Generative adversarial networks

Transforming noise

- Input: Random noise
  - e.g. $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- Output: Image
- Objective
  - Output should look good

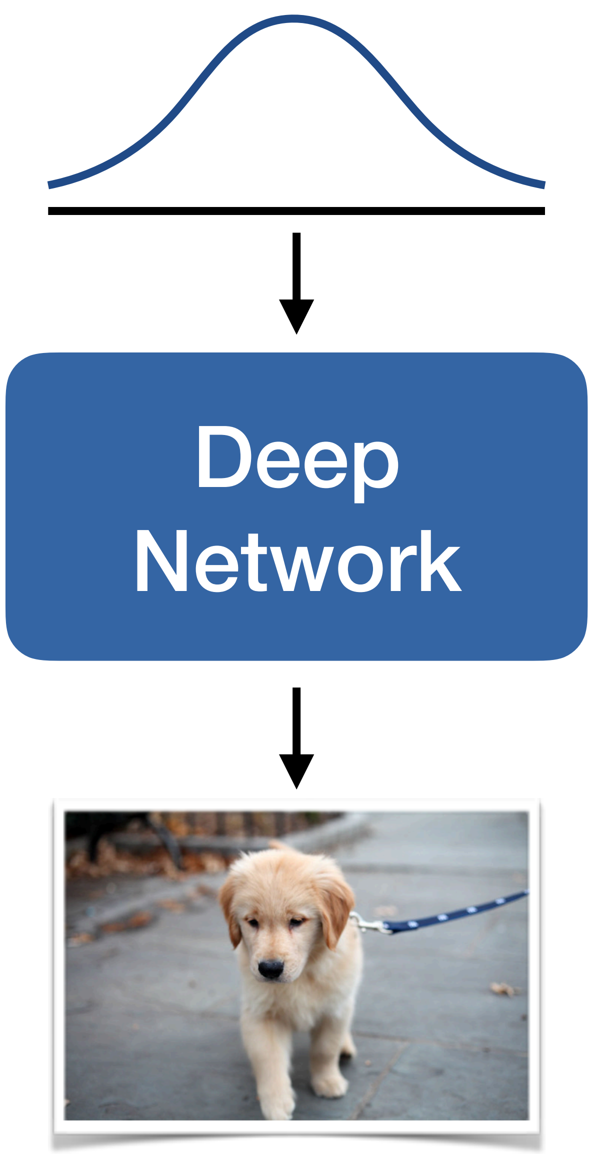

Judging output quality

- A second network
  - Output looks like training data or not?

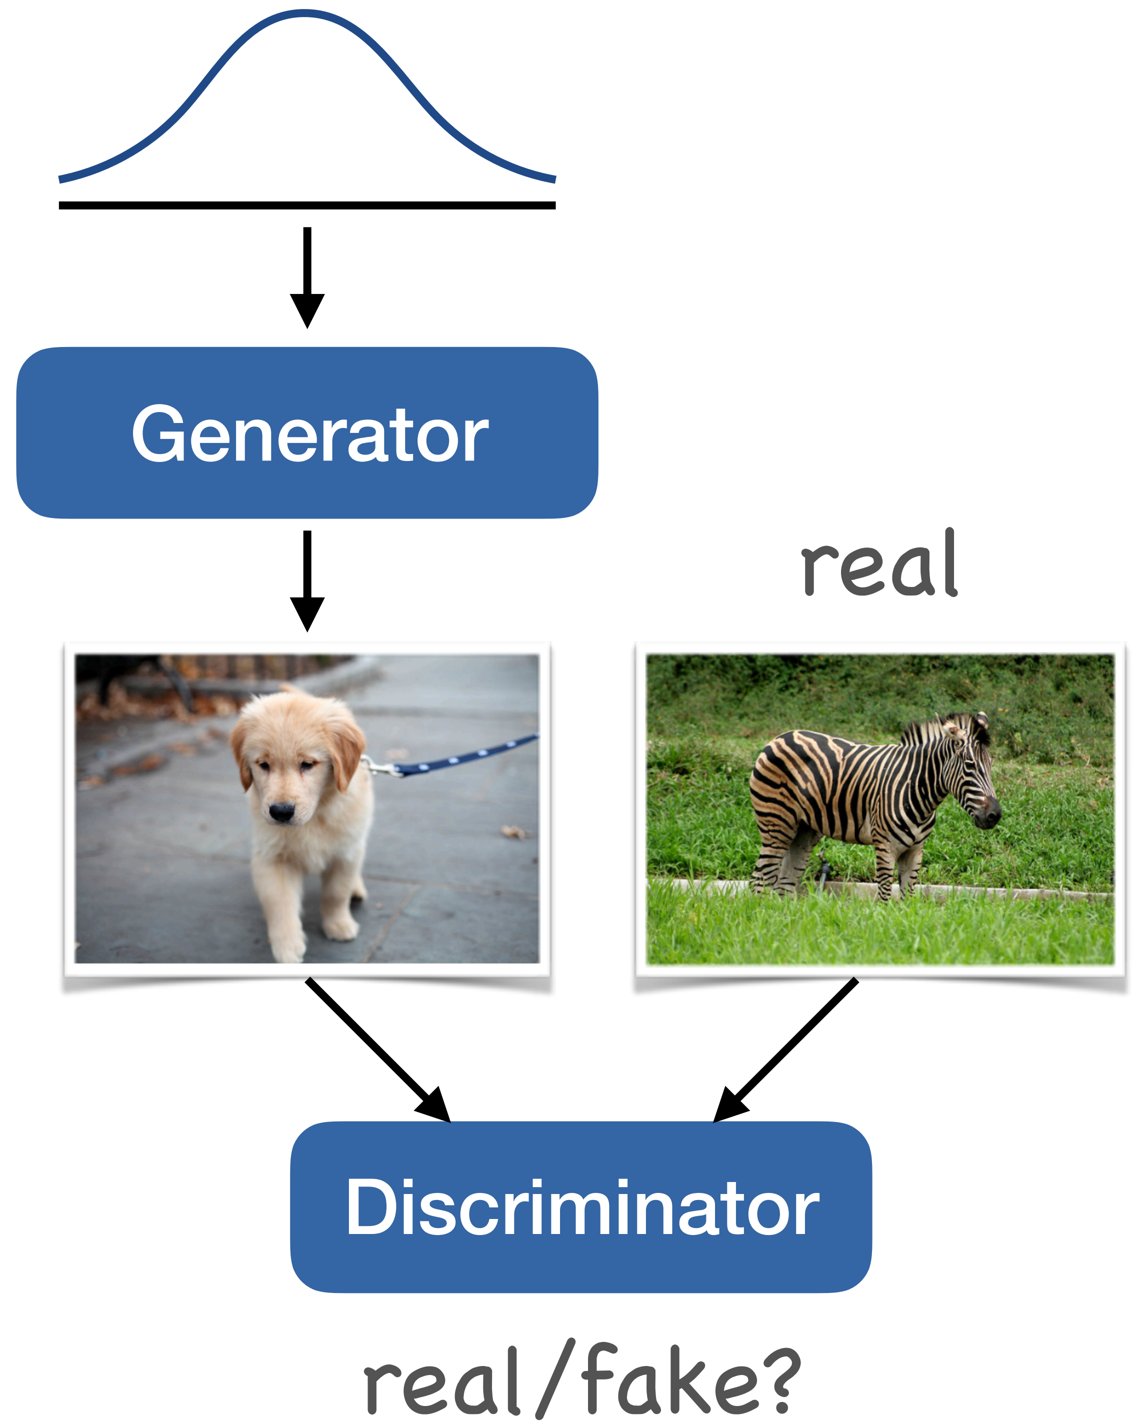

Generative Adversarial Nets, Goodfellow et al., NIPS 2014

Adversarial objective

- Generator
  - Produce an image that fools discriminator
- Discriminator
  - Tell difference between generation and training data

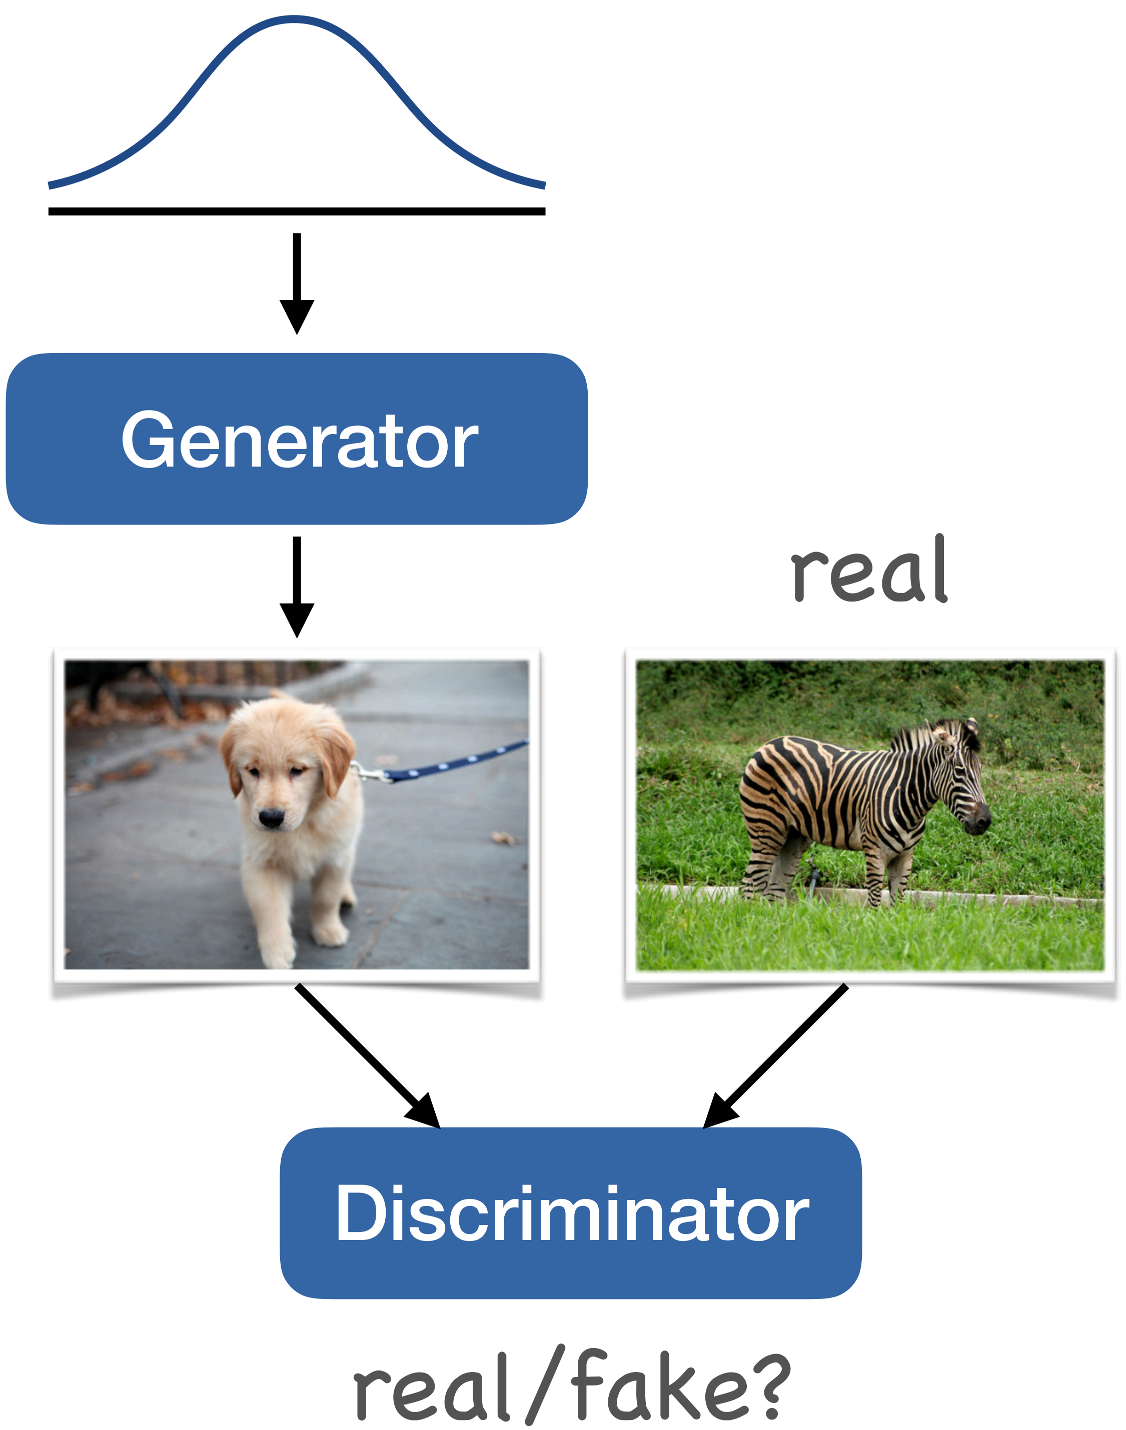

What does this optimize?

- Discriminator log-likelihood
  - Jensen-Shanon divergence data distribution and generator distribution

GANs work!

- Sampling is easy
- Learn pixel-distance
- Loss on distributions

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

Applications - Super resolution

- Learn to up-sample images
- Input: Low-res image
- Output: HD image
- Loss
    - Reconstruct HD
    - GAN for sharp reconstruction

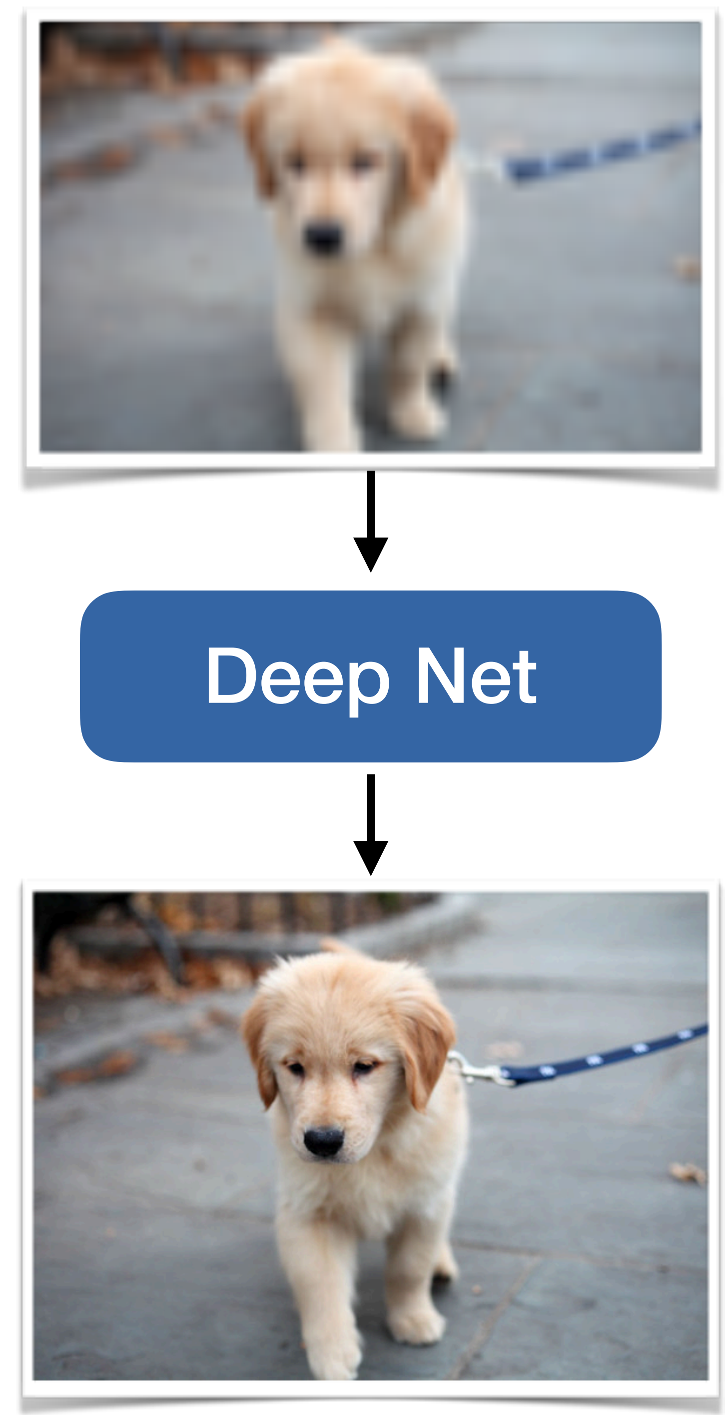

Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, Ledig et al., CVPR 2017

Applications - Text to image

- Input: text
- Output: Image
- Loss:
  - GAN
    - Does text and image fit or not?

"This bird has a yellow belly and tarsus, grey back, wings, and brown throat, nape with a black face"

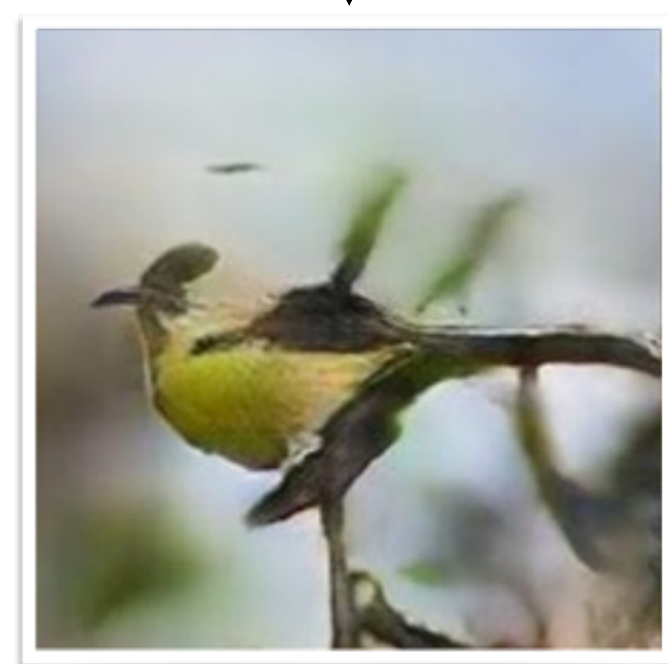

Image source: StackGAN paper

StackGAN: Text to Photo-realistic Image Synthesis with Stacked Generative Adversarial Networks, Zhang et al., ICCV 2017