CAP5415 Computer Vision

Yogesh S Rawat

yogesh@ucf.edu

HEC-241

Autoencoder

Lecture 10

Autoencoder

- Reproduce the input
  - Via learning features
- Unsupervised learning
  - Efficient way to learn features
  - Still need a loss function – implicit supervision
- Supervised learning
  - Need labels/annotations

Autoencoder

- Encoder – decoder
- Encoding
  - Key idea

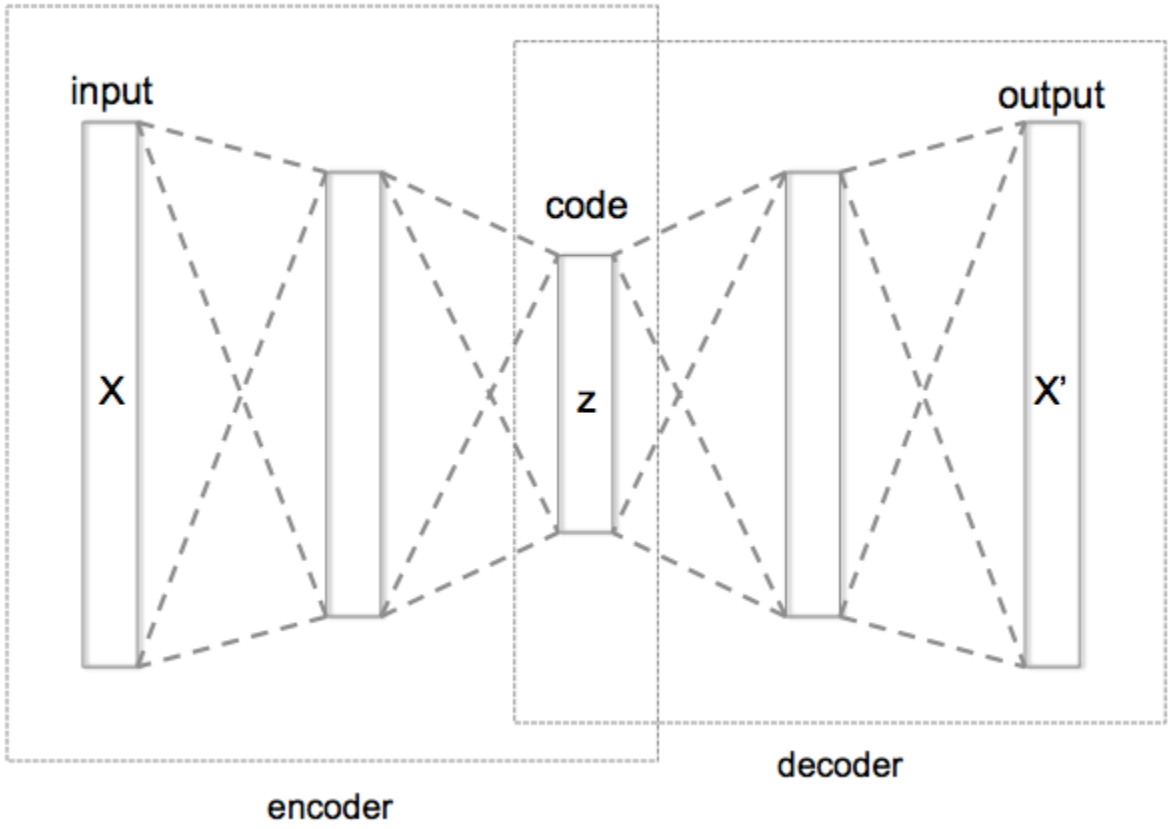

Autoencoder

- Compare PCA/SVD
  - PCA produce smaller set of vectors
  - Approximate the input vectors via linear combination.
  - Very efficient for certain applications.
- Autoencoder
  - Can learn nonlinear dependencies
  - Can use convolutional layers
  - Can use transfer learning

Autoencoder

- Encoder: h = f(x)
  - Compress input into a latent-space
  - Usually smaller dimension
- Decoder: r = g(f(x))
  - Reconstruct input from the latent space

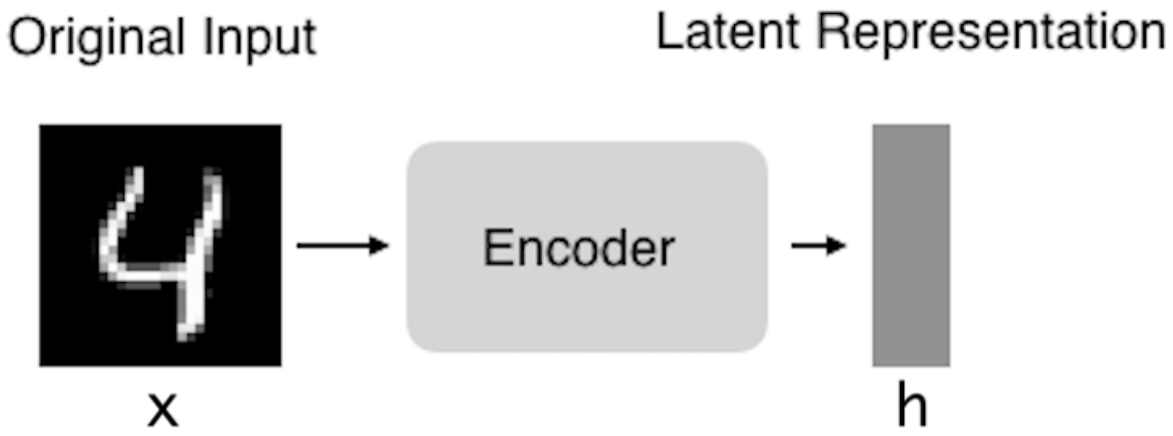

Autoencoder

- Encoder: h = f(x)
  - Compress input into a latent-space
  - Usually smaller dimension
- Decoder: r = g(f(x))
  - Reconstruct input from the latent space

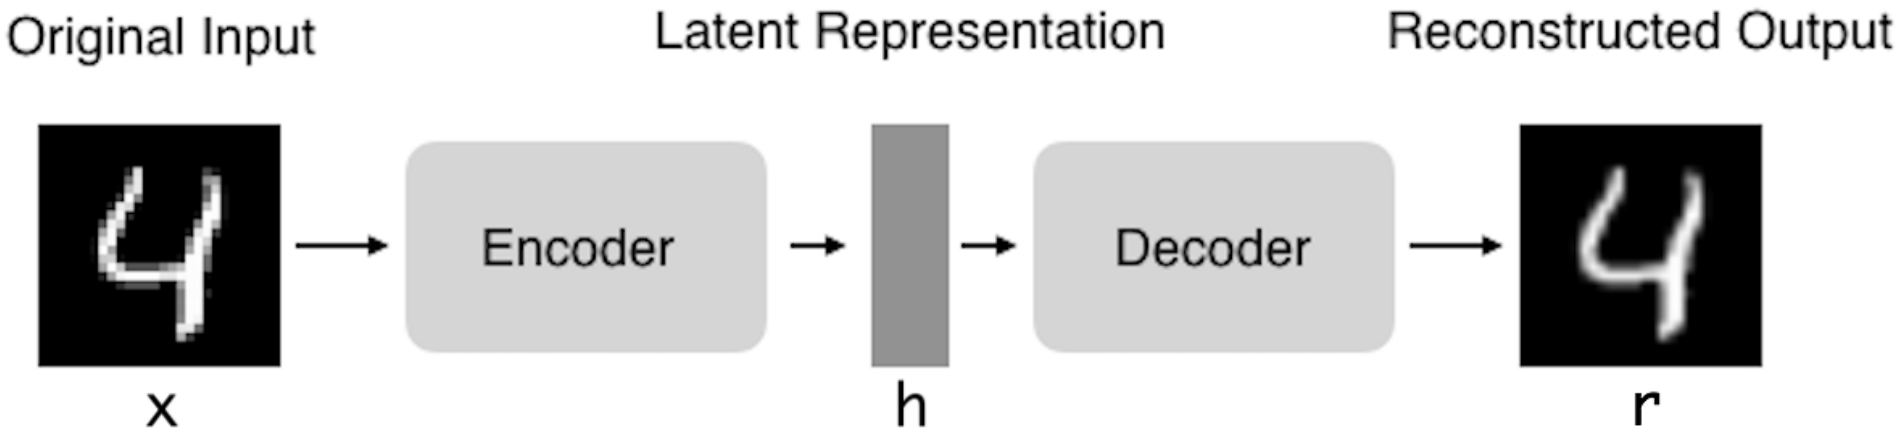

Autoencoder

- Shallow

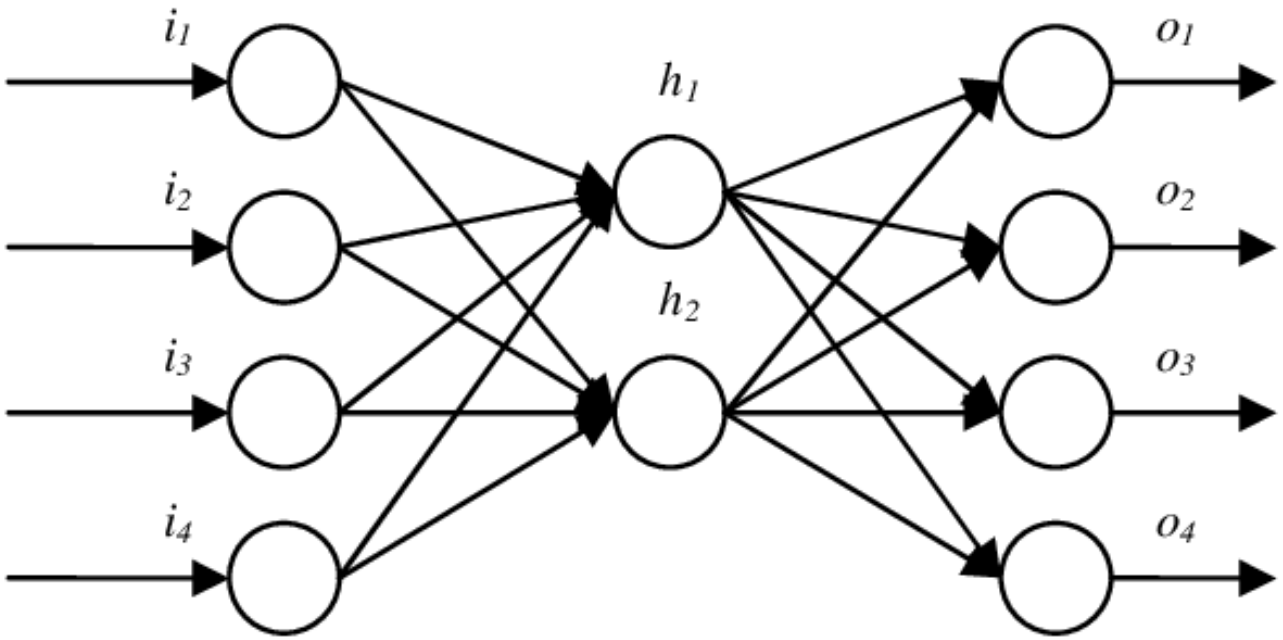

Autoencoder

- Deep

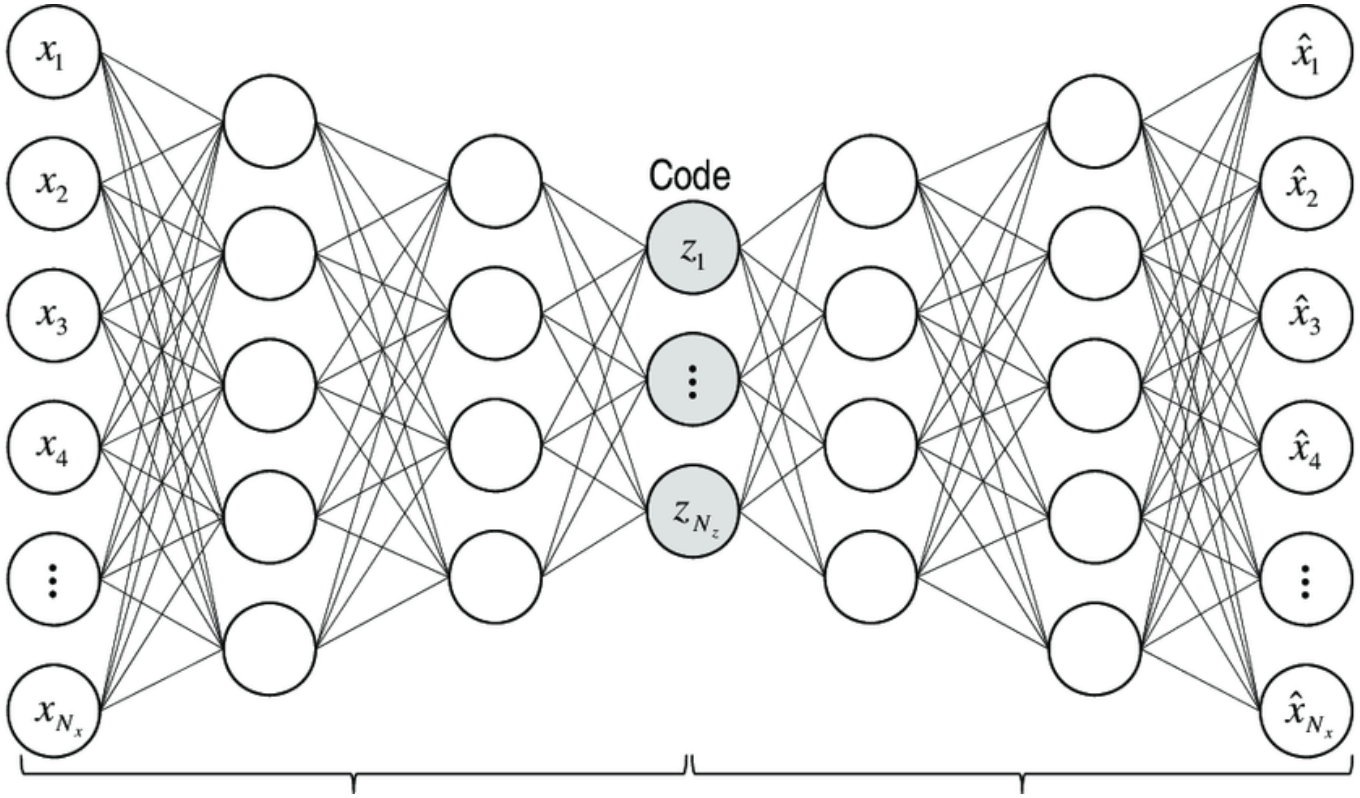

Autoencoder

- CNN

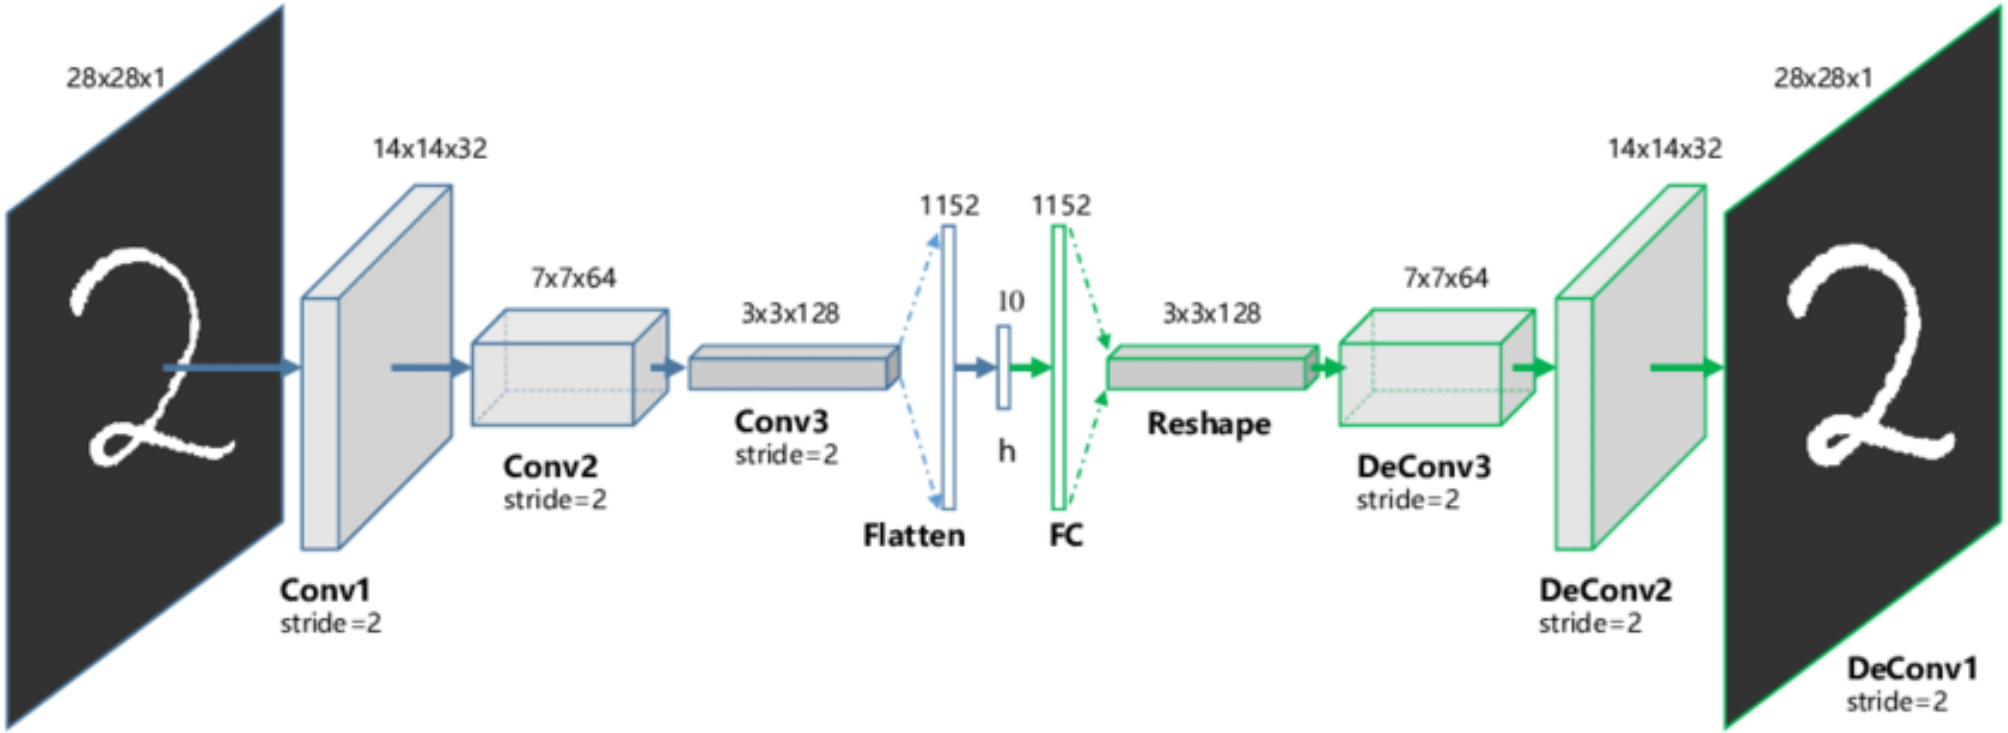

Autoencoder

- Reconstruction
  - Latent vector of size 2
  - Compression from 28x28

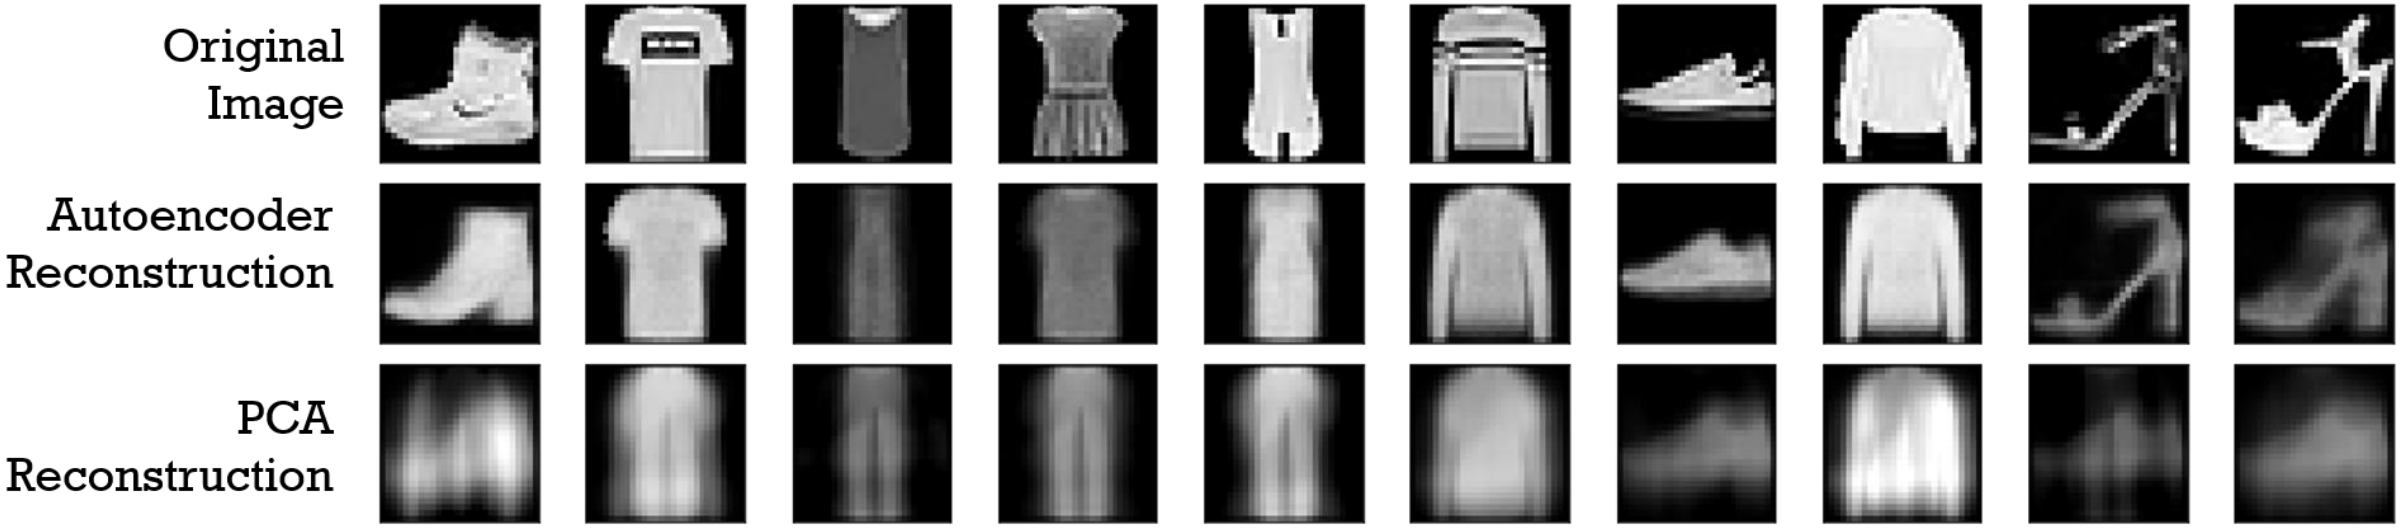

Feature learning

- Define a loss function
  - MSE, CE, etc.
- Optimize

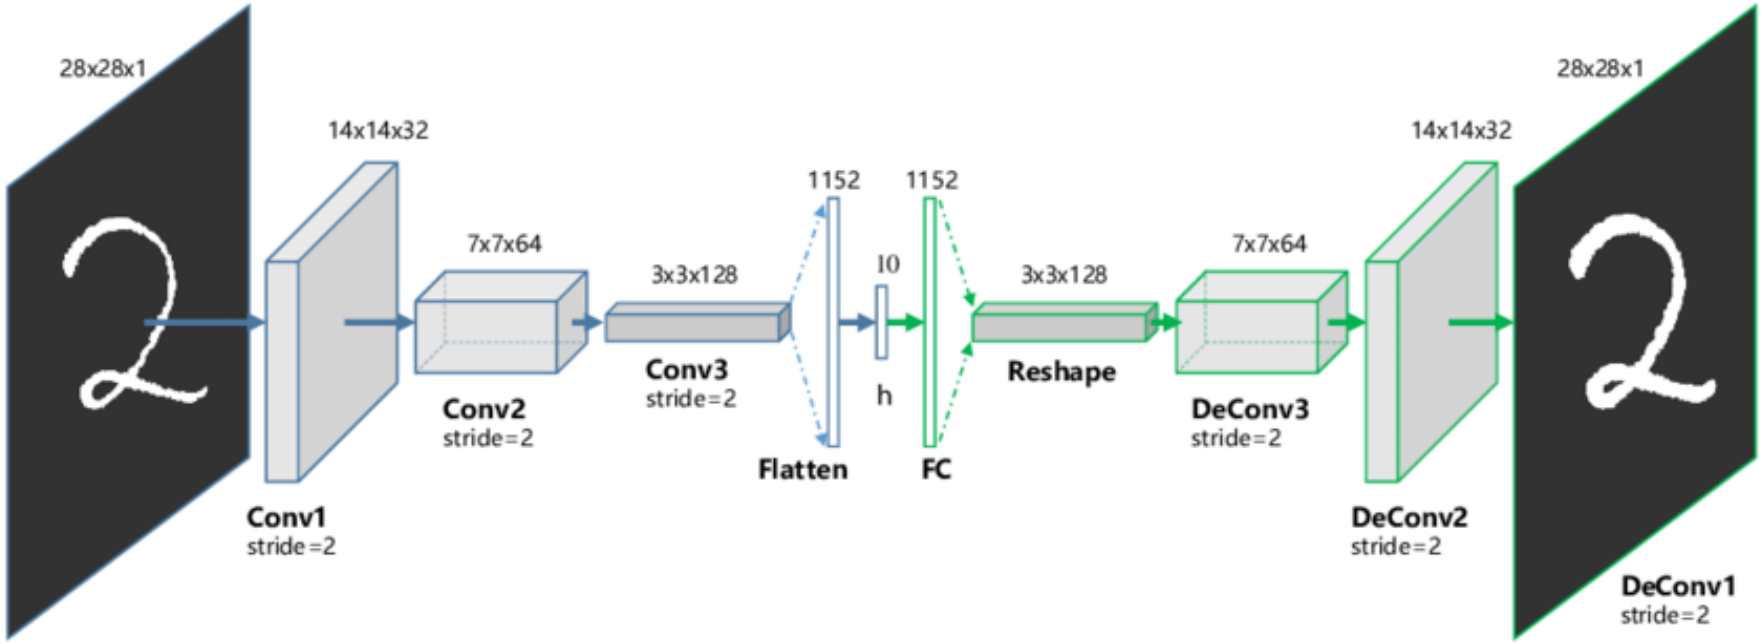

Feature learning

- Image retrieval
  - Dimensionality reduction helps

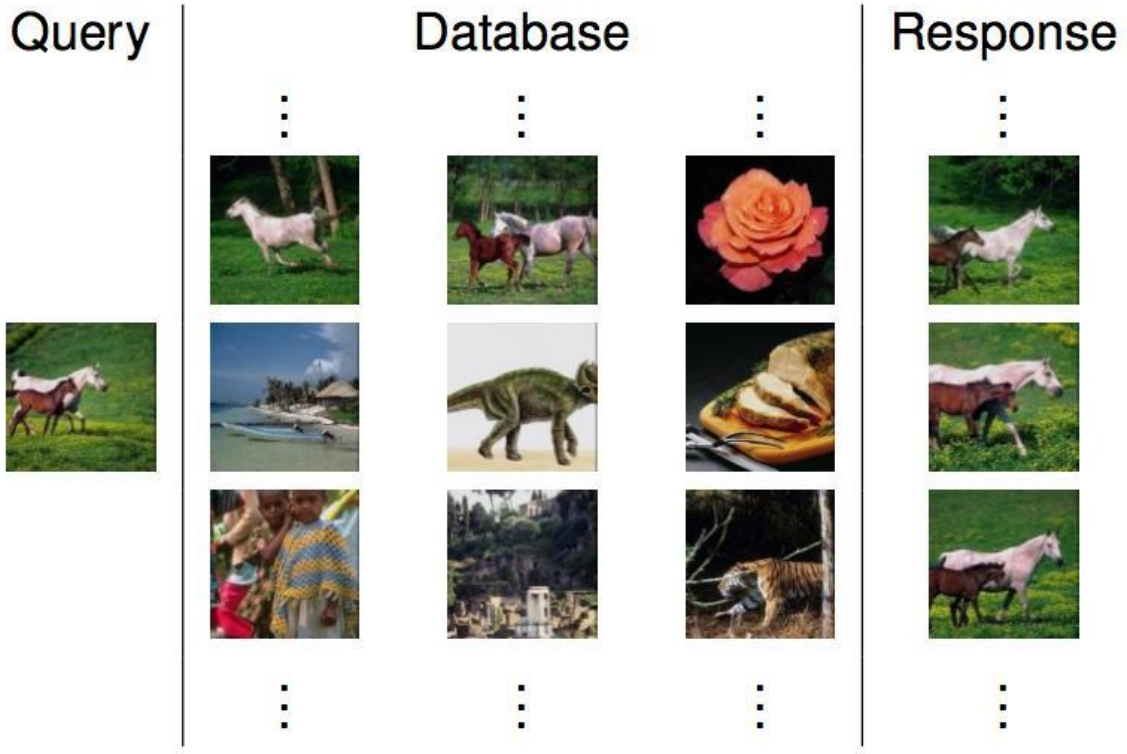

Autoencoder – application

- Denoising

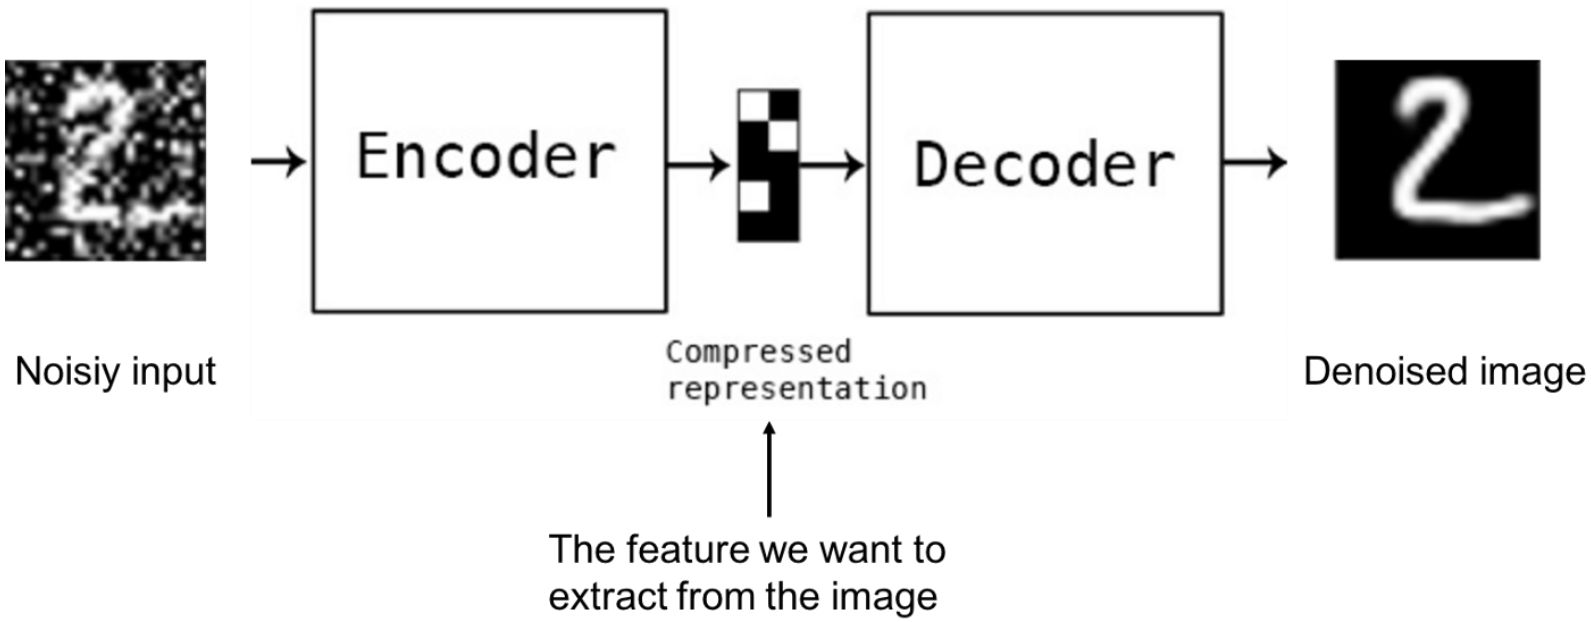

Autoencoder – application

- Image colorization





Autoencoder – application

- Anomaly detection





Properties

- Data-specific
  - Compress data similar to what they have been trained on
- Lossy
  - Outputs will be degraded compared to the original inputs
- Learned automatically from examples
  - It is easy to train
  - It will perform well on data similar to training samples
- Compare with hand-crafted features

Questions?