COMP 642 — Machine Learning March 10, 2022

Lecture 8

Lecturer: Zhaozhuo Xu *Scribe By: Abhishek Shah, Aditya Dave , Durga Parulekar, Sacheth Reddy*

Disclaimer: *These lecture notes are intended to develop the thought process and intuition in machine learning. The materials are not thoroughly reviewed and can contain errors.*

1 Supervised Machine Learning vs Unsupervised Machine Learning

1.1 Supervised Machine Learning

1.1.1 Definition

Supervised machine learning is defined by its use of labeled data sets to train algorithms that to classify data or predict outcomes accurately.

1.1.2 Data

The data in supervised machine learning is labelled

1.1.3 Goal

: The goal of supervised machine learning is to learn a function that maps the data x to the correct label y.

1.1.4 Examples

Classification, regression, object detection, semantic segmentation, image captioning, etc.

1.2 Unsupervised Machine Learning

1.2.1 Definition

Unsupervised learning is a type of algorithm that learns patterns from untagged data, instead of predicting the corresponding label for the given data.

1.2.2 Data

The data in unsupervised machine learning is not labelled

1.2.3 Goal

The goal of unsupervised machine learning is to infer the patterns within datasets without reference to known or labeled outcomes.

1.2.4 Examples

Clustering, dimensional reduction, feature learning, density estimation.

2 PCA(Dimensionality Reduction)

PCA is an unsupervised machine learning technique for reducing the dimensionality of large datasets, increasing interpretability but at the same time minimizing information loss. Basically, we are transforming a large set of variables into smaller ones that still contain most of the information from the large set.

3 Autoencoders(Feature Learning)

An autoencoder is a type of convolutional neural network used to learn data encodings in an unsupervised manner. The aim of an autoencoder is to learn a lower-dimensional representation (encoding) for a higher-dimensional data, typically for dimensionality reduction, by training the network to capture the most important parts of the input image. The simplest implementation of an autoencoder has one activation function and behaves like PCA.

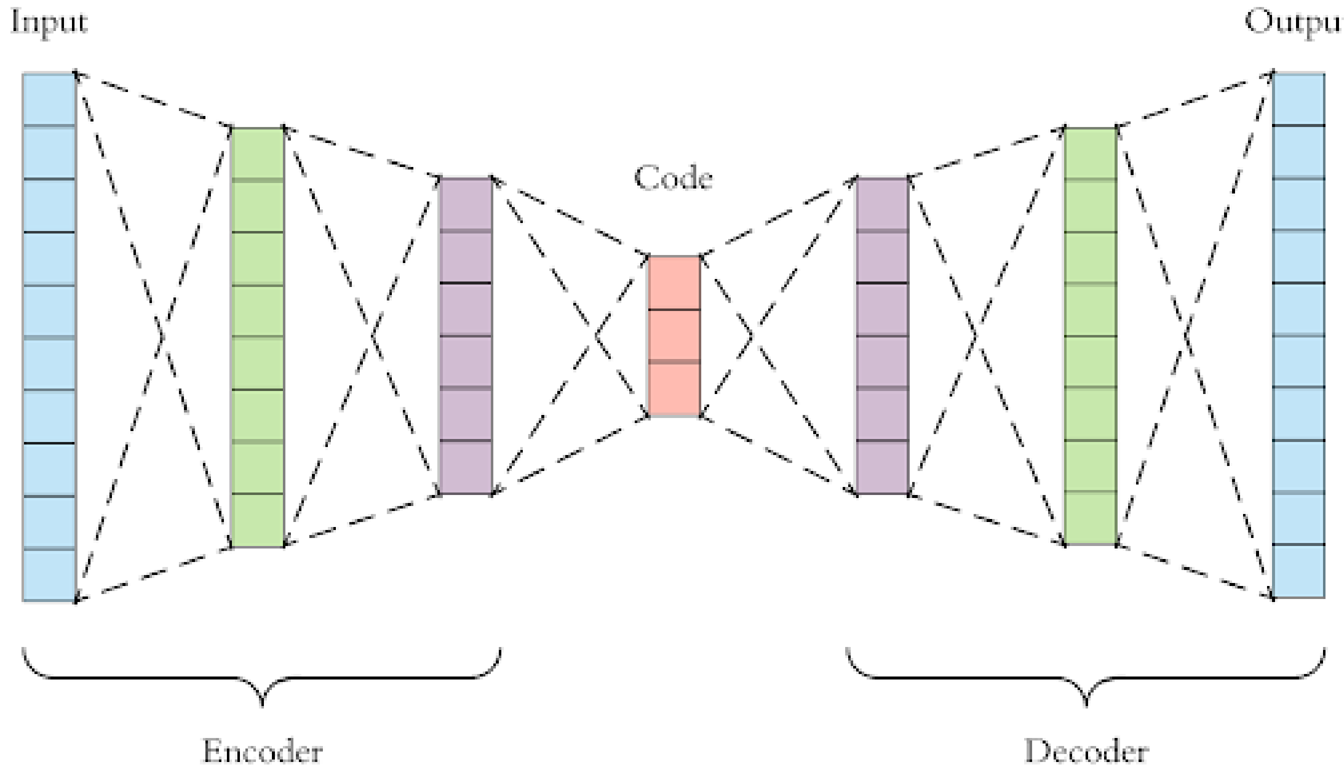

3.1 Example of Autoencoders

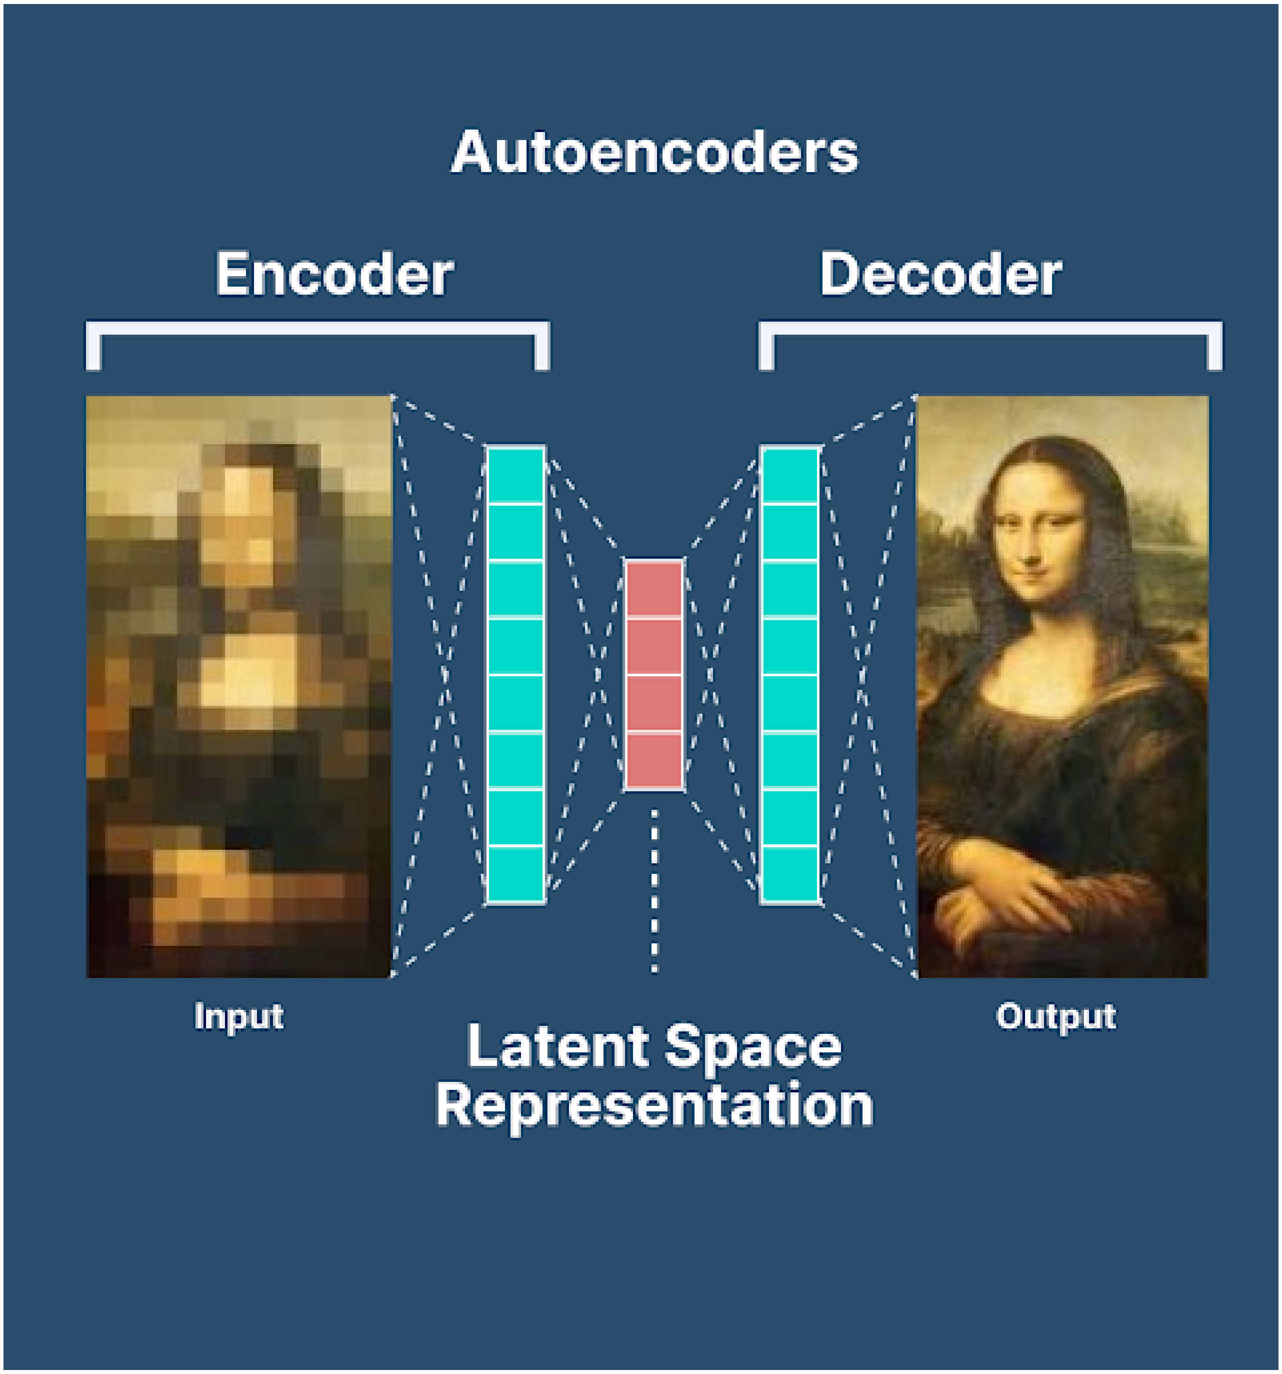

4 Generative Modeling

Generative modeling is an unsupervised machine learning technique that involves discovering regularities/patterns/similarities/features in the input data in such a way that the knowledge of them can be used to generate new samples from the input data in the same distribution.

5 Generative Adversarial Networks(GANs)

Generative modeling is an unsupervised learning task in machine learning that involves automatically discovering and learning the regularities or patterns in input data in such a way that the model can be used to generate or output new examples that plausibly could have been drawn from the original dataset. More generally, GANs are a model architecture for training a generative model, and it is most common to use deep learning models, such as convolutional neural networks in this architecture. GANs are usually trained to generate images from random noises and usually have two parts in which they work, namely the Generator that generates new samples from complex, high dimensional data, and the Discriminator that, given the training data and the sample generated by the generator classifies the sample as real or fake. The goal is that the discriminator should be able to distinguish between real and fake images with high accuracy. For example we can train a GAN model to generate digit images that look like hand-written digit images from the MNIST dataset and apart from this GANs are widely used for voice generation, image generation or video generation.

5.1 Learning in GAN

The goal of GAN is to maximize the expectation that the discriminator - D, correctly categorizes the data as either real or fake. Mathematically, the goal of learning is to minimize the following objective function

$$\min_{g} \max_{D} V(D, G) = \mathbb{E}_{x \sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))]$$

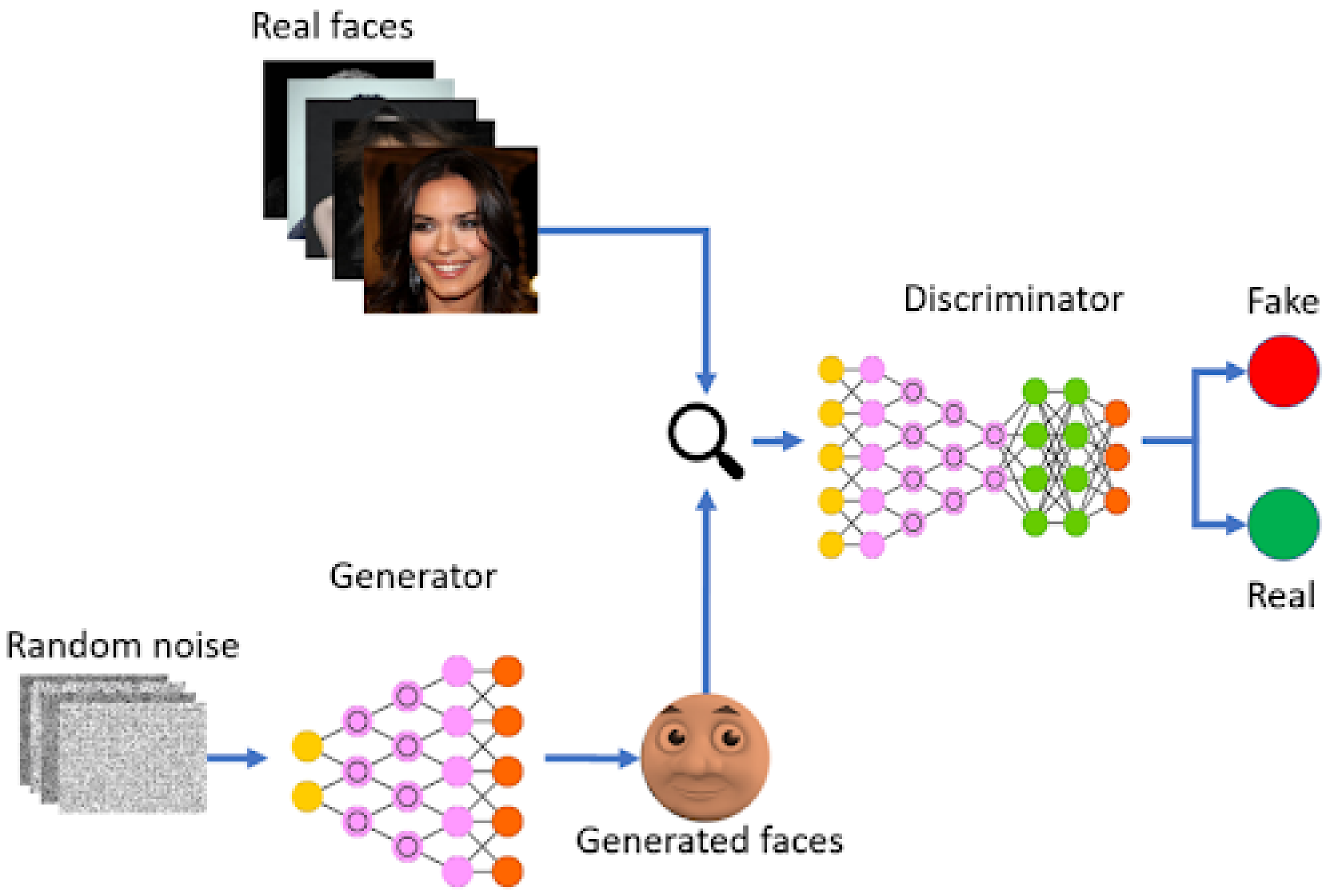

5.2 Applications

Examples of GAN applications include generating images, translating images, domain adaptation, imitation learning

References

https://www.v7labs.com/blog/autoencoders-guide

https://towardsdatascience.com/a-one-stop-shop-for-principal-component-analysis-5582fb7e0a9c

https://towardsdatascience.com/understanding-generative-adversarial-networks-gans-cd6e4651a29