Autoencoders & Dimensionality Reduction

Sauman Das

Overview

- What is the purpose of Dimensionality Reduction?
- What is an autoencoder?
- How to train an autoencoder?
- Lab + Cool Visualization

Dimension of data is the **number of features per input**.

What is the dimensionality of the follow piece of data?

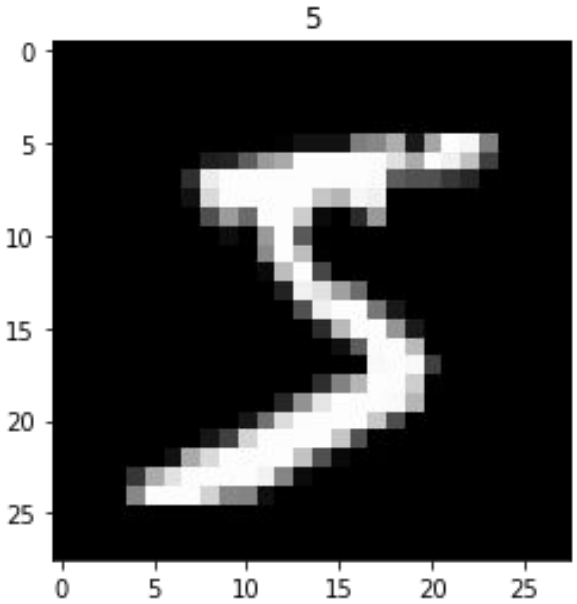

28x28 Sized Image

The Curse of Dimensionality

- The curse of dimensionality indicates that the number of samples needed to estimate an arbitrary function with a given level of accuracy **grows exponentially** with respect to the **number of input variables (i.e., dimensionality)** of the function.

Task of the Day

Represent a 28x28 image with a 2-dimensional vector

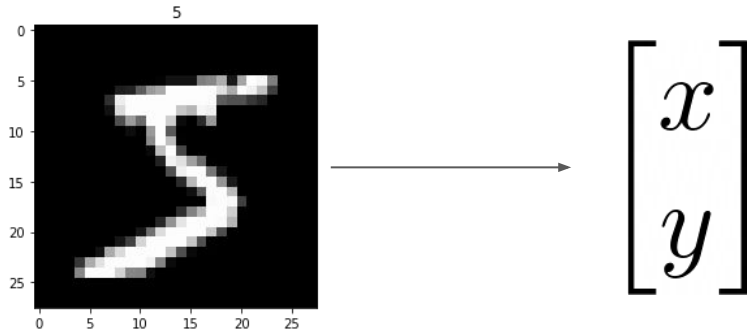

Variational Autoencoders

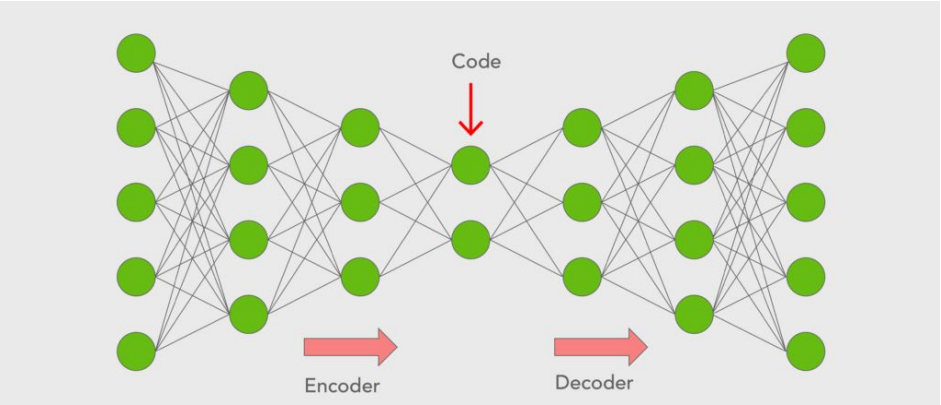

Architecture

- Input: Image
- Output: Image

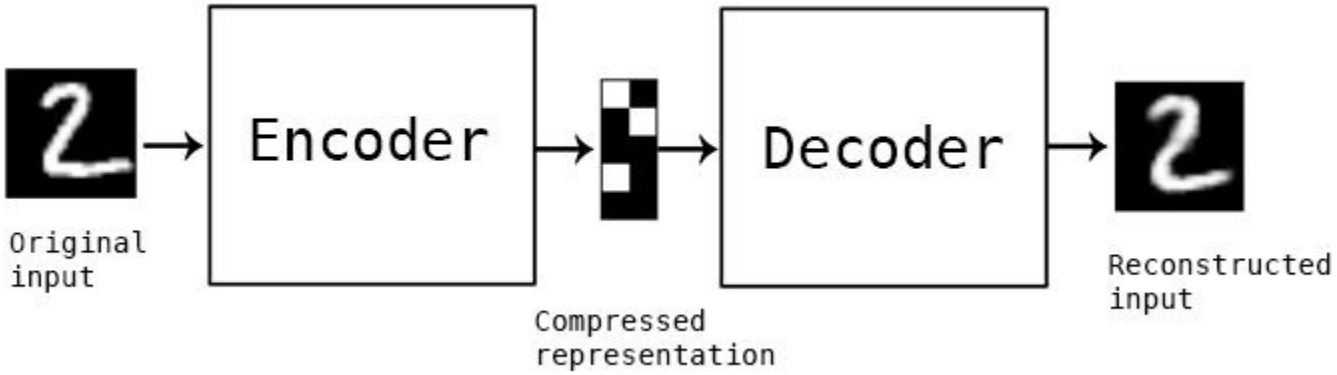

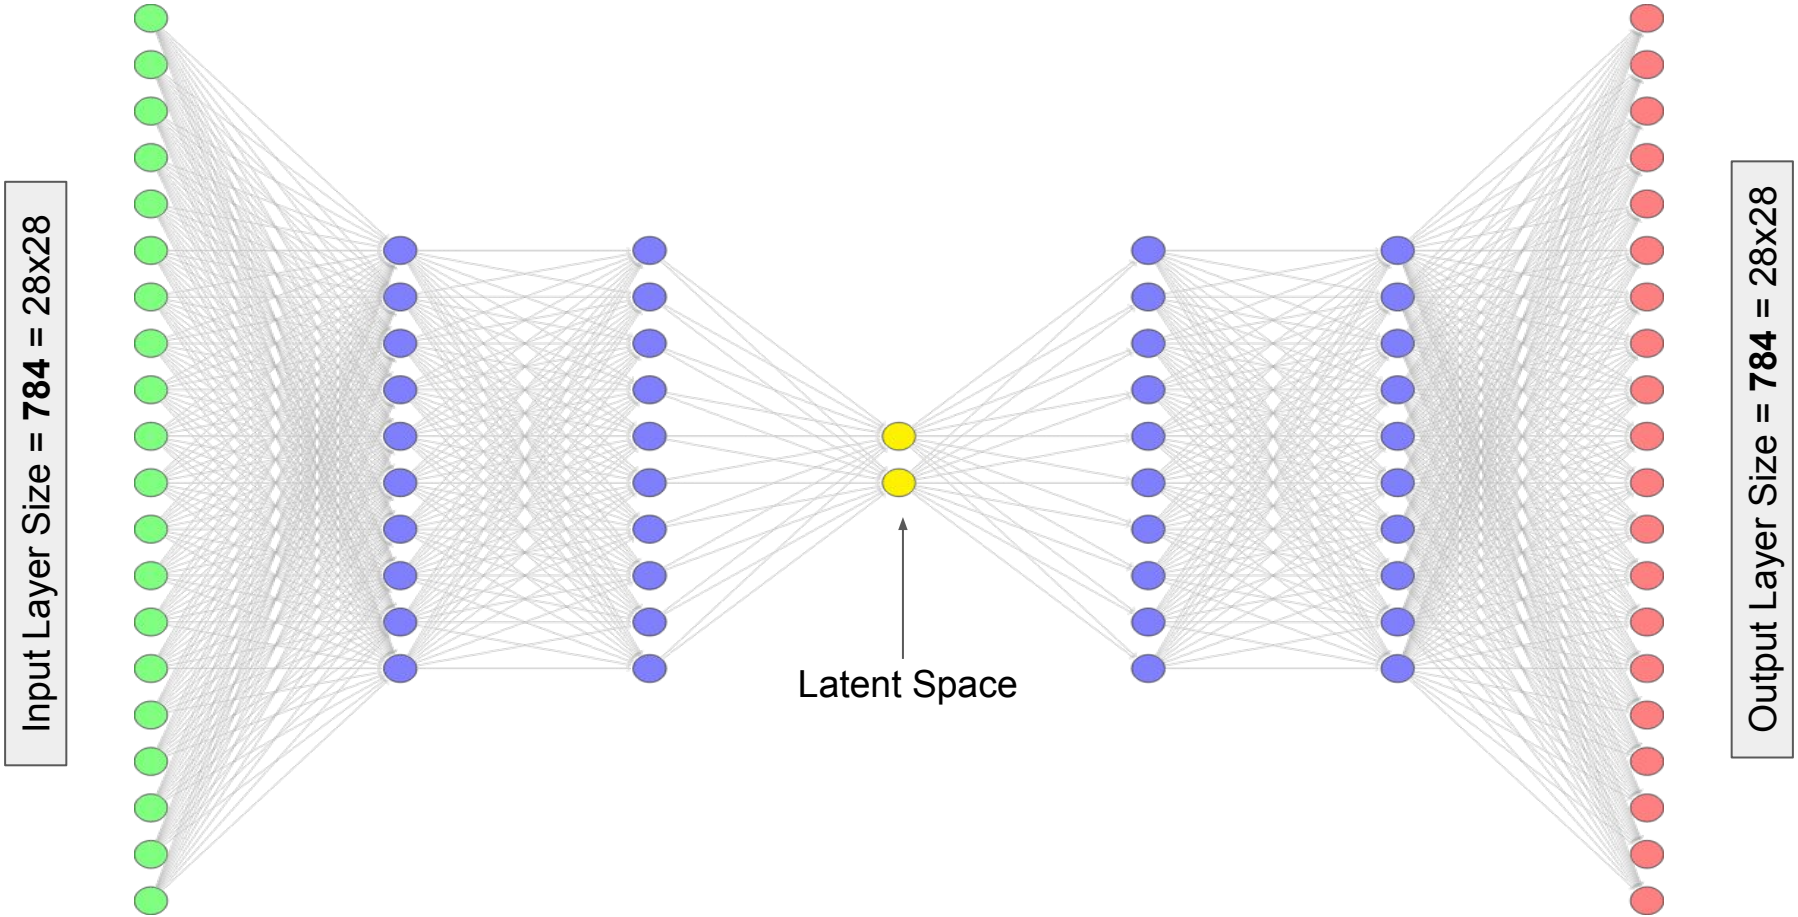
Input Layer Size = **784** = 28x28
Latent Space
Output Layer Size = **784** = 28x28

Why is reconstructing the original vector more difficult than it seems?

Why can’t the model just remember the values from the input layer?

Guess the Shadow

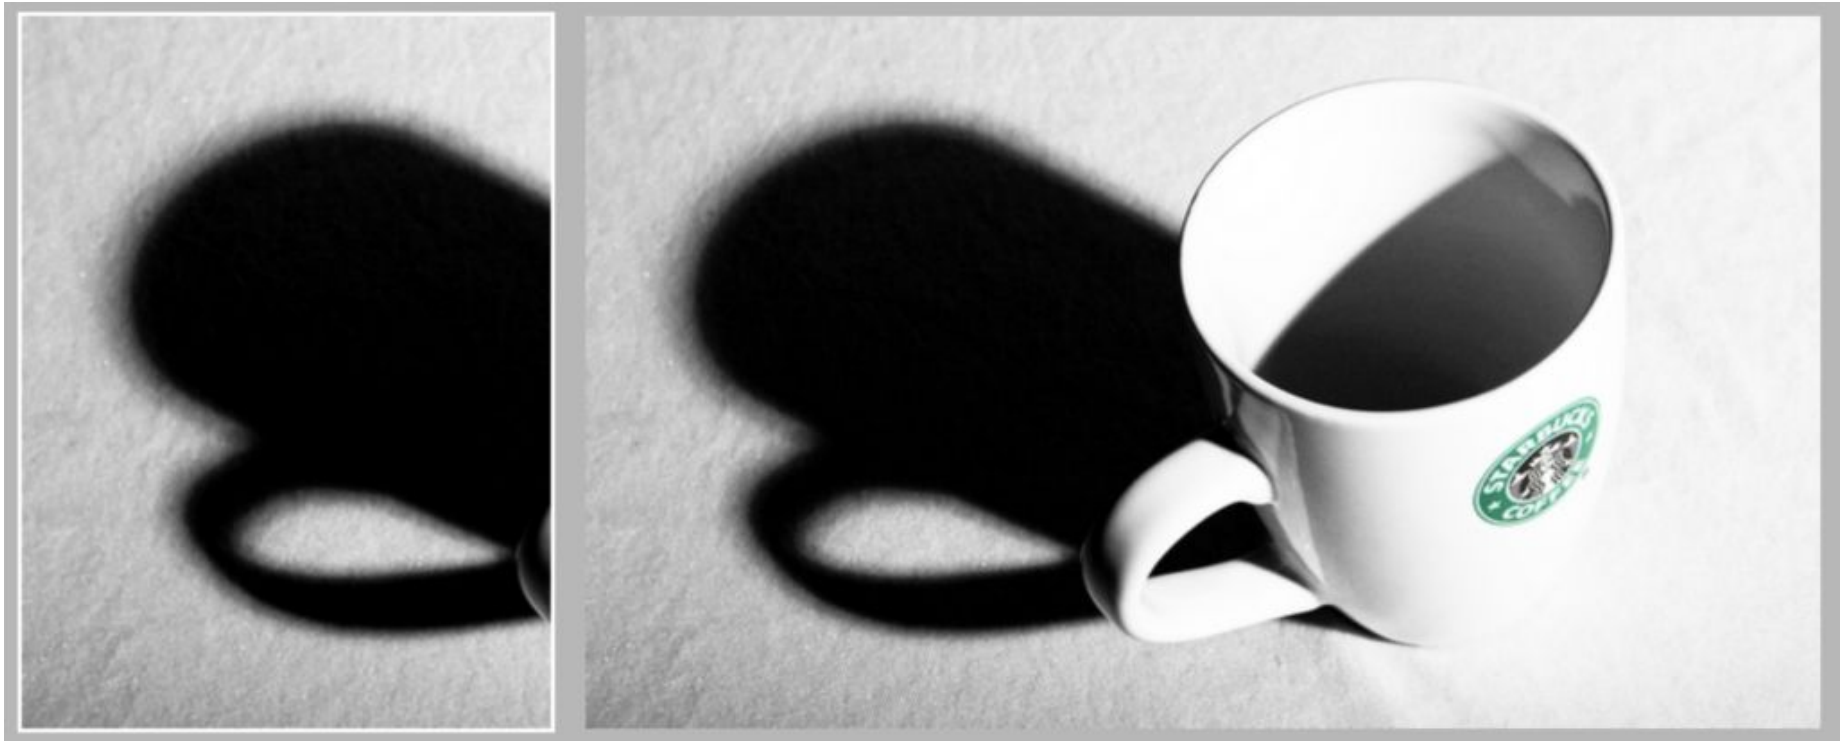

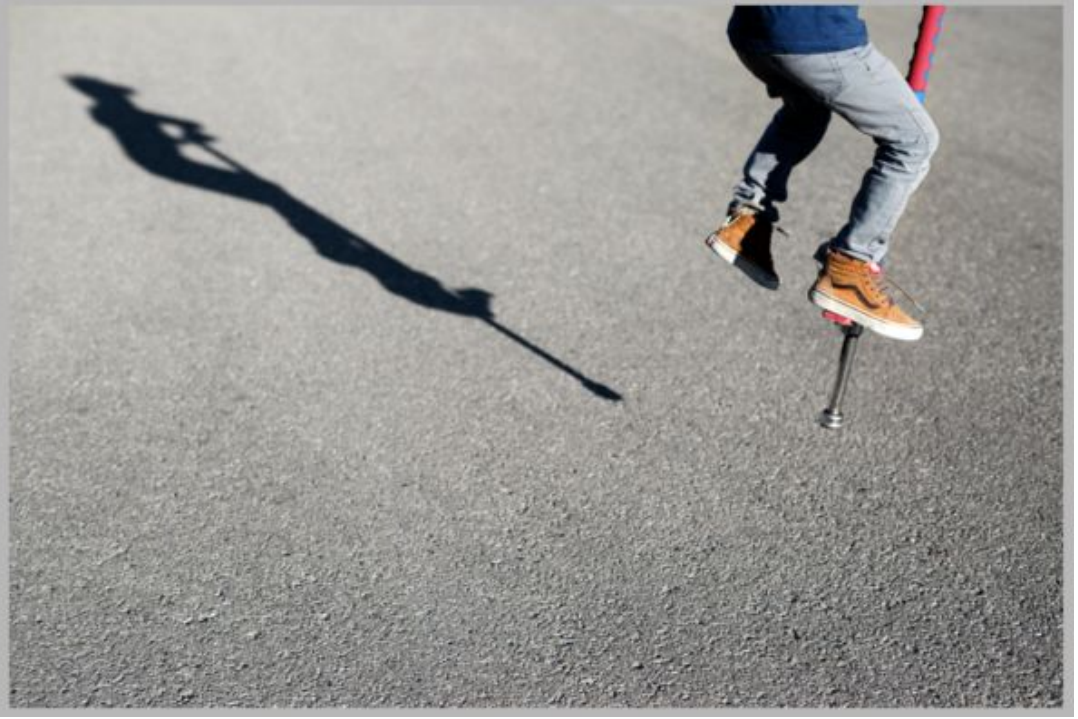

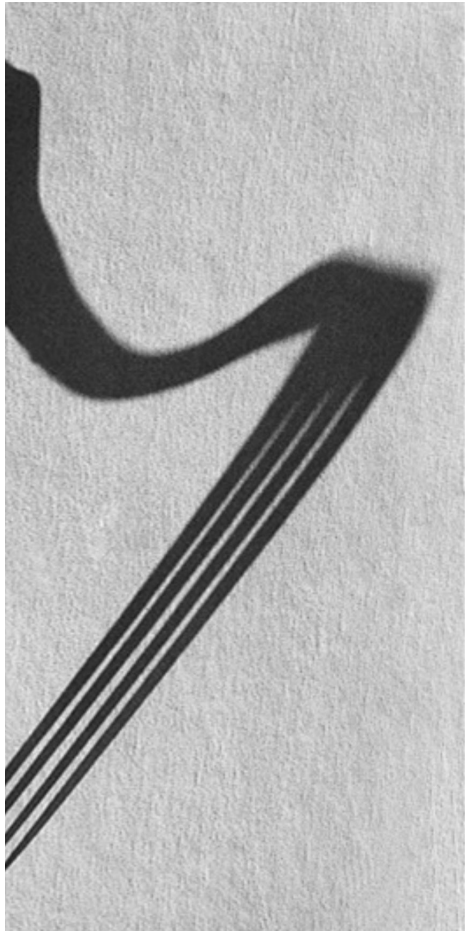

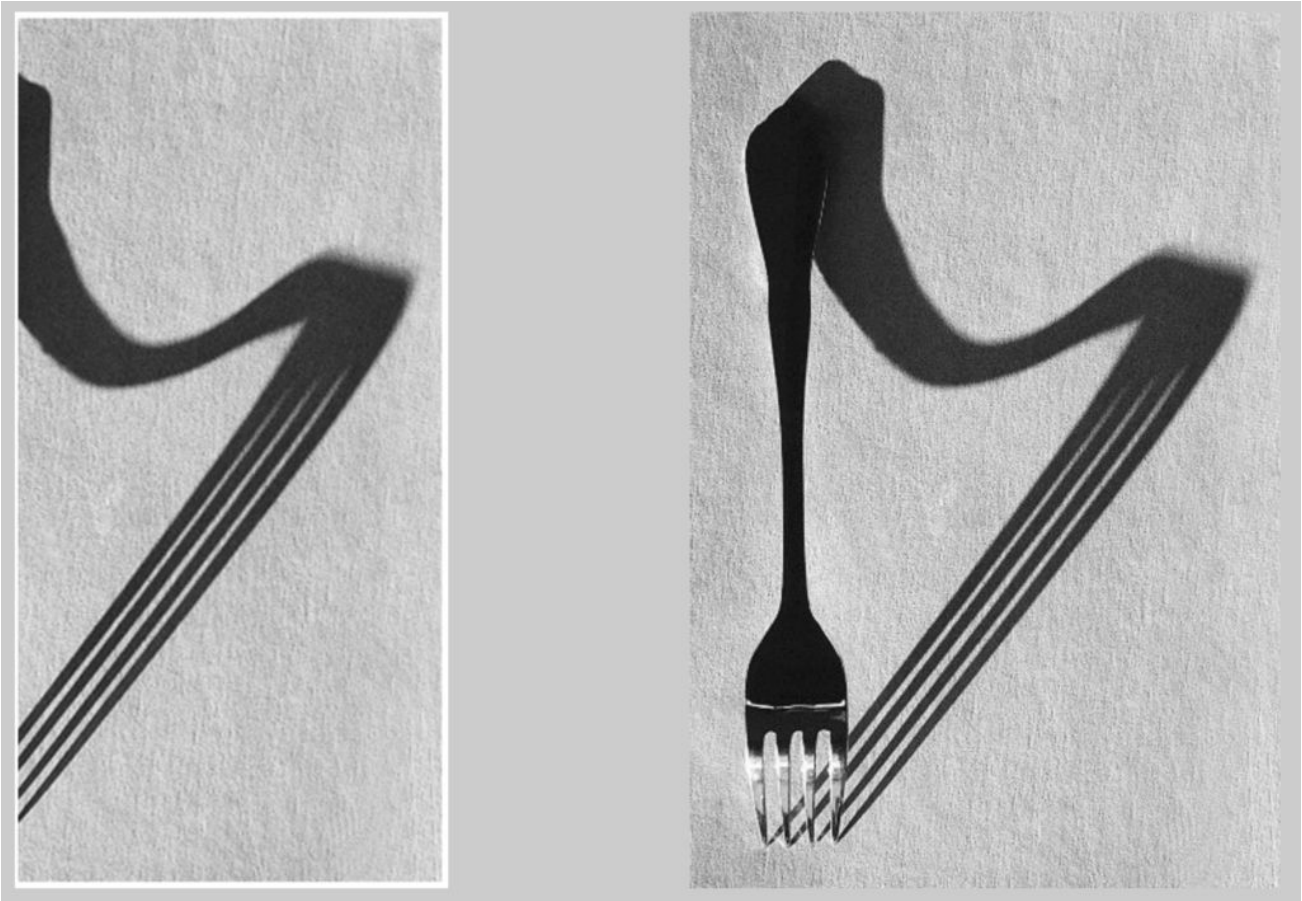

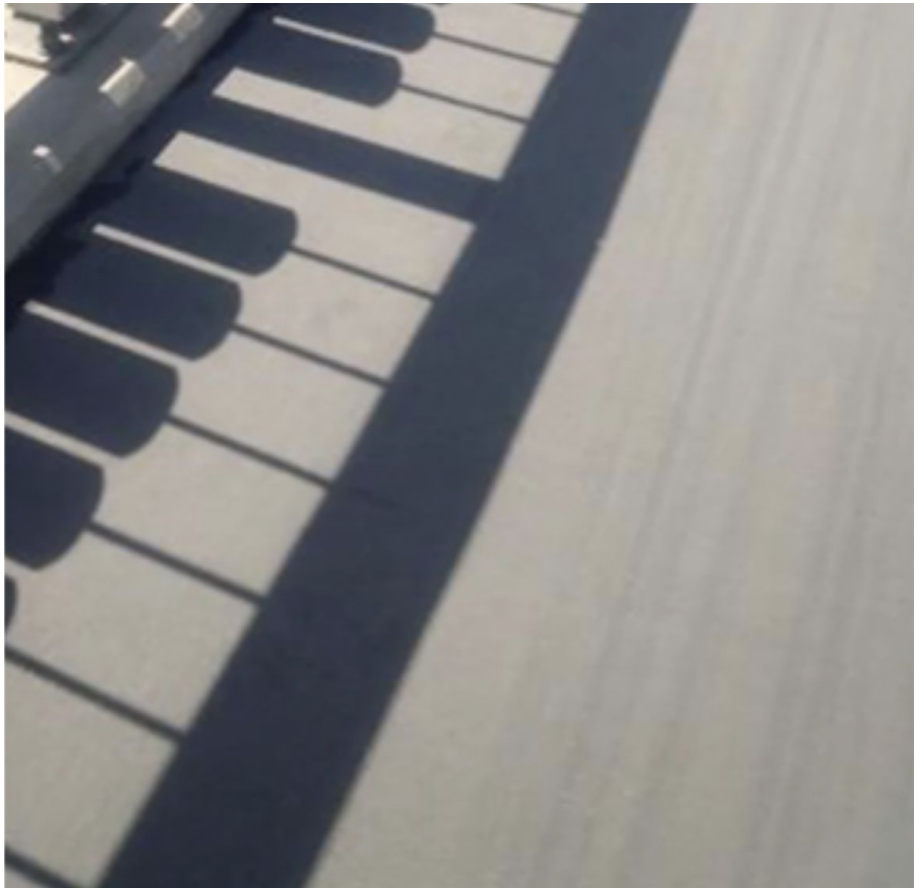

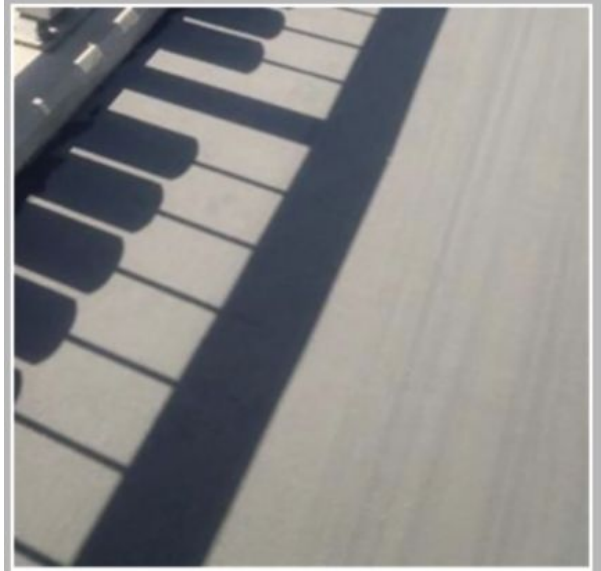

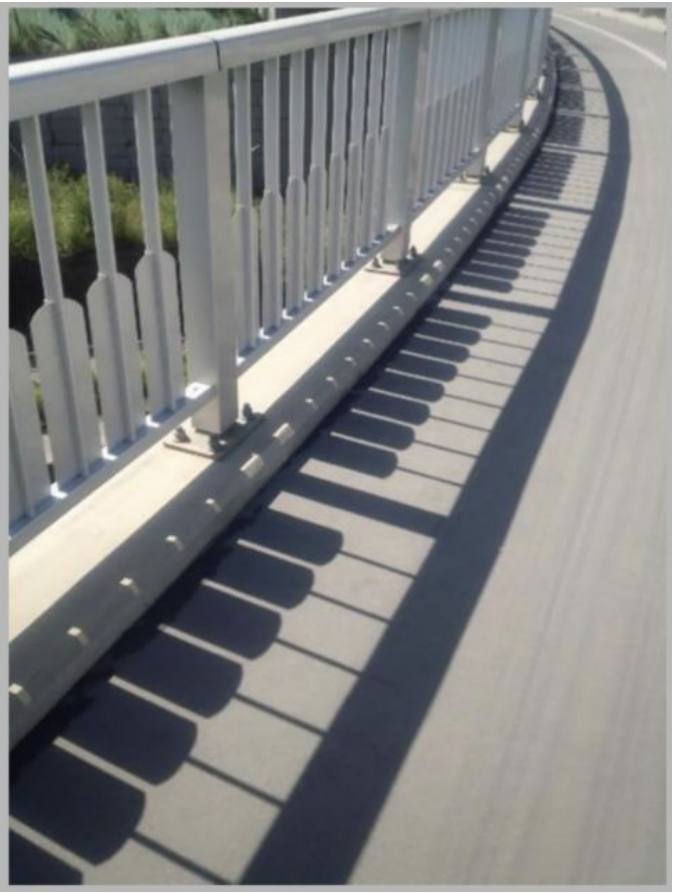

Latent Space is like the Shadow

- Even though the model starts with the original image, eventually the model has to be able to figure out the original image only from the latent space.

Encoder + Decoder

- Two Components of Autoencoder → Encoder + Decoder

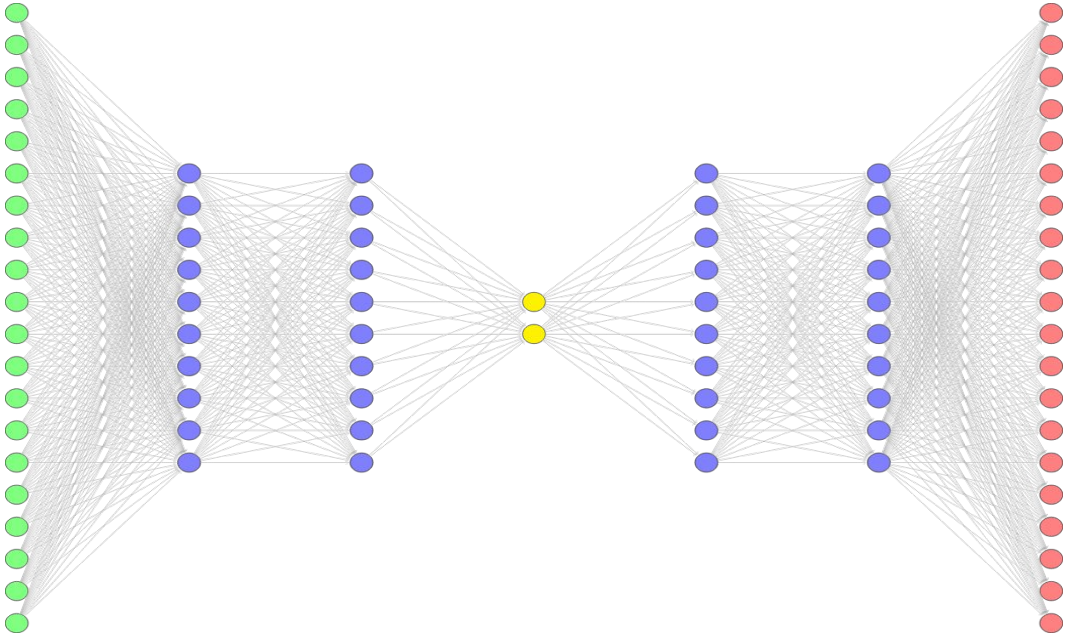

Constructing an Autoencoder in Code

- **Two Sequential Neural Network Model (i.e. Encoder, Decoder)**
  - Each layer connects to the previous layer
  - All Layers are Dense
- **Why use two separate models?**
  - At the end, we throw away the decoder. We are only concerned with getting the latent representation.

Autoencoder Lab

Lab Task

Represent a 28x28 image with a 2-dimensional vector

Other Applications of Autoencoders

- Variational Autoencoders
  - Image Generation
- Denoising Autoencoders
- Semi-Supervised Pretraining Autoencoders