Introduction to Deep Learning

Tushar B. Kute,
http://tusharkute.com

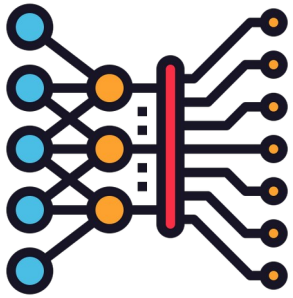

What is Deep Learning?

- Deep Learning is a subset of Machine Learning that uses mathematical functions to map the input to the output.
- These functions can extract non-redundant information or patterns from the data, which enables them to form a relationship between the input and the output.
- This is known as learning, and the process of learning is called training.

Programming Patterns

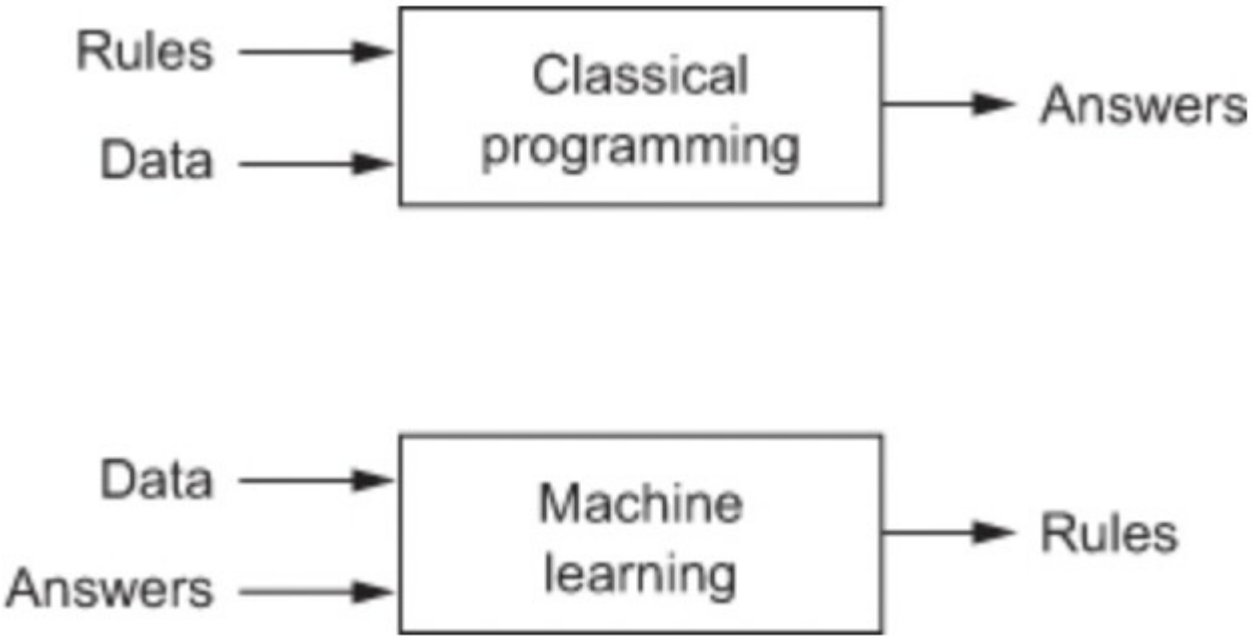

Reference: https://www.v7labs.com

Deep Learning

- Modern deep learning models use artificial neural networks or simply neural networks to extract information.
- These neural networks are made up of a simple mathematical function that can be stacked on top of each other and arranged in the form of layers, giving them a sense of depth, hence the term Deep Learning.
- Deep learning can also be thought of as an approach to Artificial Intelligence, a smart combination of hardware and software to solve tasks requiring human intelligence.

Deep Learning

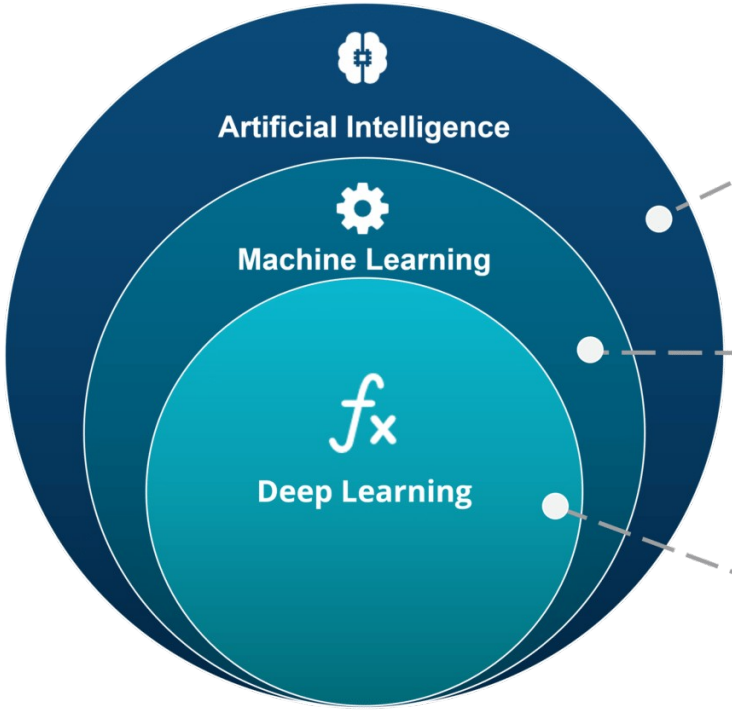

ARTIFICIAL INTELLIGENCE
A technique which enables machines to mimic human behaviour

MACHINE LEARNING
Subset of AI technique which use statistical methods to enable machines to improve with experience

DEEP LEARNING
Subset of ML which make the computation of multi-layer neural network feasible

Deep Learning

- Deep Learning was first theorized in the 1980s, but it has only become useful recently because:
  - It requires large amounts of labeled data
  - It requires significant computational power (high performing GPUs)

Neuron

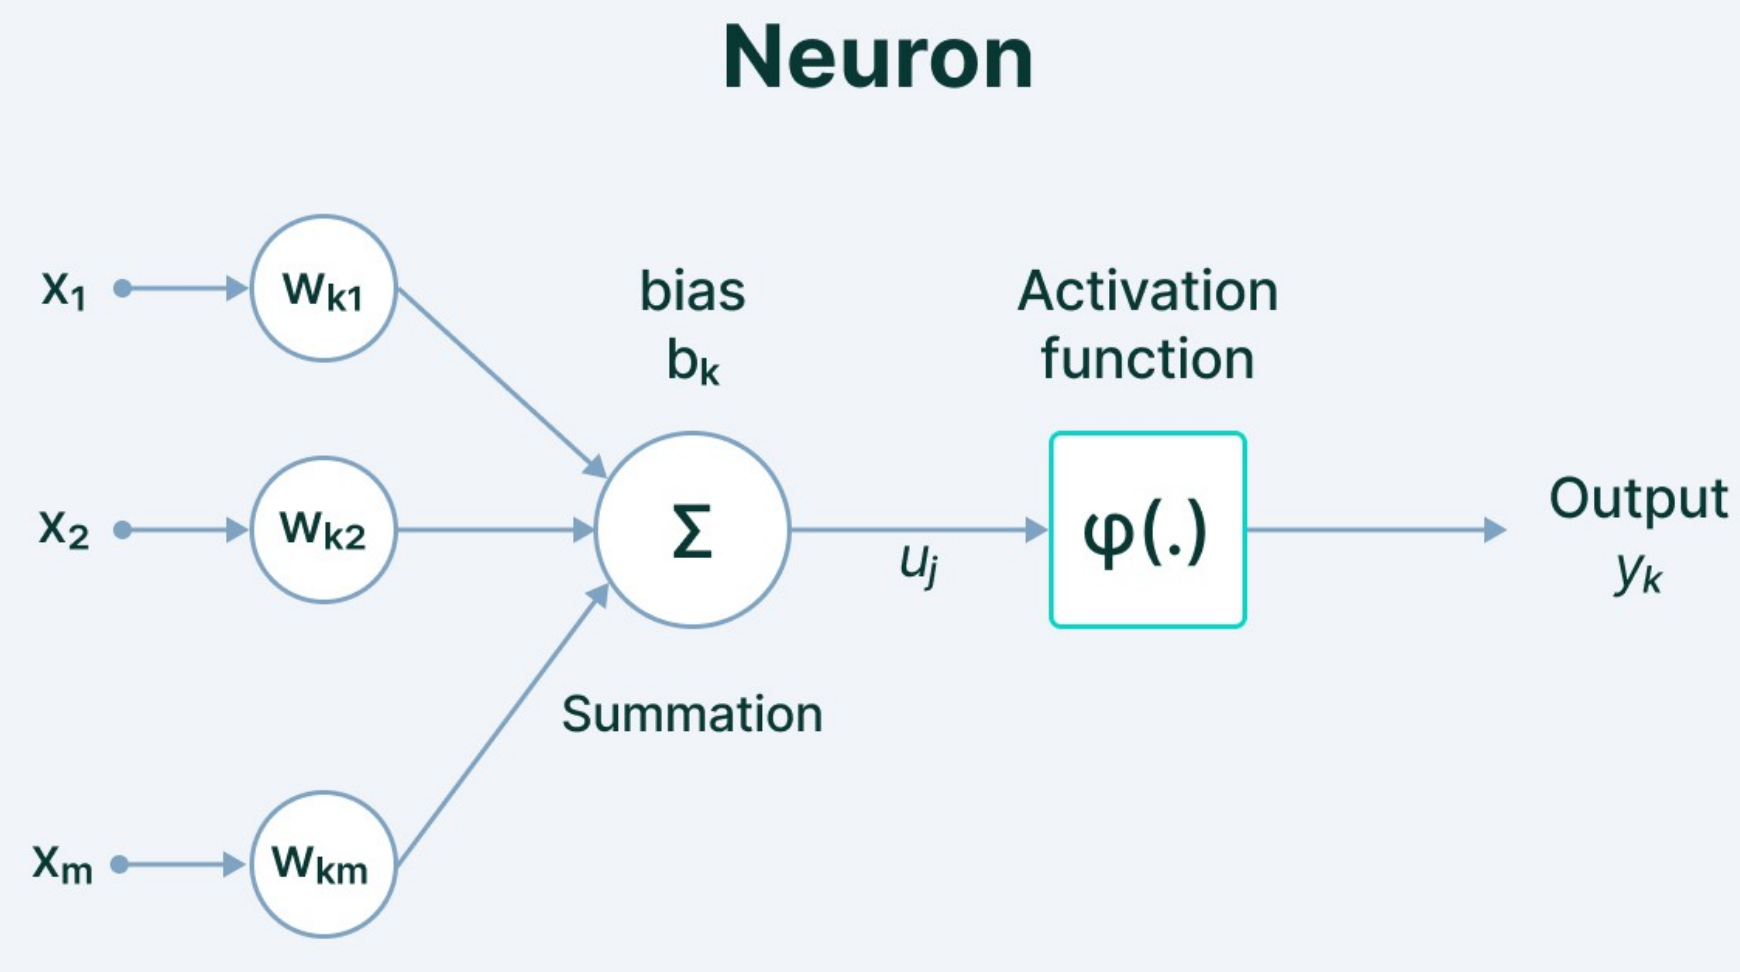
Neuron
x1
xm
wk1
wk2
wkm
bias
bk
Σ
Summation
uj
Activation
function
φ(.)
Output
yk

Neuron

- The neuronal perception of deep learning is generally motivated by two main ideas:
  - It is assumed that the human brain proves that intelligent behavior is possible, and—by reverse engineering, it is possible to build an intelligent system
  - Another perspective is that to understand the working of the human brain and the principles that underlie its intelligence is to build a mathematical model that could shed light on the fundamental scientific questions.

Deep Learning vs. Machine Learning

- Deep Learning can essentially do everything that machine learning does, but not the other way around.
- For instance, machine learning is useful when the dataset is small and well-curated, which means that the data is carefully preprocessed.
- Data preprocessing requires human intervention. It also means that when the dataset is large and complex, machine learning algorithms will fail to extract information, and it will underfit.

Deep Learning vs. Machine Learning

- Generally, machine learning is alternatively termed shallow learning because it is very effective for smaller datasets.
- Deep learning, on the other hand, is extremely powerful when the dataset is large.
- It can learn any complex patterns from the data and can draw accurate conclusions on its own. In fact, deep learning is so powerful that it can even process unstructured data - data that is not adequately arranged like text corpus, social media activity, etc.
- Furthermore, it can also generate new data samples and find anomalies that machine learning algorithms and human eyes can miss.

Deep Learning vs. Machine Learning

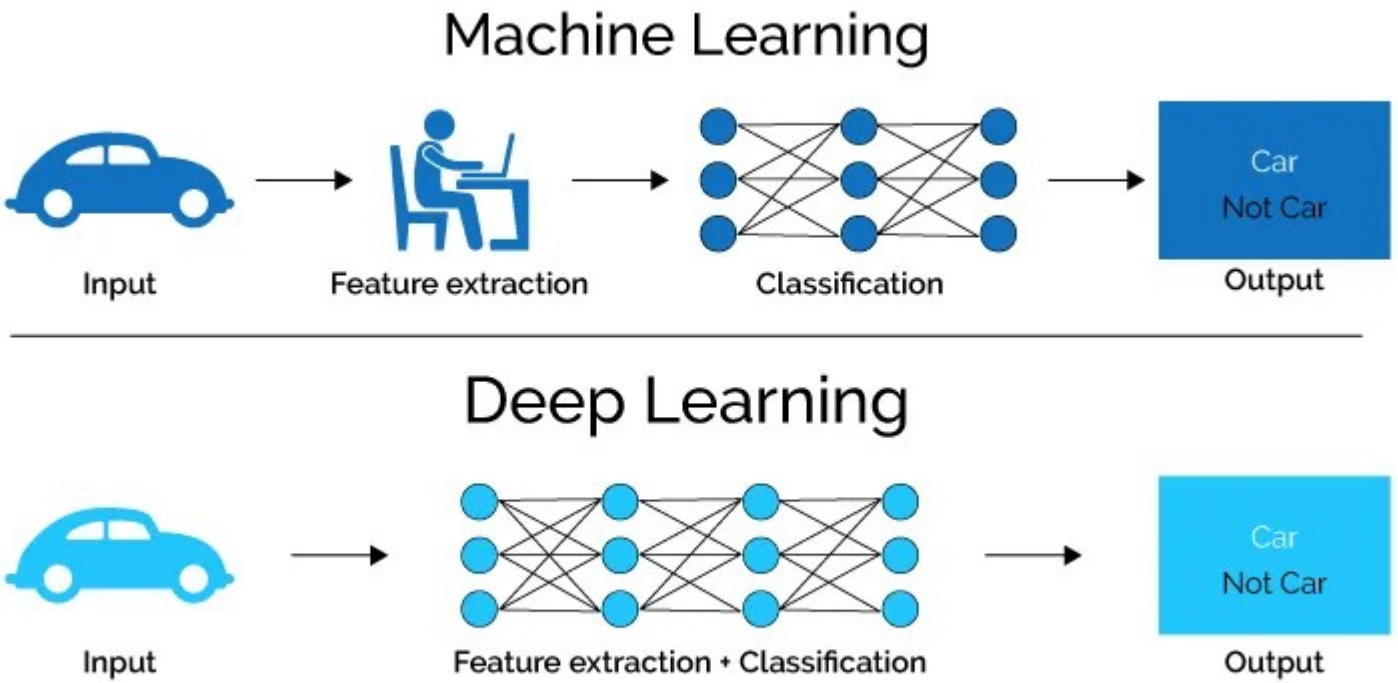

Why Deep Learning ?

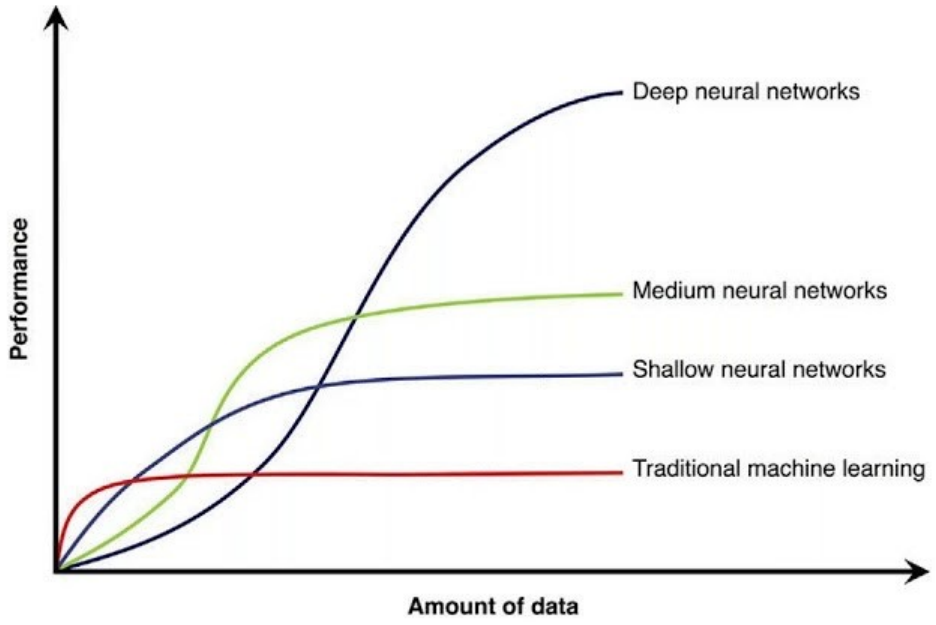

Why Now?

- *Algorithm Advancements*
- *GPU Computing*
- *Availability of Larger Training Data*

Deep Learning vs. Machine Learning

- On the downside, deep learning is computationally expensive compared to machine learning, which also means that it requires a lot of time to process.
- Deep Learning and Machine Learning are both capable of different types of learning: Supervised Learning (labeled data), Unsupervised Learning (unlabeled data), and Reinforcement Learning.
- But their usefulness is usually determined by the size and complexity of the data.

A Quick Summary

- Machine learning requires data preprocessing, which involves human intervention.
- The neural networks in deep learning are capable of extracting features; hence no human intervention is required.
- Deep Learning can process unstructured data.
- Deep Learning is usually based on representative learning i.e., finding and extracting vital information or patterns that represent the entire dataset.
- Deep learning is computationally expensive and time-consuming.

How does Deep Learning work?

- Deep Neural Networks have multiple layers of interconnected artificial neurons or nodes that are stacked together.
- Each of these nodes has a simple mathematical function - usually a linear function that performs extraction and mapping of information.
- There are three layers to a deep neural network: the input layer, hidden layers, and the output layer.

How does Deep Learning work?

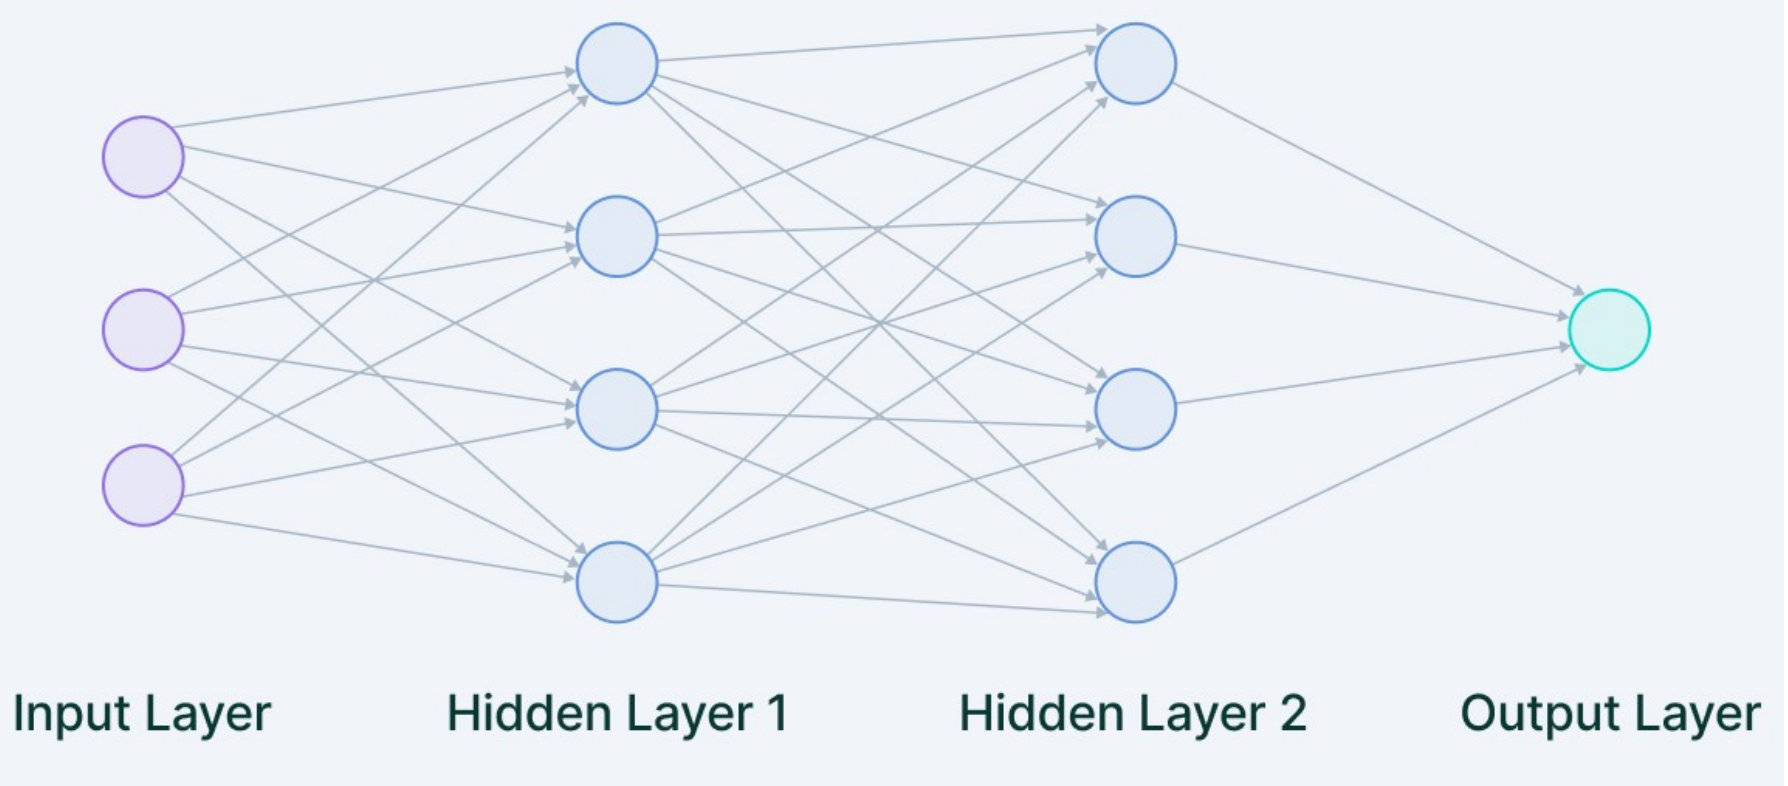

Types of Neural Network

- Artificial Neural Network
- Convolutional Neural Network
- Recurrent Neural Network
- Generative Adversarial Network

CNN

- The Convolutional Neural Networks or CNNs are primarily used for tasks related to computer vision or image processing.
- CNNs are extremely good in modeling spatial data such as 2D or 3D images and videos.
- They can extract features and patterns within an image, enabling tasks such as image classification or object detection.

CNN

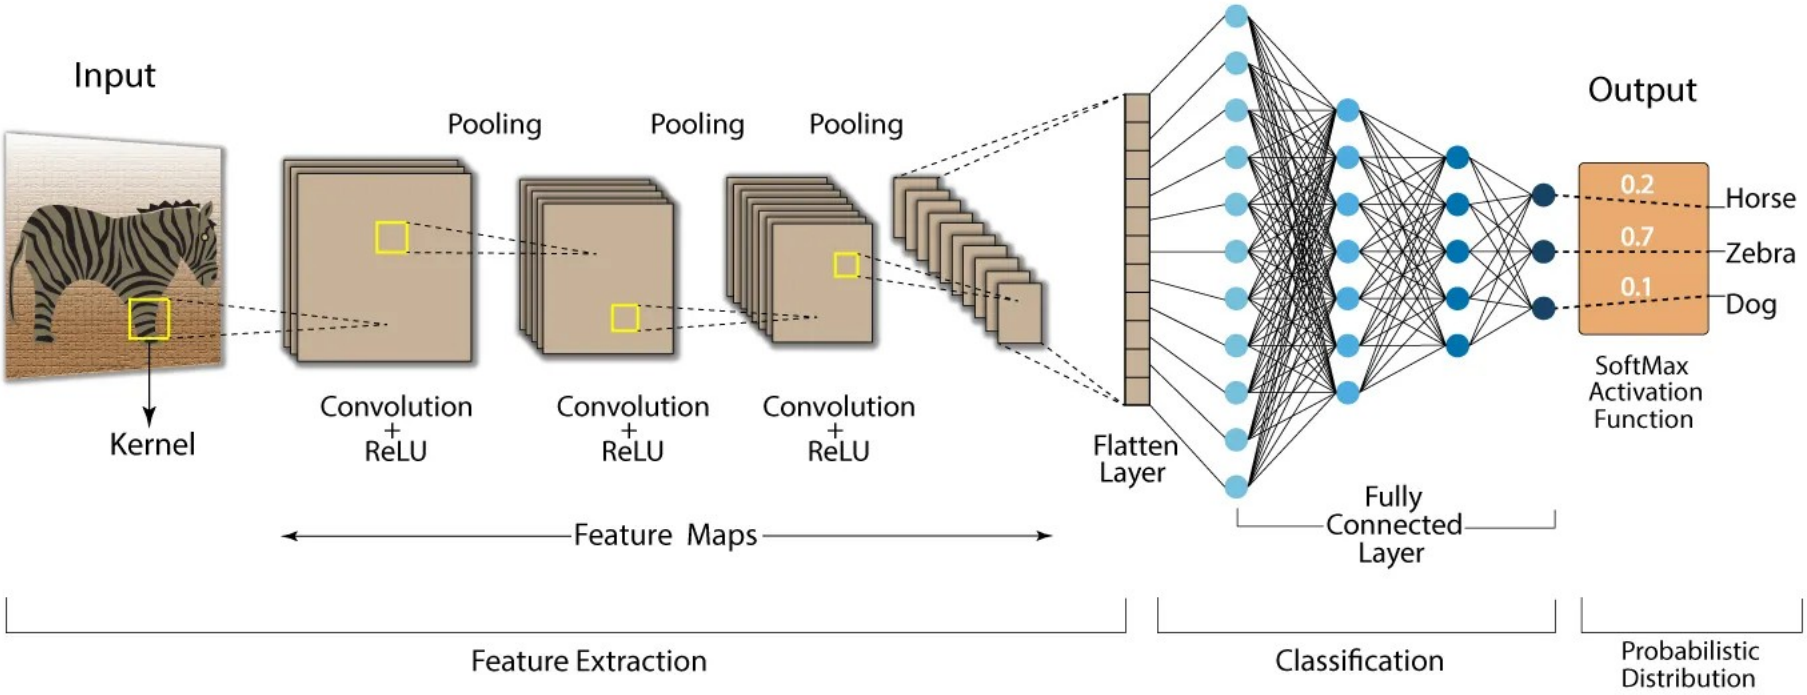
Input
Pooling
Pooling
Pooling
Output
0.2
Horse
0.7
Zebra
0.1
Dog
Kernel
Convolution + ReLU
Convolution + ReLU
Convolution + ReLU
Flatten Layer
SoftMax Activation Function
Feature Maps
Fully Connected Layer
Feature Extraction
Classification
Probabilistic Distribution

RNN

- The Recurrent Neural Networks or RNN are primarily used to model sequential data, such as text, audio, or any type of data that represents sequence or time.
- They are often used in tasks related to natural language processing (NLP).

RNN

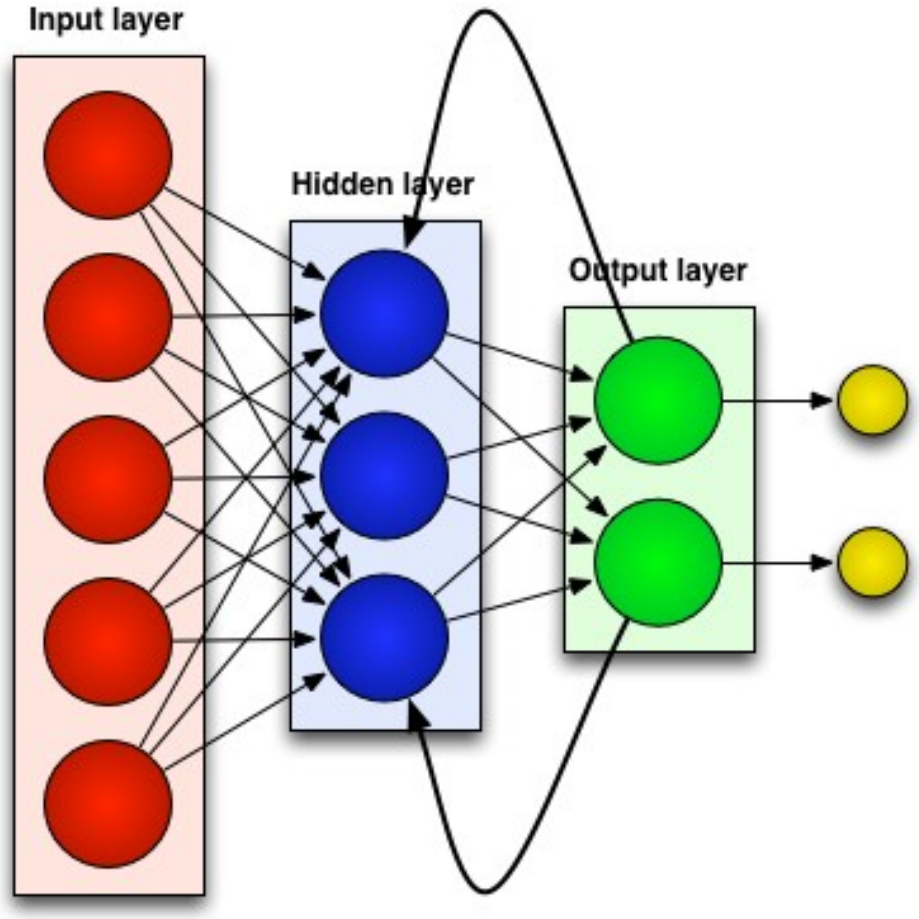

GAN

- Generative adversarial networks or GANs are frameworks that are used for the tasks related to unsupervised learning.
- This type of network essentially learns the structure of the data, and patterns in a way that it can be used to generate new examples, similar to that of the original dataset.

GAN

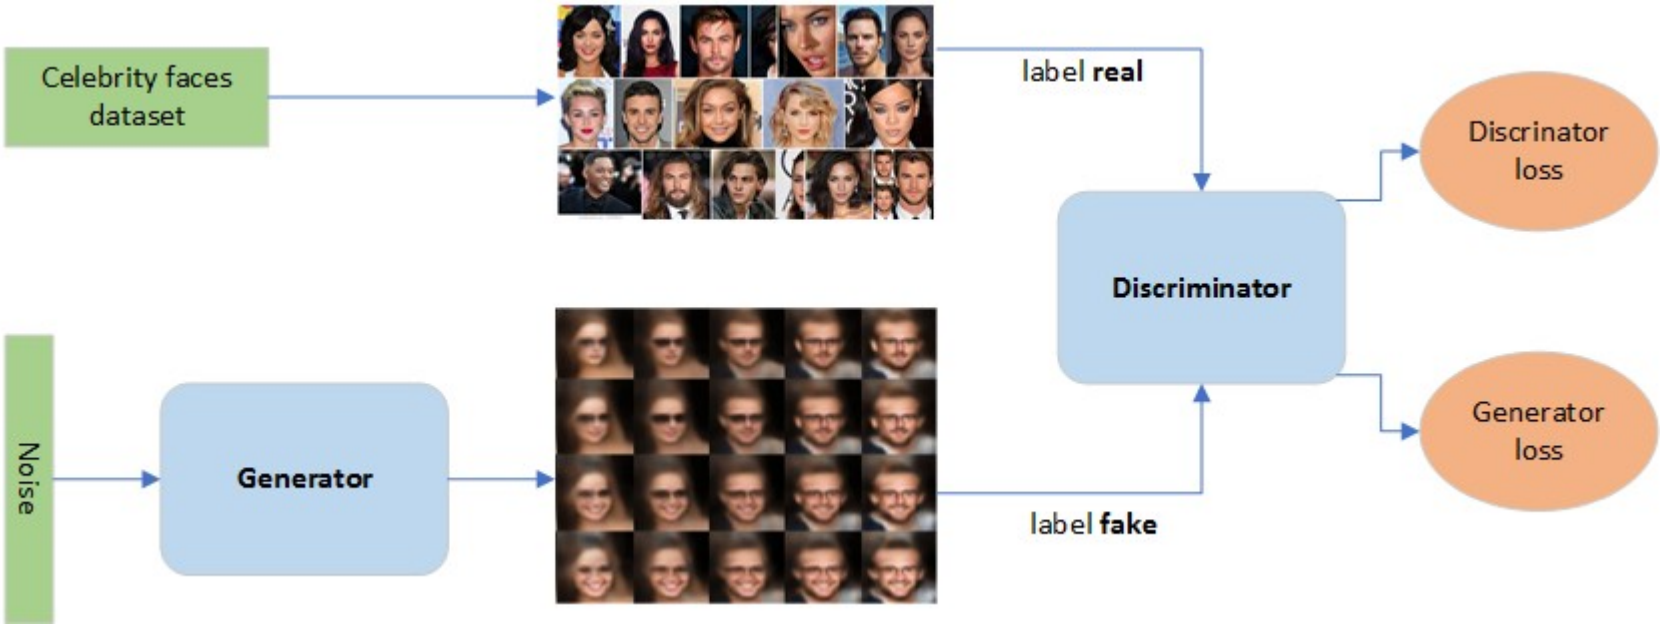
Celebrity faces dataset
label real
Discrinator loss
Discriminator
Noise
Generator
Generator loss
label fake

Transformers

- Transformers are the new class deep learning model that is used mostly for the tasks related to modeling sequential data, like that in NLP.
- It is much more powerful than RNNs and they are replacing them in every task.
- Recently, transformers are also being applied in computer vision tasks and they are proving to be quite effective than the traditional CNNs.

Deep Learning: Limitations

- Data availability
- The complexity of the model
- Lacks global generalization
- Incapable of Multitasking
- Hardware dependence

Deep Learning: Limitations

- Data availability
  - Deep learning models require a lot of data to learn the representation, structure, distribution, and pattern of the data.
  - If there isn't enough varied data available, then the model will not learn well and will lack generalization (it won't perform well on unseen data).
  - The model can only generalize well if it is trained on large amounts of data.

Deep Learning: Limitations

- The complexity of the model
  - Designing a deep learning model is often a trial and error process.
  - A simple model is most likely to underfit, i.e. not able to extract information from the training set, and a very complex model is most likely to overfit, i.e., not able to generalize well on the test dataset.
  - Deep learning models will perform well when their complexity is appropriate to the complexity of the data.

Deep Learning: Limitations

- Lacks global generalization
  - A simple neural network can have thousands to tens of thousands of parameters.
  - The idea of global generalization is that all the parameters in the model should cohesively update themselves to reduce the generalization error or test error as much as possible. However, because of the complexity of the model, it is very difficult to achieve zero generalization error on the test set.
  - Hence, the deep learning model will always lack global generalization which can at times yield wrong results.

Deep Learning: Limitations

- Incapable of Multitasking
  - Deep neural networks are incapable of multitasking.
  - These models can only perform targeted tasks, i.e., process data on which they are trained. For instance, a model trained on classifying cats and dogs will not classify men and women.
  - Furthermore, applications that require reasoning or general intelligence are completely beyond what the current generation's deep learning techniques can do, even with large sets of data.

Deep Learning: Limitations

- Hardware dependence
  - As mentioned before, deep learning models are computationally expensive.
  - These models are so complex that a normal CPU will not be able to withstand the computational complexity.
  - However, multicore high-performing graphics processing units (GPUs) and tensor processing units (TPUs) are required to effectively train these models in a shorter time.
  - Although these processors save time, they are expensive and use large amounts of energy.

Deep Learning: Applications

Thank you

@mitu_skillologies

/mITuSkillologies

@mitu_group

/company/mitu-skillologies

MITUSkillologies

Web Resources
https://mitu.co.in
http://tusharkute.com

contact@mitu.co.in

tushar@tusharkute.com