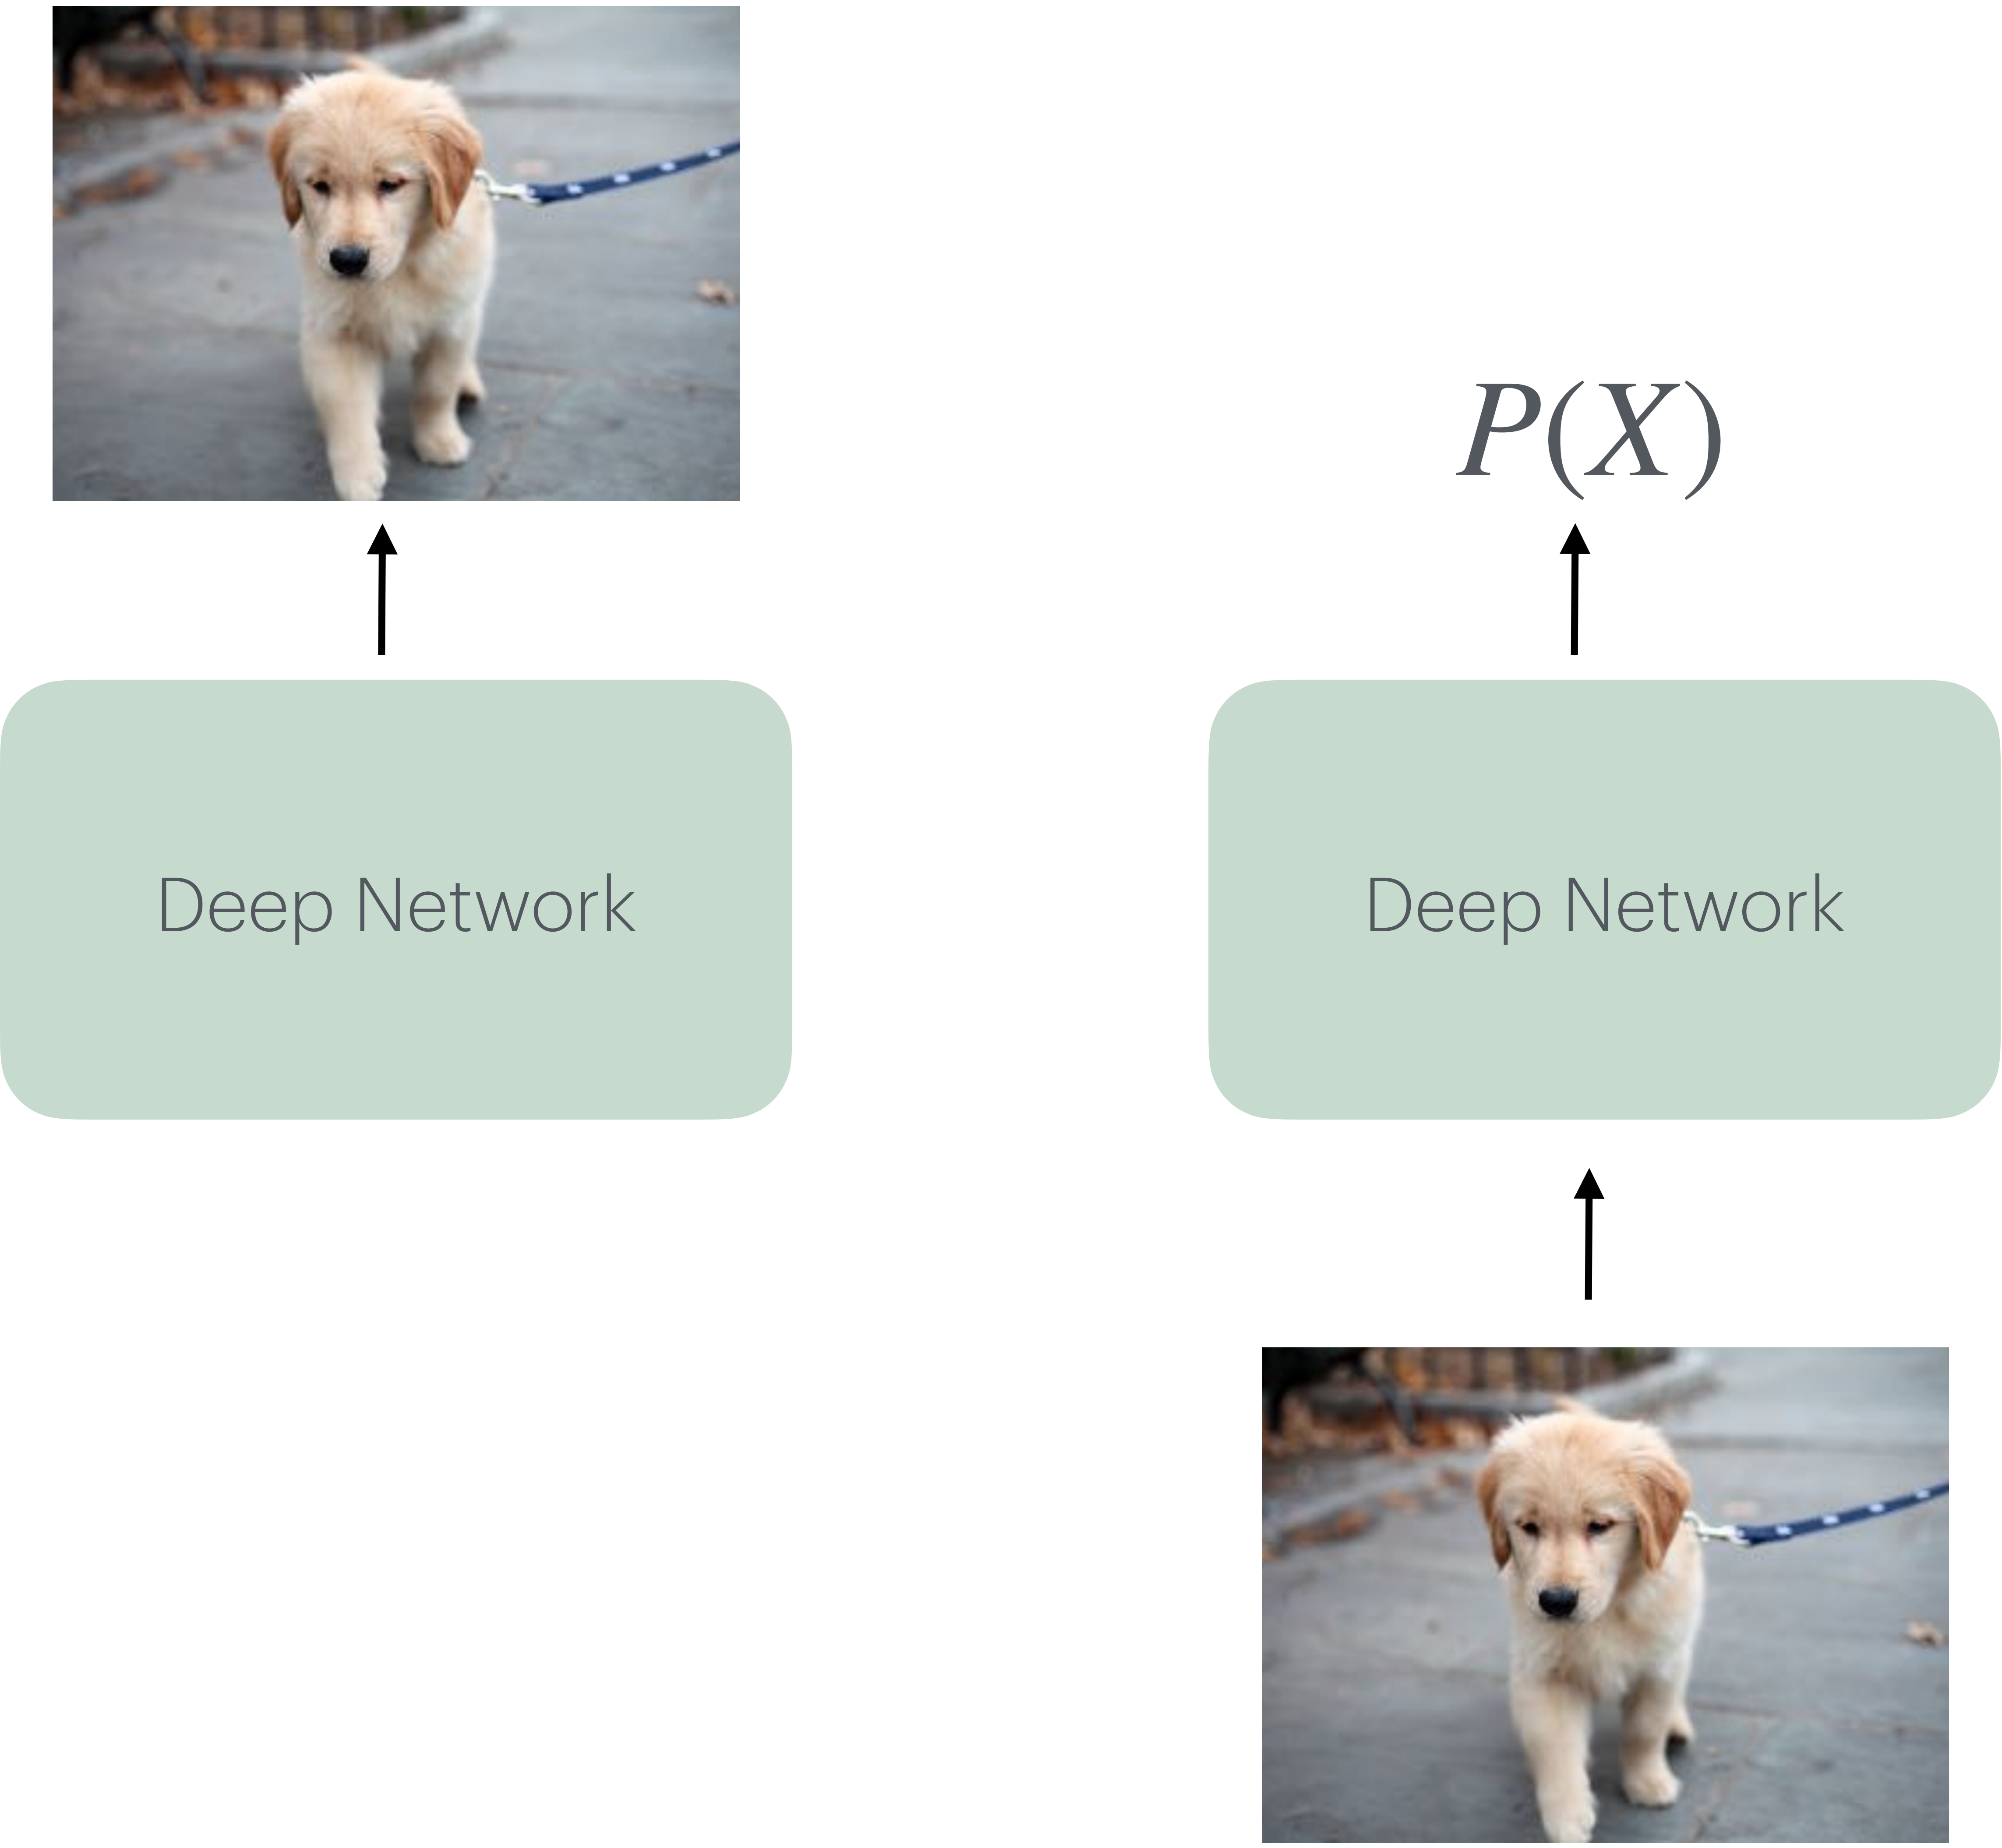

Generative modeling is hard

- Density estimation $P(X = x)$
  - How to ensure $\sum_x P(x) = 1$ for all x
  - Impossible to compute (in general)
- Sampling $x \sim P(X)$
  - What is the input to the network?

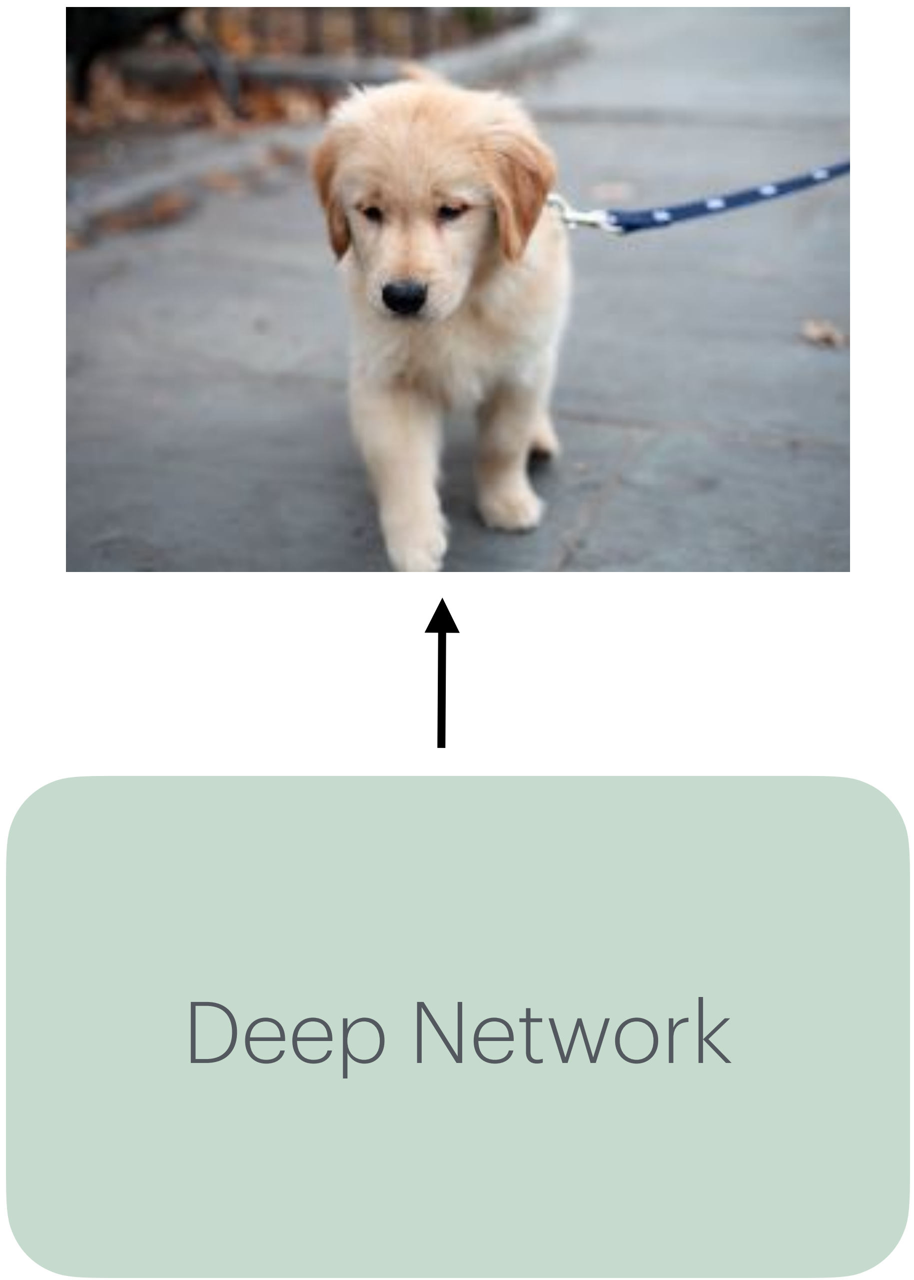

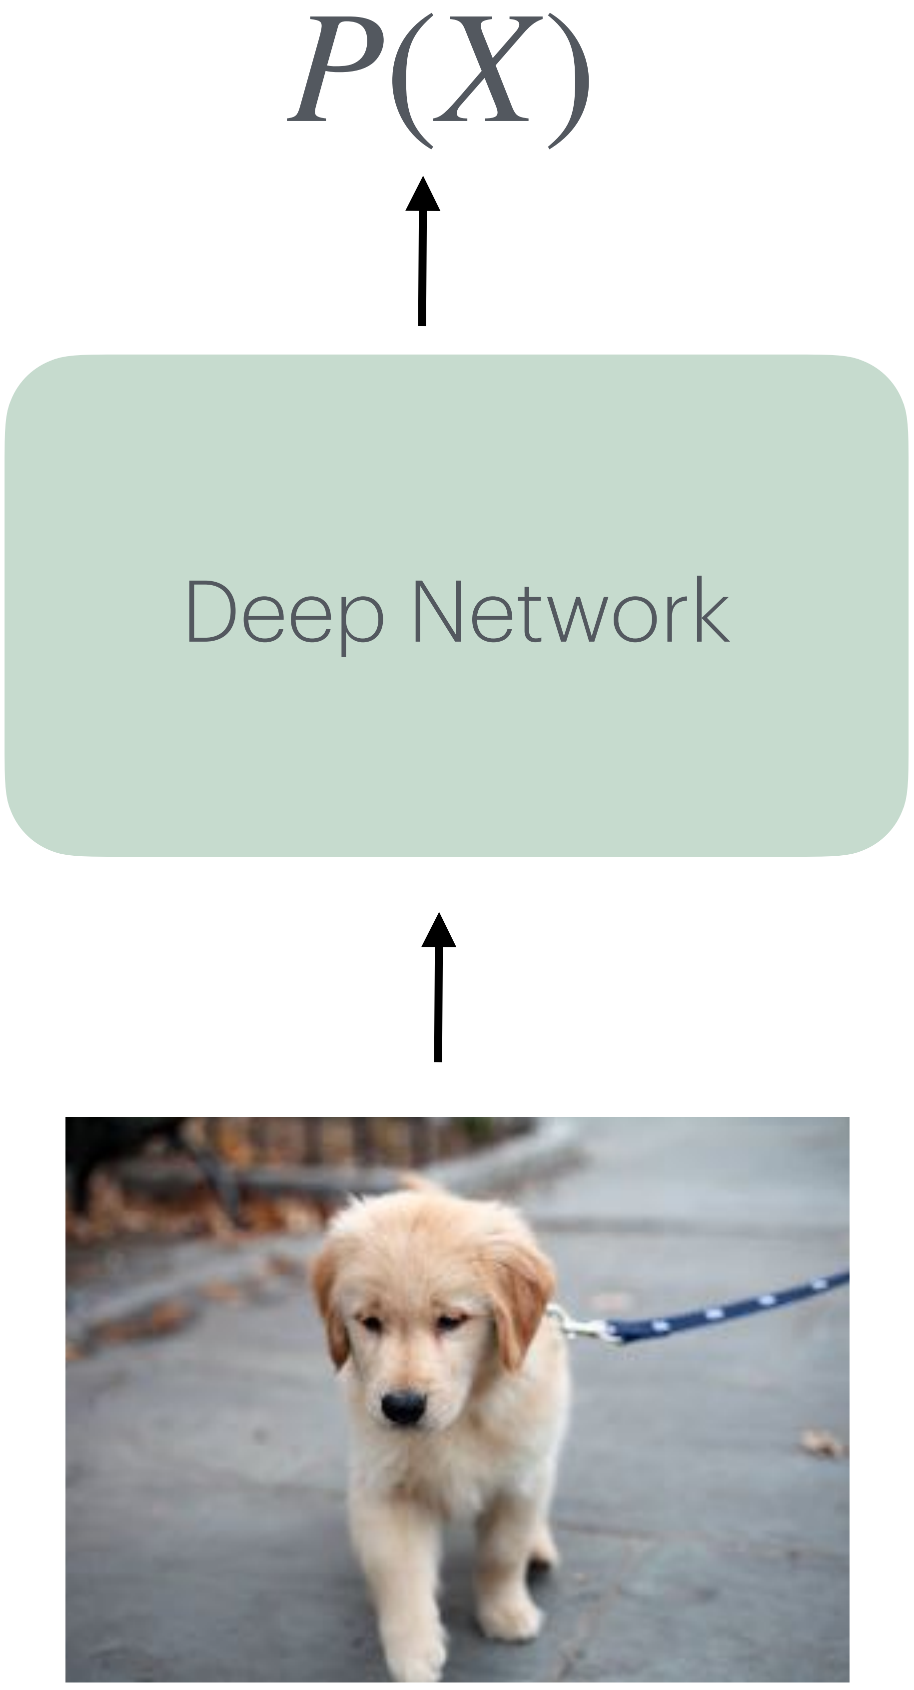

Generative models

Two kinds of models

Sampling based $x \sim P(X)$

- Sample $z \sim P(Z)$
- Learn transformation
  - $P(x \mid z)$ or $f : z \rightarrow x$

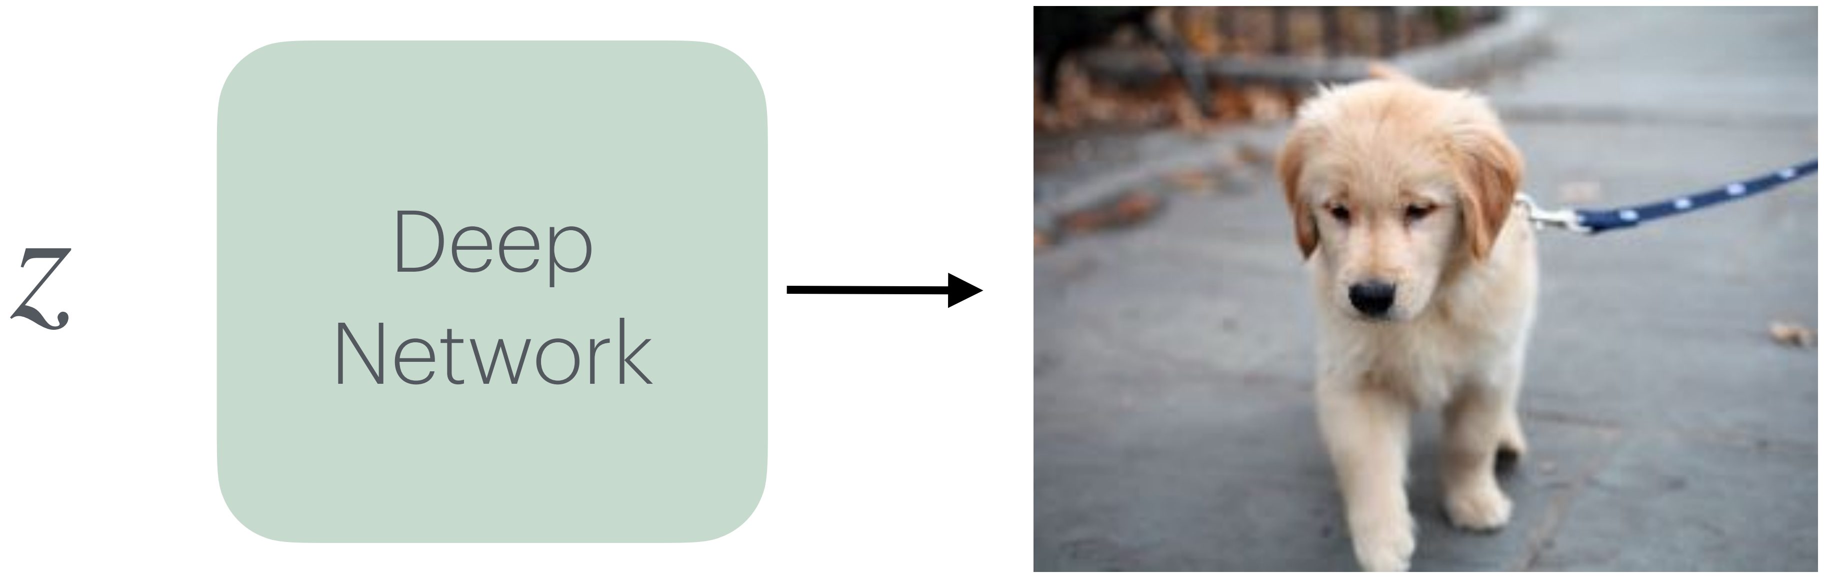

Density estimation based $P(X)$

- Learn special form of $P(X)$
- Model specific sampling / generation

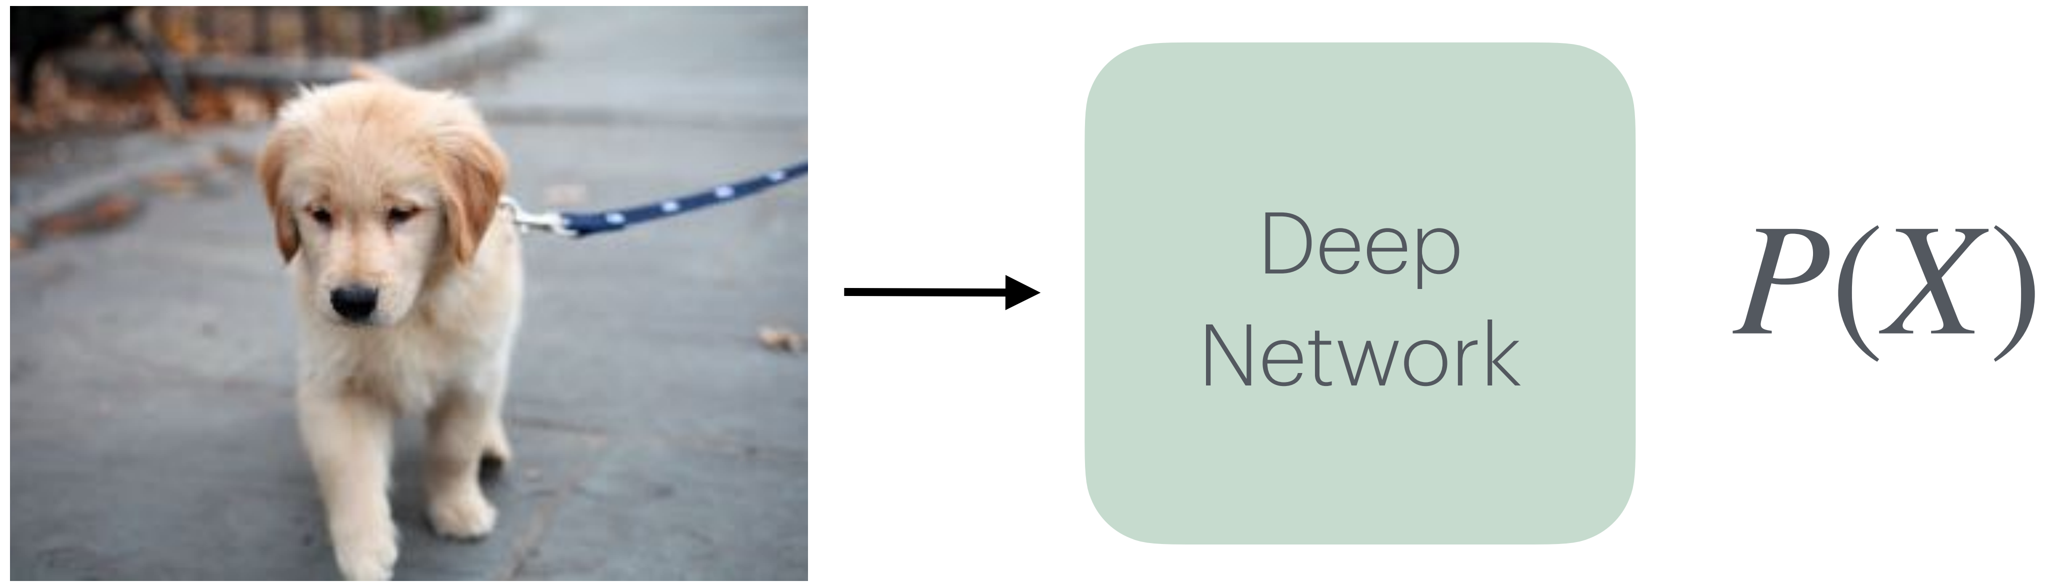

Variational auto-encoder

A "probabilistic" auto-encoder

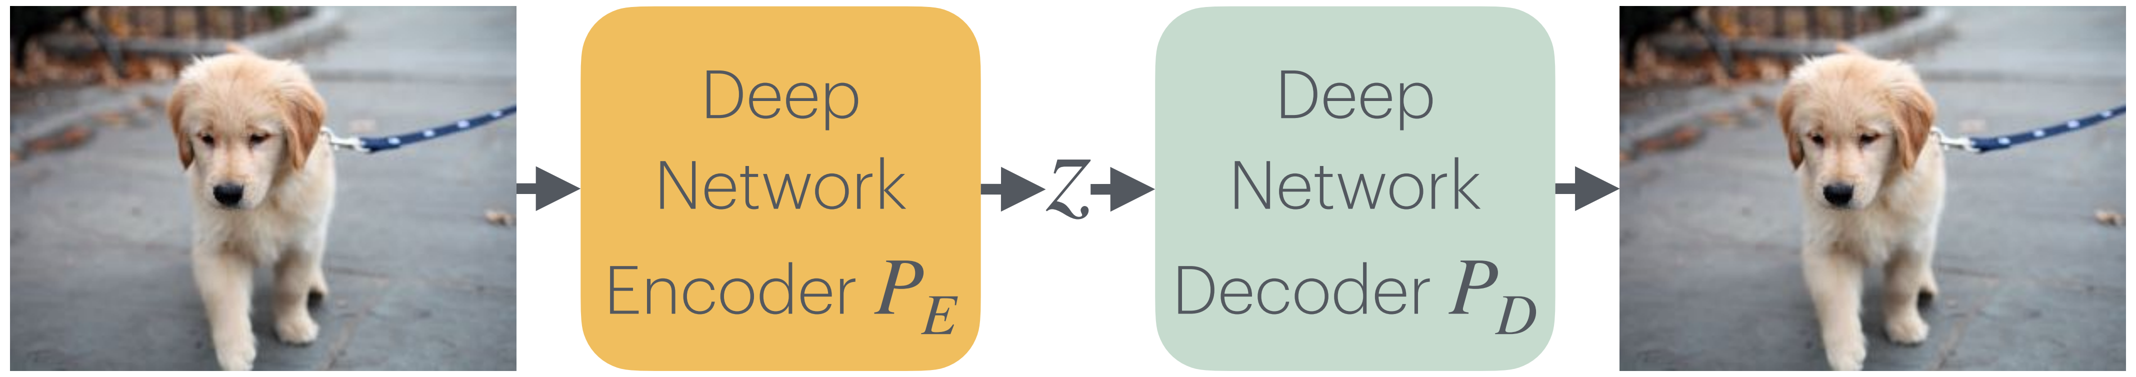

- Learn a model of $P(x) = P_D(x \mid z)P(z)$ with $P(z) = \mathcal{N}(z; 0,1)$
  - Approximate $Q \approx P_E$
  - Maximize $P(x)$ of data

[1] Auto-Encoding Variational Bayes, Kingma et al., ICLR 2014

Generative Adversarial Networks

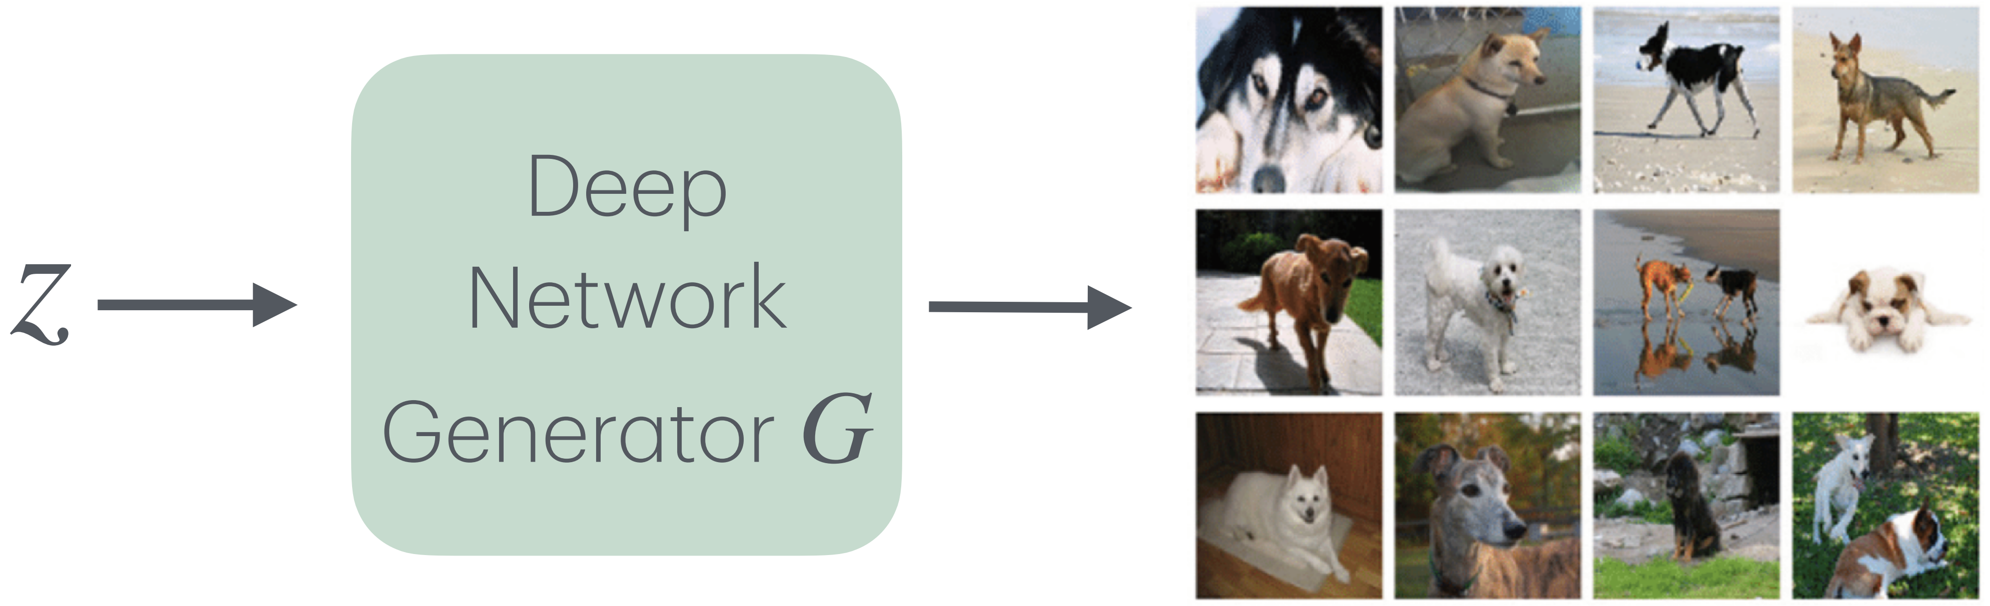

- Learn a model of $P(x)$
  - Sampling distribution (model)
    - $x = G(z); z \sim \mathcal{N}(0,1)$
  - Match data distribution
    - $x \sim \mathcal{D}$

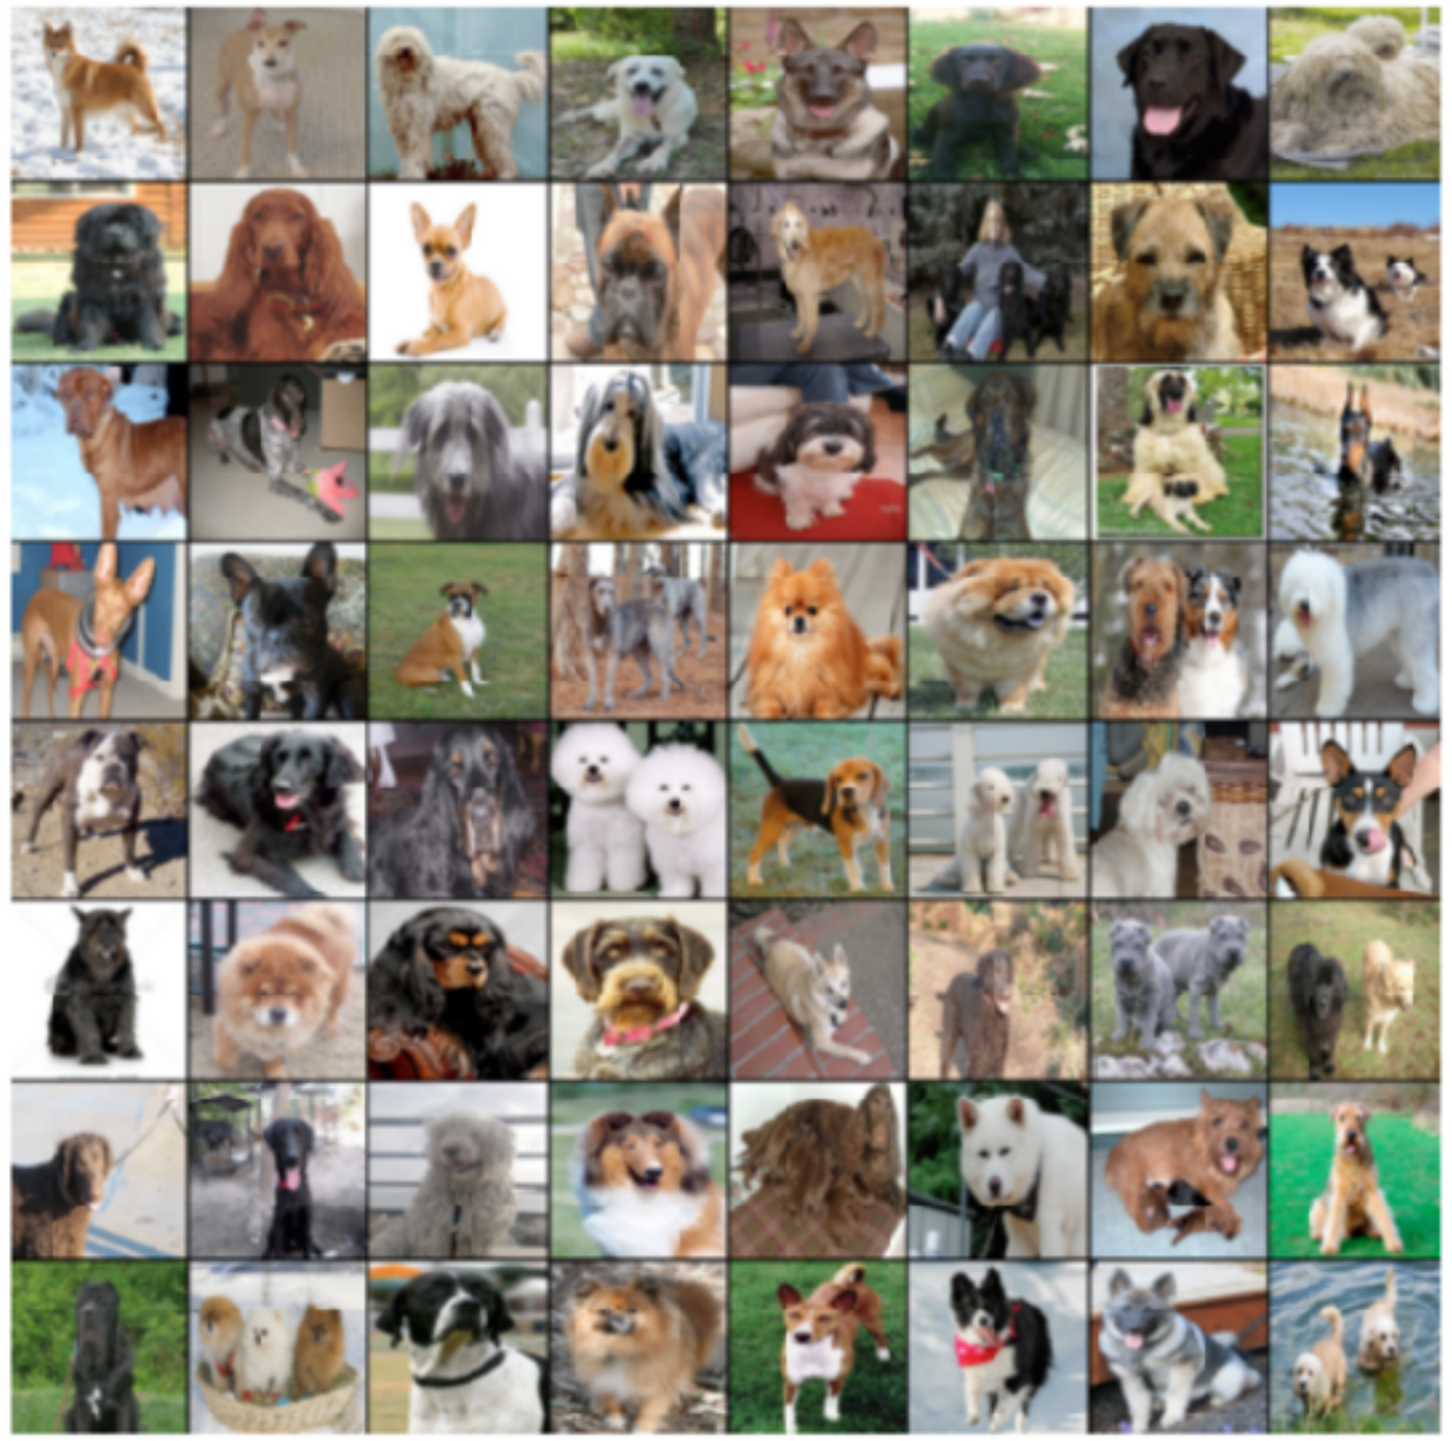

Columbia University Dogs dataset [3]

[2] Generative Adversarial Networks, Goodfellow et al., NeurIPS 2014
[3] Dog breed classification using part localization, Liu et al., ECCV 2012

Generative Adversarial Networks

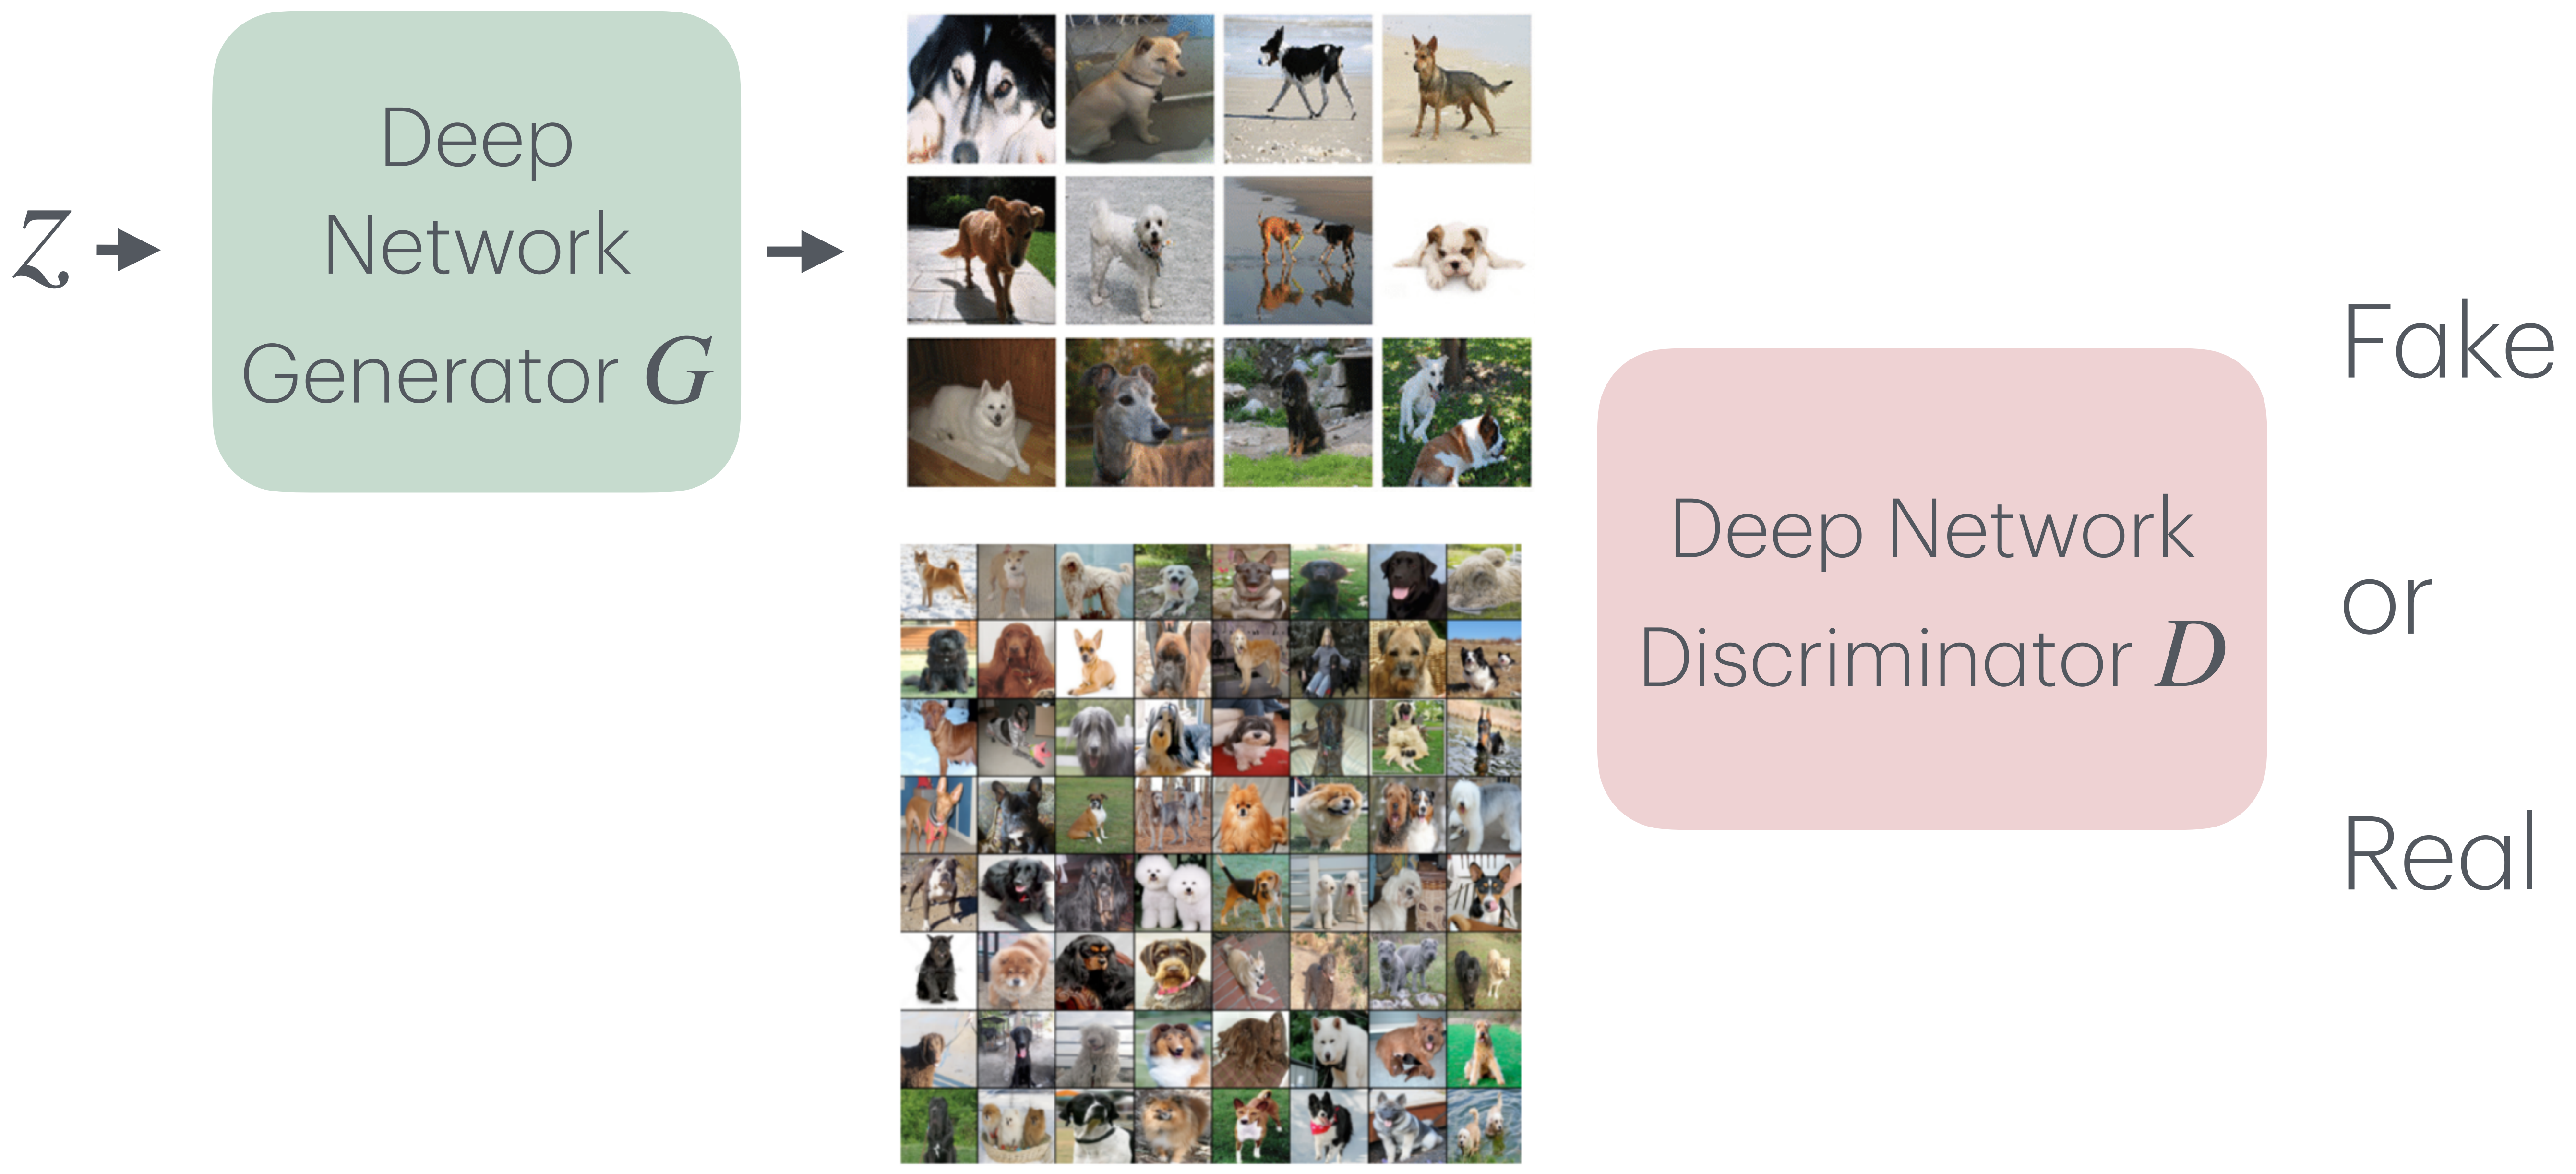

- Objective: Match sampling distributions as two player game
  - Payer 1: Generator G
    - Generate images from noise
  - Payer 2: Discriminator D
    - Tell real from fake images

$$\min_G \max_D E_{x\sim G}\left[\log D(x)\right] + E_{x\sim\mathcal{D}}\left[\log(1-D(x))\right] =$$

$$\min_G \max_D E_{z\sim\mathcal{N}}\left[\log D(G(x))\right] + E_{x\sim\mathcal{D}}\left[\log(1-D(x))\right]$$

[2] Generative Adversarial Networks, Goodfellow et al., NeurIPS 2014

Generative Adversarial Networks

$z \rightarrow$ Deep Network Generator G $\rightarrow$

Deep Network Discriminator D

Fake

or

Real

- Minimizes Jensen-Shannon-Divergence
  $JSD(P, Q) = \max_M D_{KL}(P \,|\, M) + D_{KL}(\mathscr{D} \,|\, M)$

$$\min_G \max_D E_{x \sim G}\left[\log D(x)\right] + E_{x \sim \mathscr{D}}\left[\log(1 - D(x))\right] =$$

$$\min_G \max_D E_{z \sim \mathscr{N}}\left[\log D(G(x))\right] + E_{x \sim \mathscr{D}}\left[\log(1 - D(x))\right]$$

[2] Generative Adversarial Networks, Goodfellow et al., NeurIPS 2014

Generative Adversarial Networks

Optimization

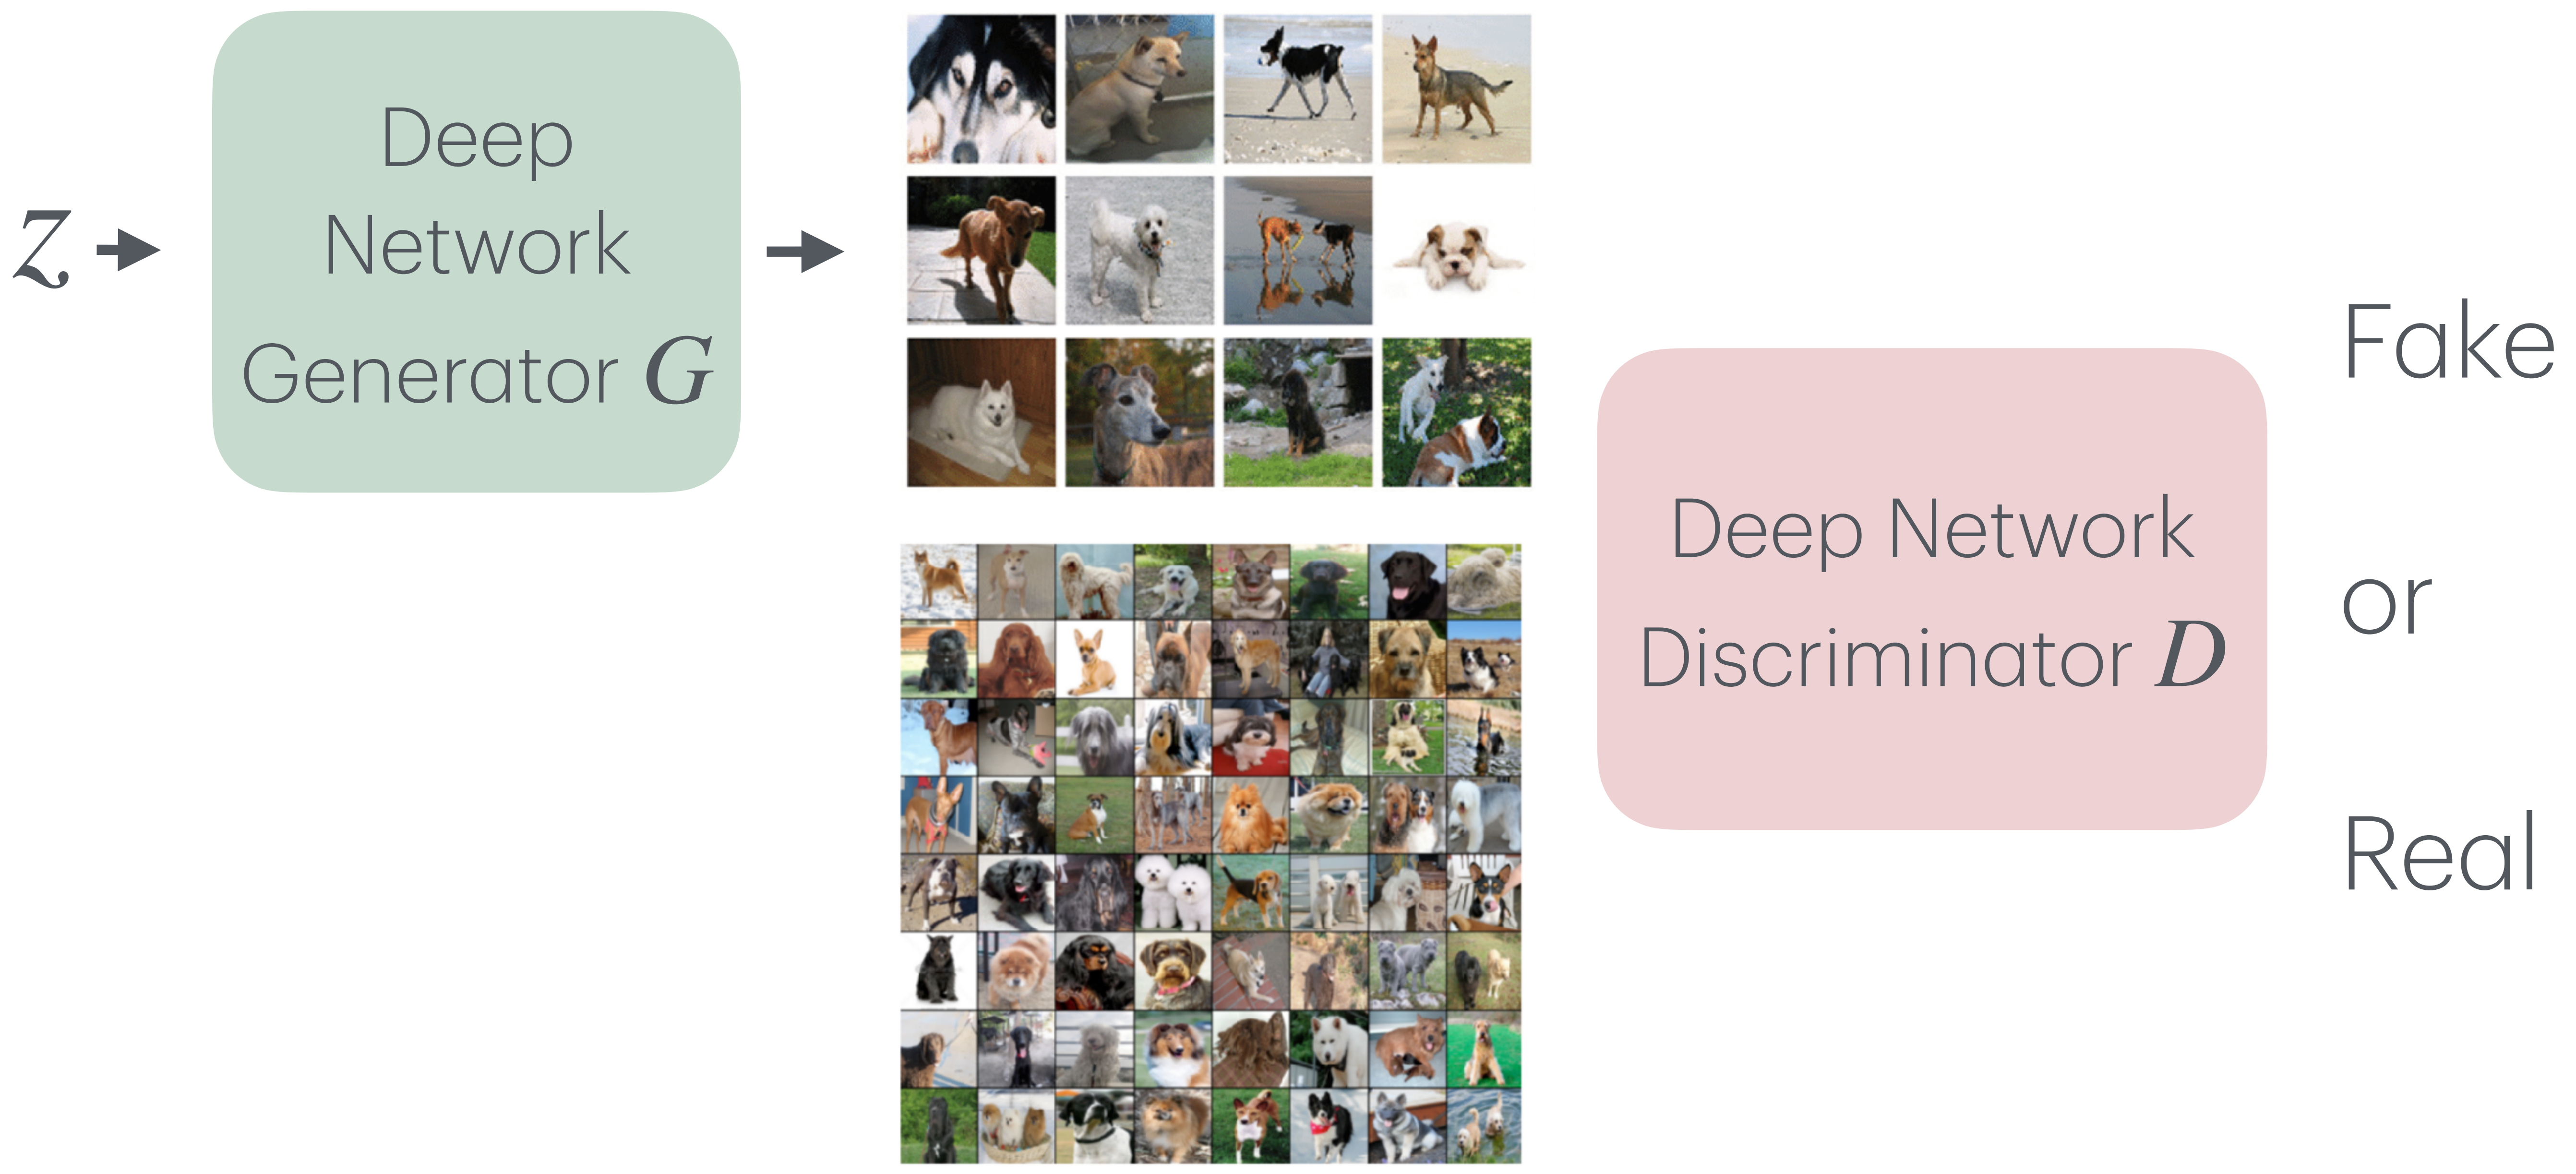

$$\min_{G}\max_{D} E_{z\sim\mathcal{N}}\left[\log D(G(x))\right] + E_{x\sim\mathscr{D}}\left[\log(1 - D(x))\right]$$

- Hard to optimize
  - Mathematically correct
    - For every step in G, run a full optimizer pass for D
  - In practice
    - One step G, one step D
    - Pray for convergence

[2] Generative Adversarial Networks, Goodfellow et al., NeurIPS 2014

GANs work!

- Sampling is easy
- Learned pixel-distance
  - No blurriness
- Loss on distributions
  - Requires a few tricks

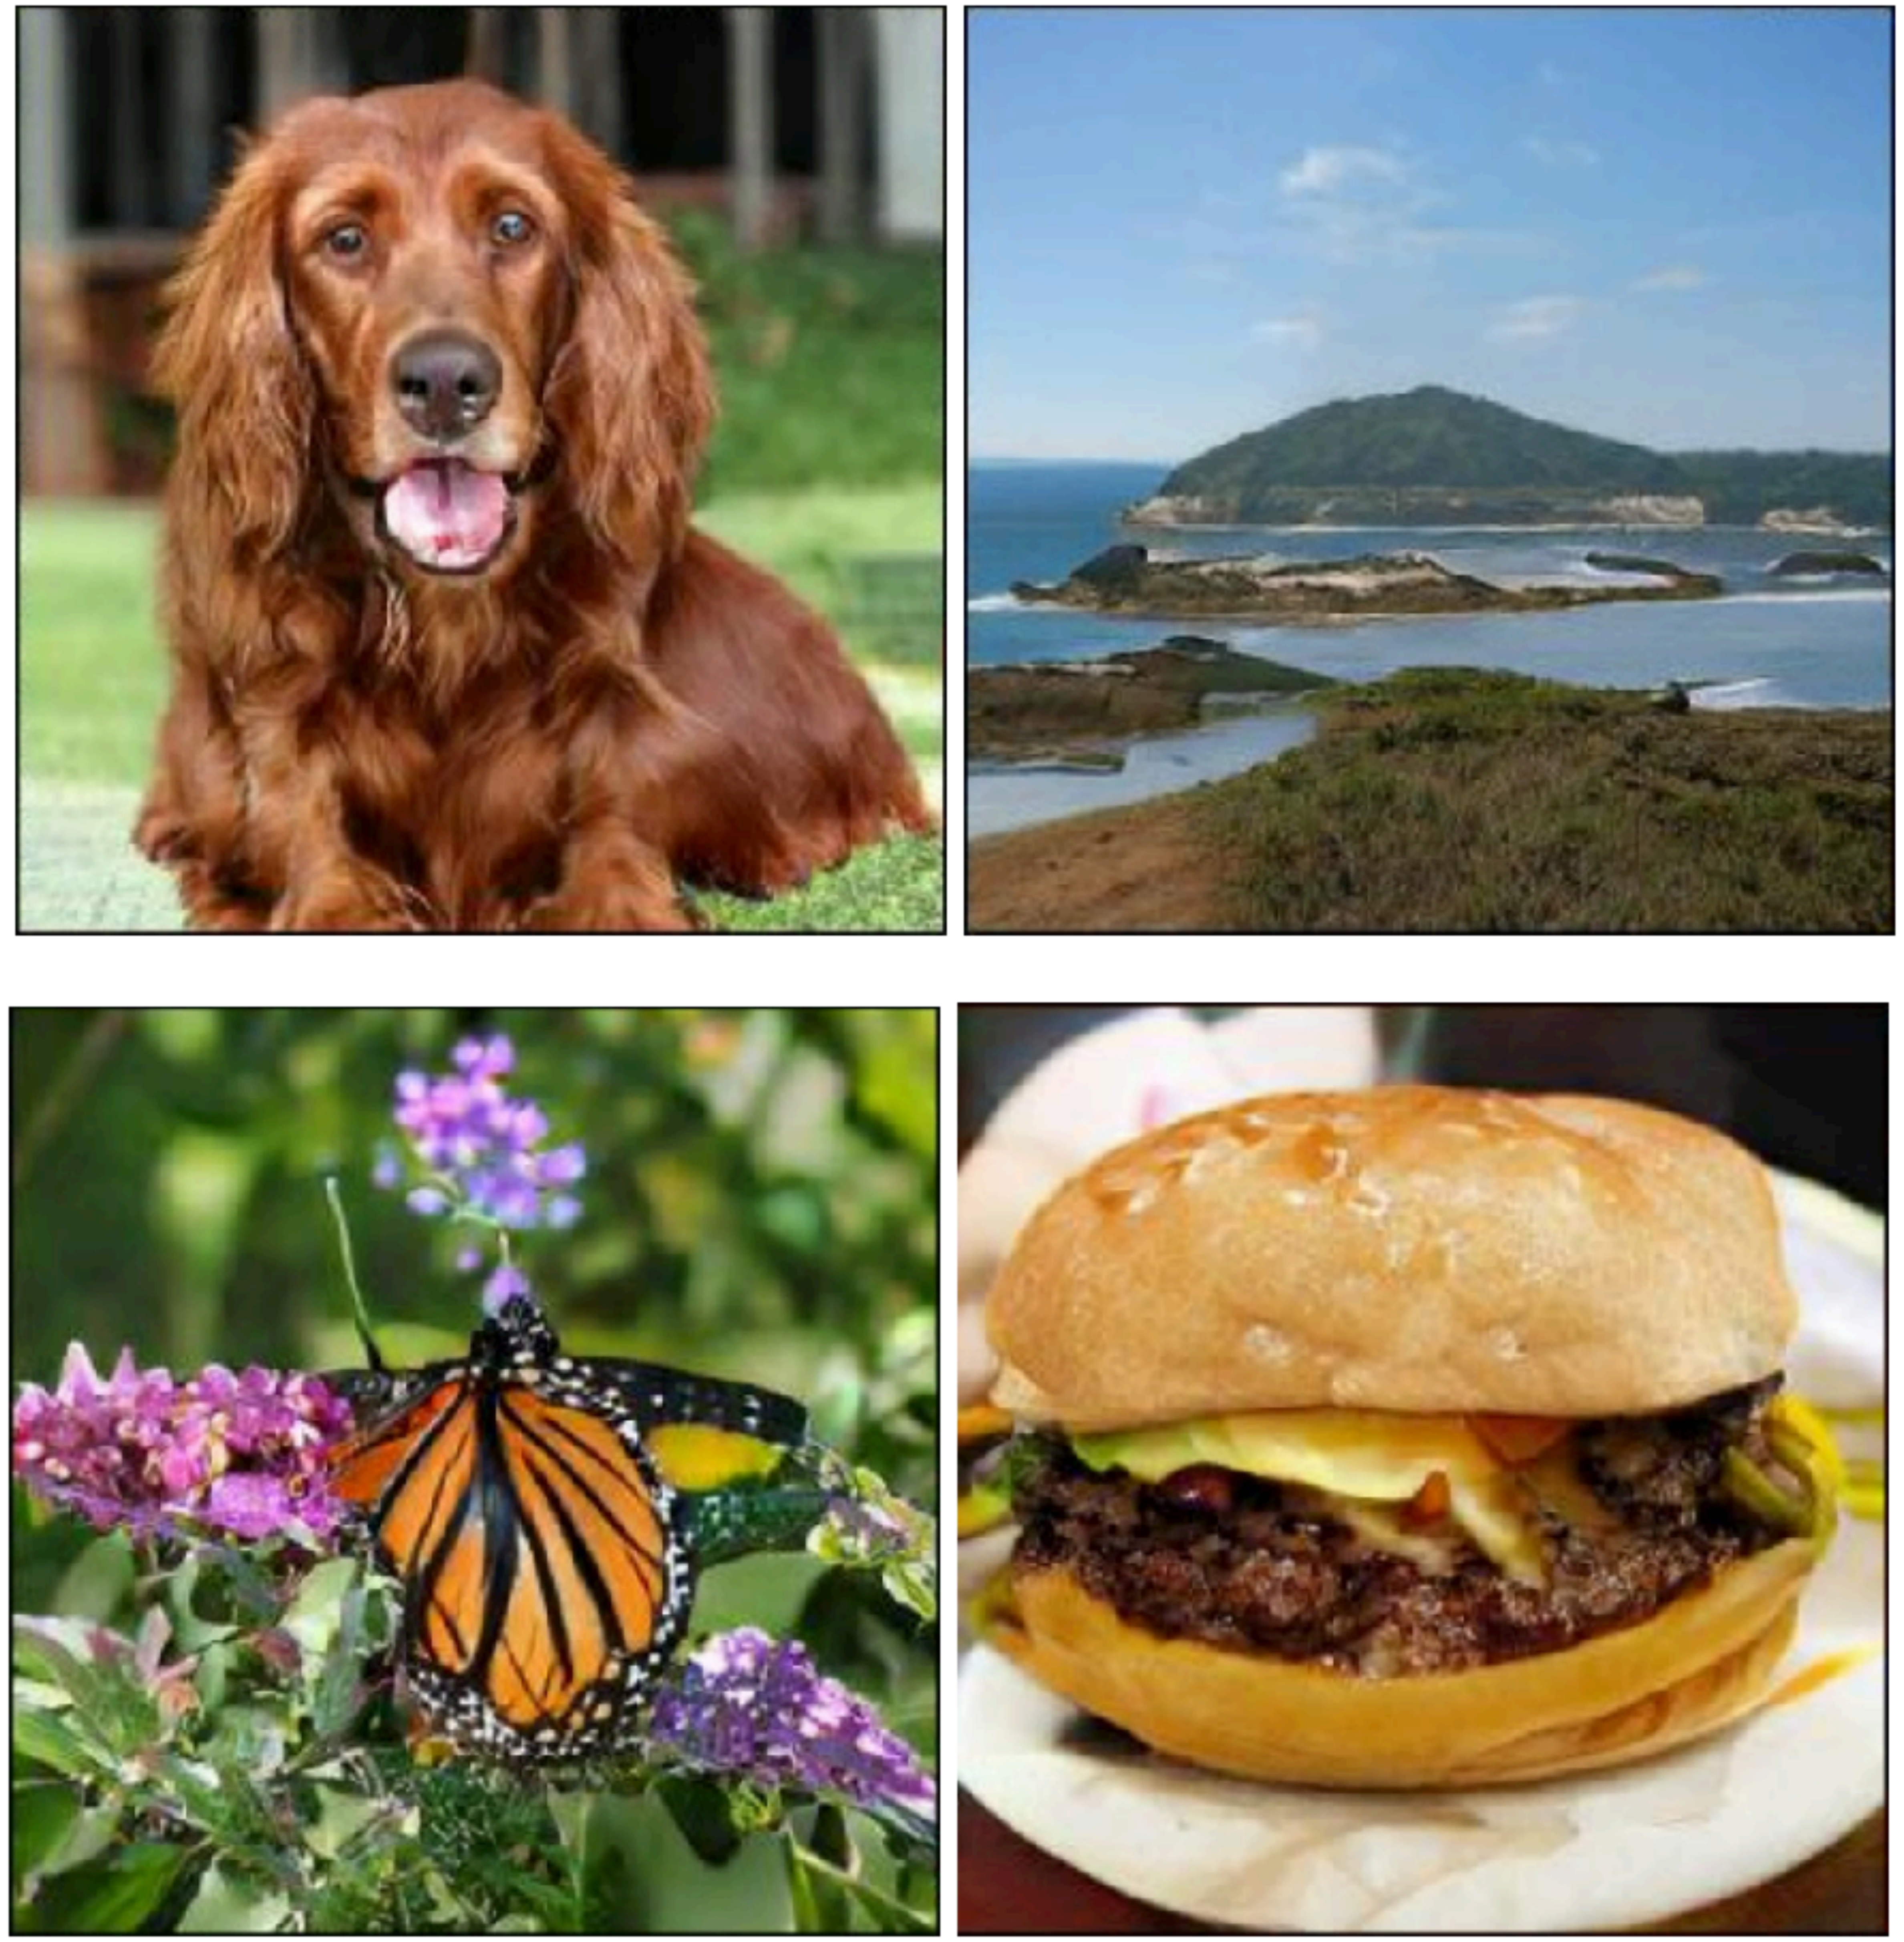

[4] Large Scale GAN Training for High Fidelity Natural Image Synthesis, Brock et al., ICLR 2019

Application

Super resolution

- Supervised training task
  - Take high-res image, downsample
- GANs as loss
  - Leads to sharper reconstruction than just reconstruction

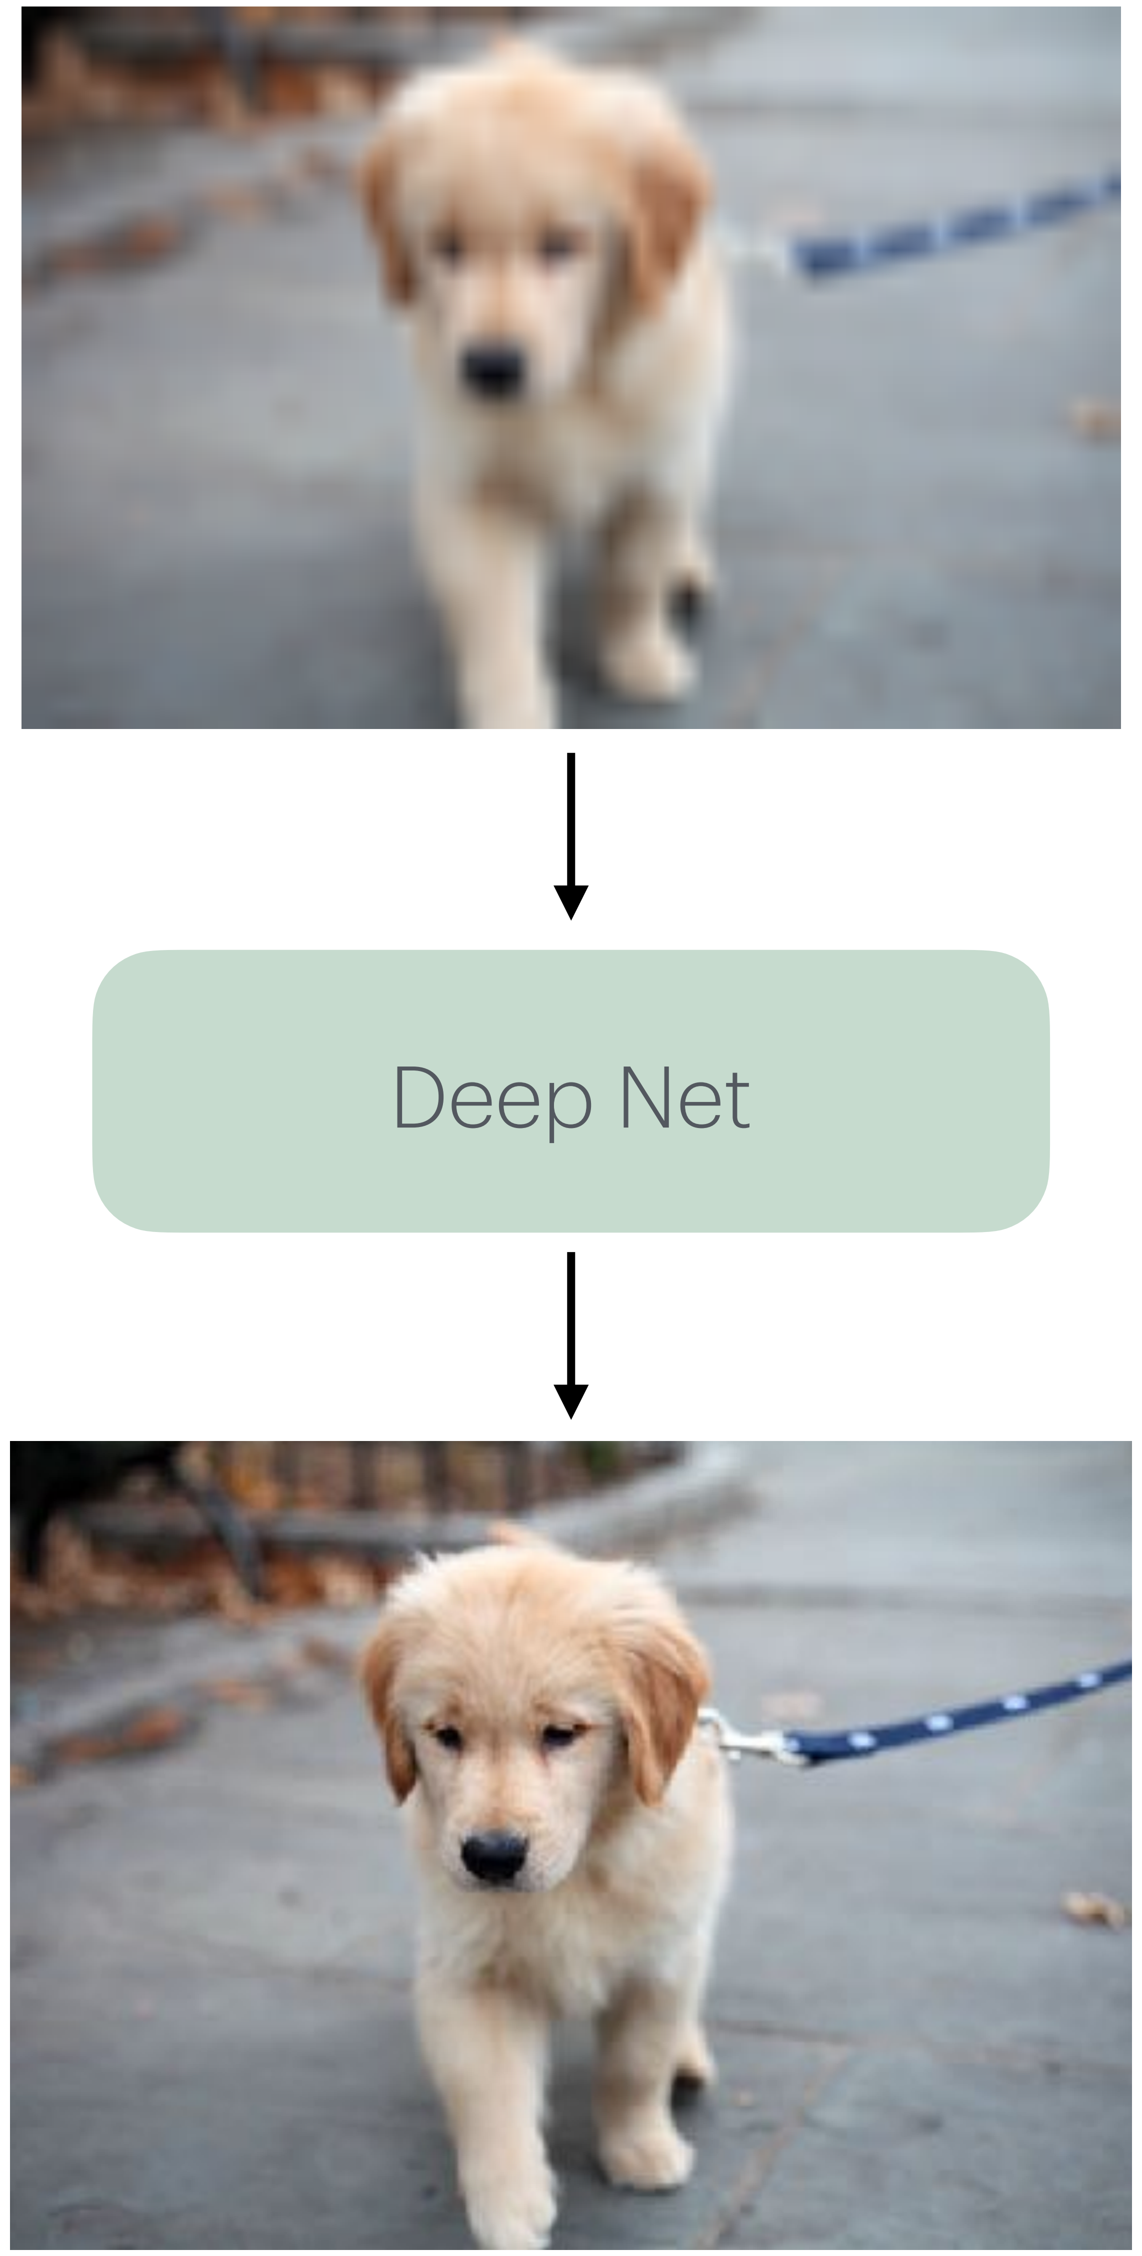

[5] Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, Ledig et al., CVPR 2017

Generative Adversarial Networks

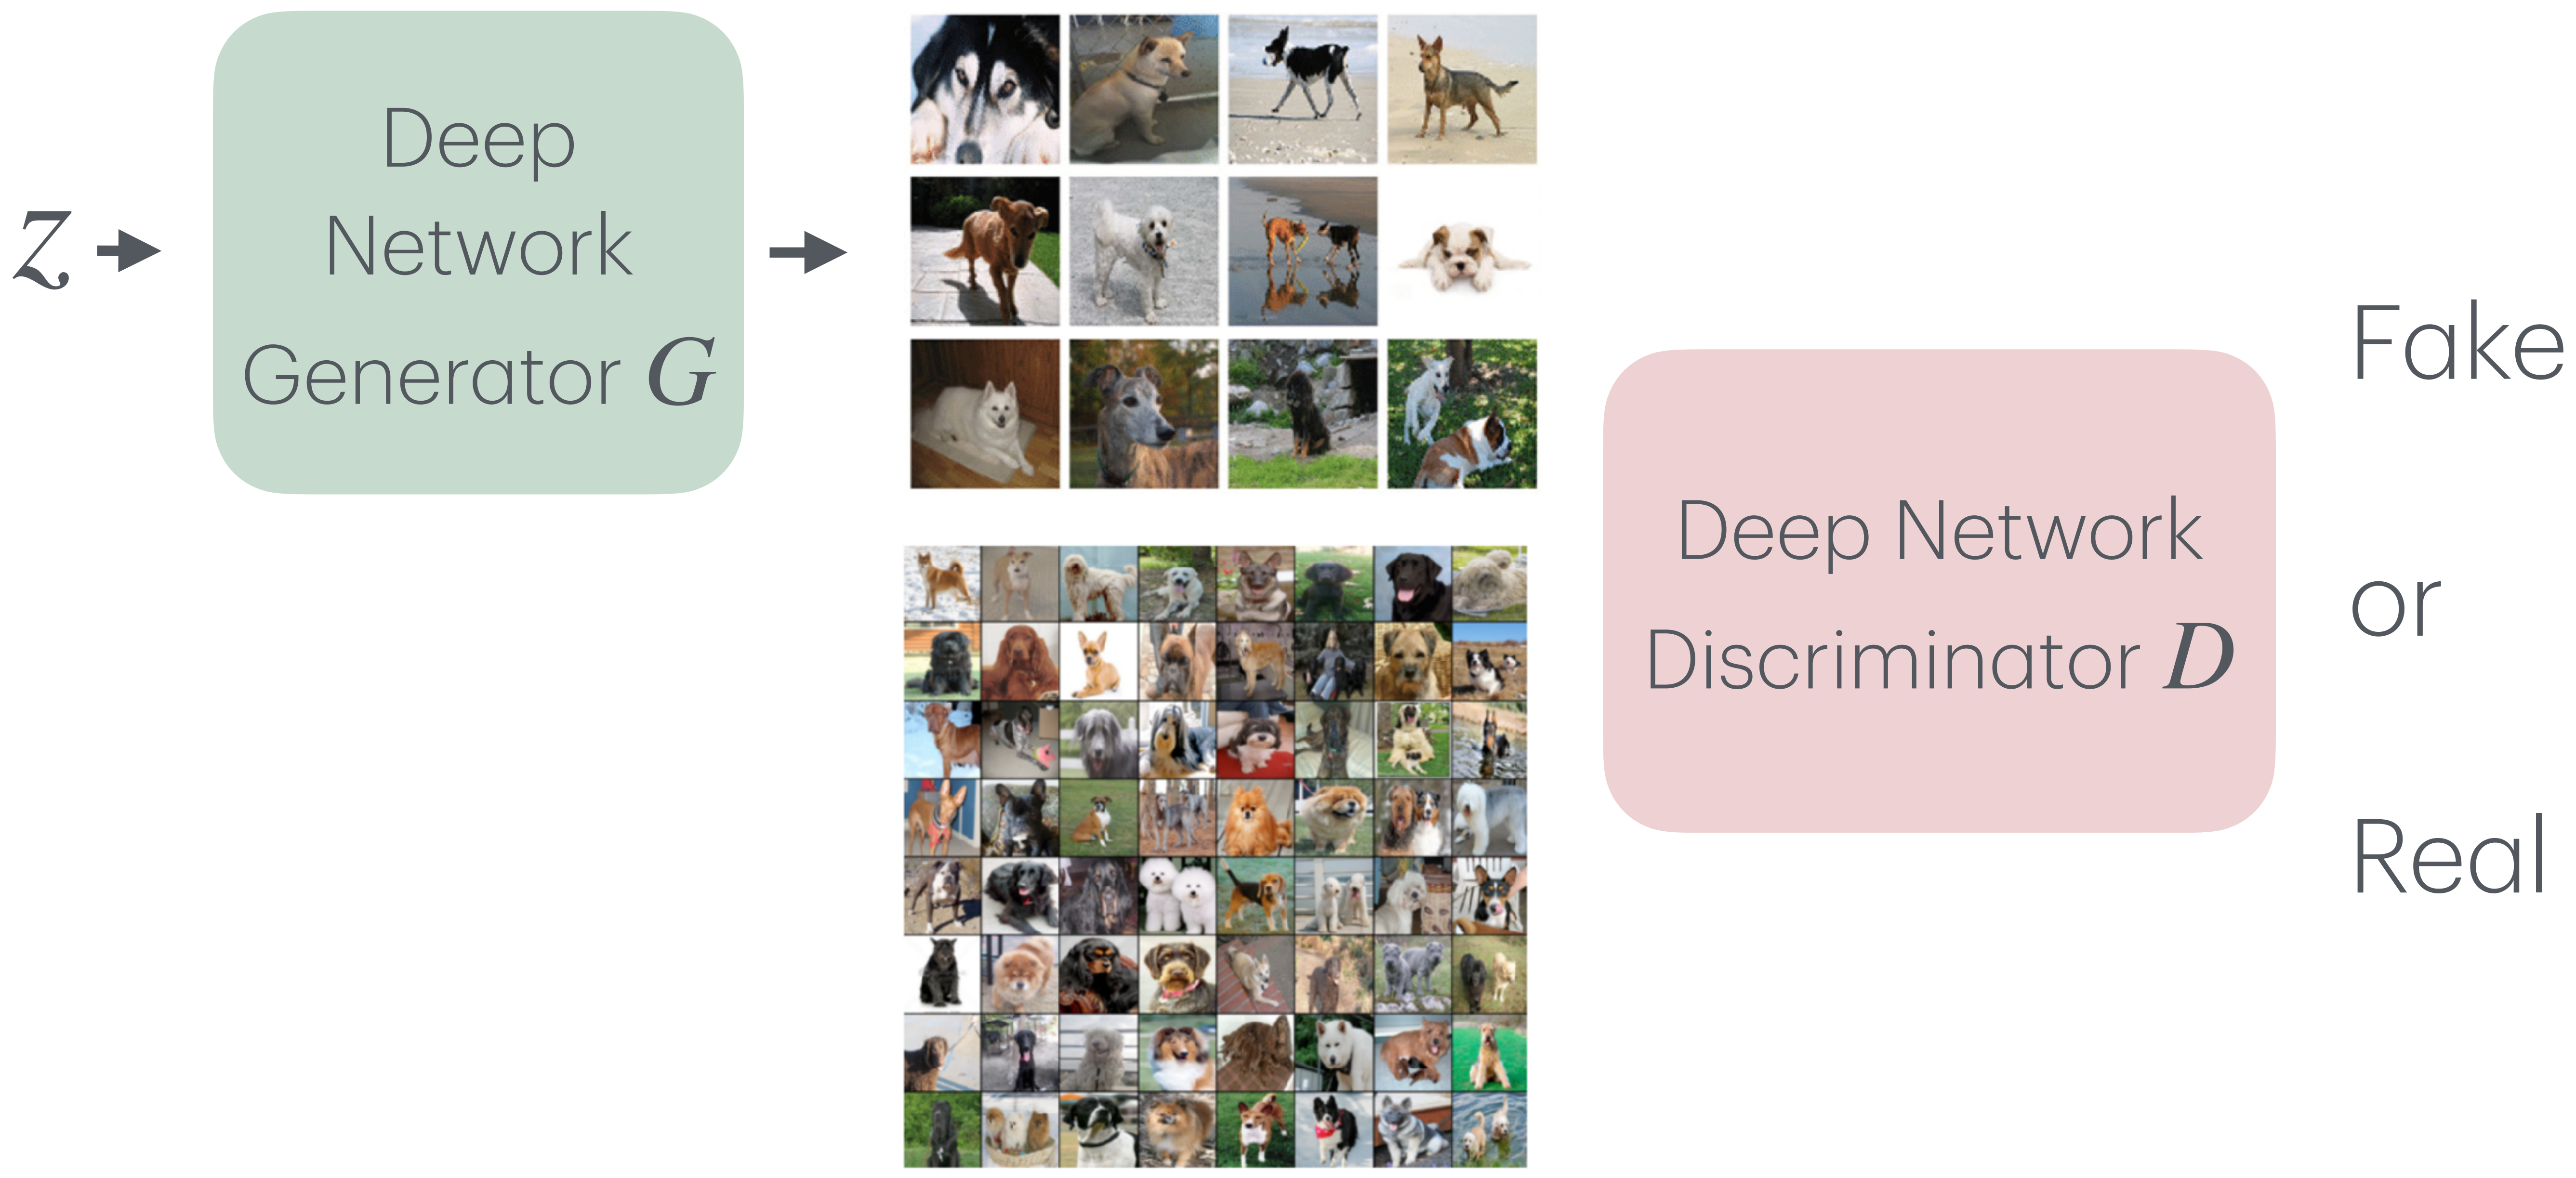

- Two player game
  - Generator: Transform noise into images
  - Discriminator: Tell difference between real and generated images
    - Leads to sharp images
    - Used in conditioned generation

[2] Generative Adversarial Networks, Goodfellow et al., NeurIPS 2014

Generative models

Two kinds of models

Sampling based $x \sim P(X)$

- Sample $z \sim P(Z)$
- Learn transformation
  - $P(x \mid z)$ or $f : z \to x$

z

Deep Network

Density estimation based $P(X)$

- Learn special form of $P(X)$
- Model specific sampling / generation

Deep Network

$P(X)$

References

- [1] Auto-Encoding Variational Bayes, Kingma et al., ICLR 2014 (link)
- [2] Generative Adversarial Networks, Goodfellow et al., NeurIPS 2014 (link)
- [3] Dog breed classification using part localization, Liu et al., ECCV 2012 (link)
- [4] Large Scale GAN Training for High Fidelity Natural Image Synthesis, Brock et al., ICLR 2019 (link)
- [5] Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, Ledig et al., CVPR 2017 (link)