


Autoencoders

- Autoencoders are designed to reproduce their input, especially for images.
 - Key point is to reproduce the input from a learned encoding.

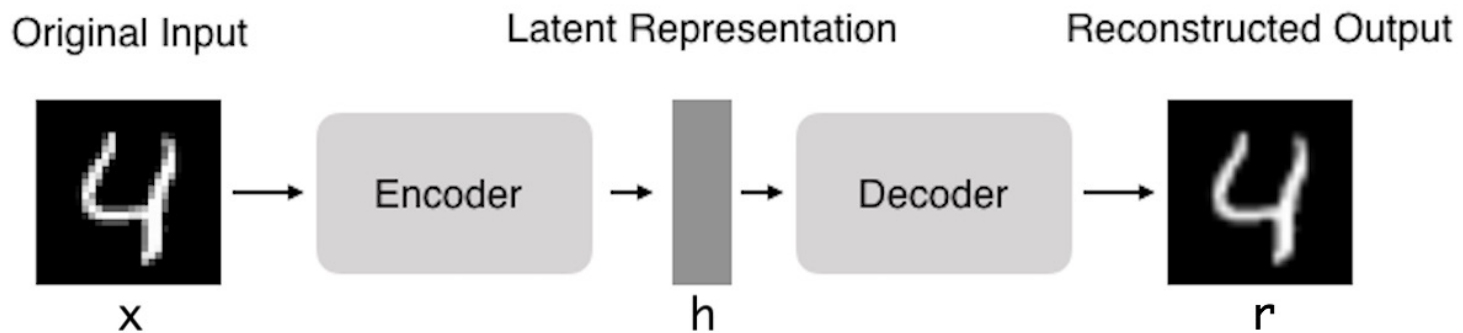


Autoencoders

- Compare PCA/SVD
 - PCA takes a collection of vectors (images) and produces a usually smaller set of vectors that can be used to approximate the input vectors via linear combination.
 - Very efficient for certain applications.
 - Fourier and wavelet compression is similar.
- Neural network autoencoders
 - Can learn nonlinear dependencies
 - Can use convolutional layers


Autoencoders: structure

- Encoder: compress input into a latent-space of usually smaller dimension. $h = f(x)$
- Decoder: reconstruct input from the latent space. $r = g(f(x))$ with r as close to x as possible



Autoencoders: applications

- Denoising: input clean image + noise and train to reproduce the clean image.




Autoencoders: Applications

- Watermark removal



Bottleneck layer (undercomplete)

- Suppose input images are $n \times n$ and the latent space is $m < n \times n$.
- Then the latent space is not sufficient to reproduce all images.
- Needs to learn an encoding that captures the important features in training data, sufficient for approximate reconstruction.



Simple bottleneck

- `input_img = Input(shape=(784,))`
- `encoding_dim = 32`
- `encoded = Dense(encoding_dim, activation='relu')(input_img)`
- `decoded = Dense(784, activation='sigmoid')(encoded)`
- `autoencoder = Model(input_img, decoded)`
- Maps 28x28 images into a 32 dimensional vector.
- Can also use more layers and/or convolutions.



Denoising autoencoders

- Basic autoencoder trains to minimize the loss between x and the reconstruction $g(f(x))$.
- Denoising autoencoders train to minimize the loss between x and $g(f(x+w))$, where w is random noise.
- Same possible architectures, different training data.
- [Kaggle has a dataset on damaged documents.](#)



