

Autoencoder

© 2019 Philipp Krähenbühl and Chao-Yuan Wu

Autoencoder


- Image to image
 - with bottleneck



Encoder

□□□□ Bottleneck

Decoder



Reducing the Dimensionality of Data with Neural Networks, Hinton and Salakhutdinov, Science, 2006

What does an autoencoder learn?


- Compression
- Invertible mapping
- Does it learn to understand the image?
- Only in the limit / best compression



Encoder

□□□□ Bottleneck

Decoder



How do we sample from an autoencoder?

- Random samples?
 - Like never seen during training
 - Might produce garbage

Alternative that works

- Deep image prior
- Learn decoder of autoencoder
- Fixed random input
- Learns to denoise

