ARTIFICIAL NEURAL NETWORK

INTRODUCTION TO NEURAL NETWORKS

- The neural network, or artificial neural network, was inspired by and modeled after the biological neural network.
- These networks, like the human brain, learn to perform specific tasks without being explicitly programmed.
- A neural network is composed of a series of neurons that are connected to form a type of network, hence the name neural network.
- A neuron, or an artificial neuron, is the fundamental unit of a neural network.

INTRODUCTION TO PERCEPTRON

- In the Artificial Neural Network(ANN), the perceptron is a convenient model of a biological neuron, it was the early algorithm of binary classifiers in supervised machine learning.
- The purpose behind the designing of the perceptron model was to incorporate visual inputs, organizing subjects or captions into one of two classes and dividing classes through a line.

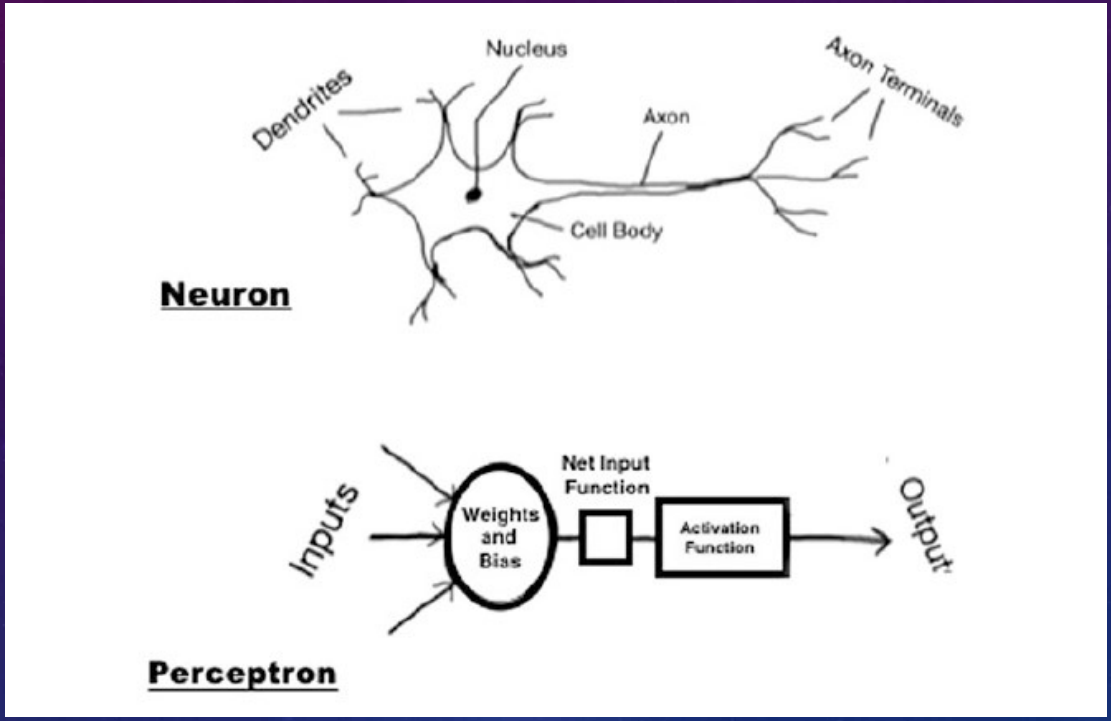

A perceptron is nothing but an artificial neuron. In deep learning, the terms perceptron and neuron are used interchangeably.

COMPARISON – BIOLOGICAL VS ARTIFICIAL

Biological Neuron	Artificial Neuron
It receives information in the form of electrical signals.	It receives information in the form of numerical values.
Literally speaking, the brain consists of about 86 billion biological neurons.	A neural network can consist of a maximum of about 1,000 artificial neurons.
The general composition is a cell body, dendrites, and axon.	The general composition is the weights and bias, the net input function, and the activation function.

DIGGING DEEPER INTO NEURAL NETWORKS

- Deep learning goes a step further in machine learning. It allows the machine to begin thinking on its own in order to make decisions and carry out certain tasks.
- Neural networks are used to develop and train deep learning models.
- For example, consider a very simple neural network, which consists of an input layer, an output layer, and one layer of neurons, known as the hidden layer.

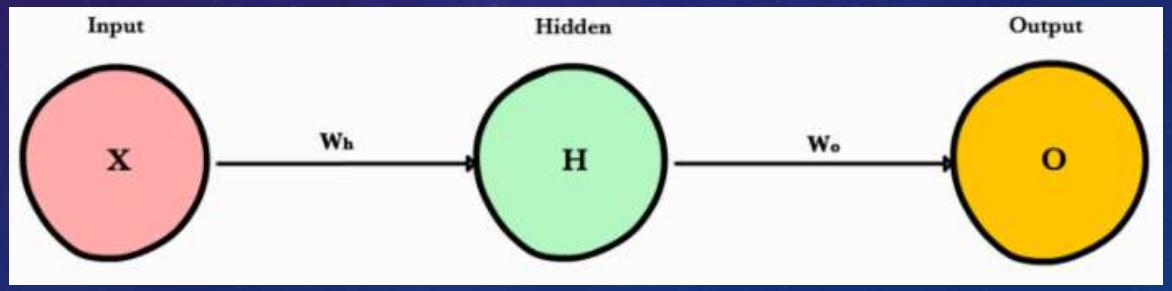

DIGGING DEEPER INTO NEURAL NETWORKS

The basic function of these three sections is as follows:

1. The input layer, as the name implies, is made of the input signals that will be further transmitted into the neural network.
2. The hidden layer is where all the important computations occur. The input that is fed to it is taken, calculations are performed on it, and then this input is sent to the next layer, which is the output layer. The hidden layer can have any number of neurons within it. There can also be more than one hidden layer, depending on our requirements and arrangement.
3. The output layer, as the name suggests, contains the output signals. These are nothing but the final results of all the calculations performed by the hidden layer/s.

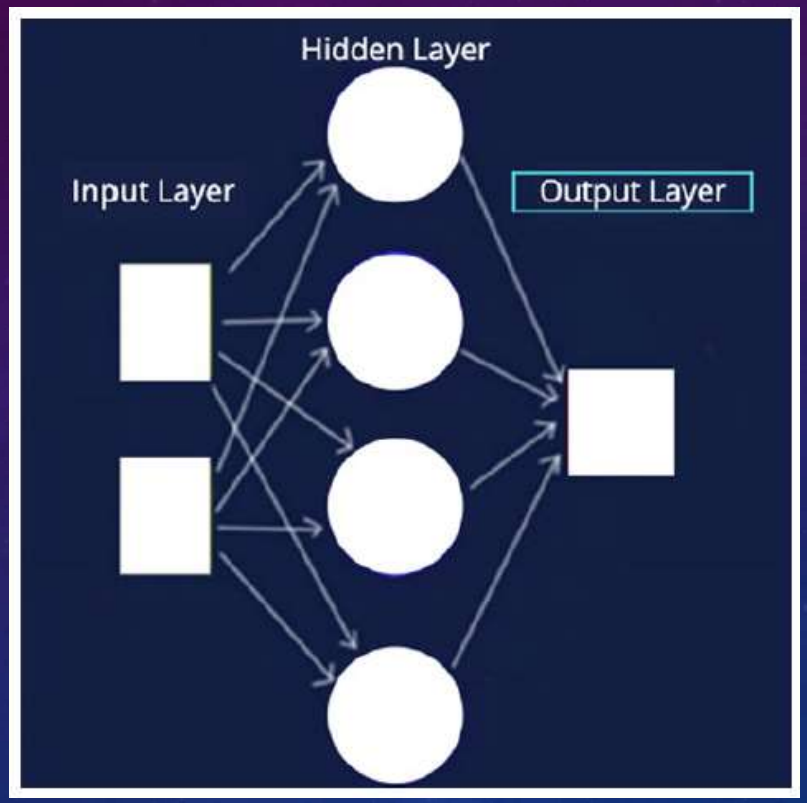
Hidden Layer
Input Layer
Output Layer

THE PROCESS

1. The numerical input signals are passed into the neural network's hidden layers.
2. The net input function is calculated with the weights and the bias that are generated during the training.
3. The activation function is applied to the net input function.
4. The result is then produced as the output of the neural network.

NEURAL NETWORK ACTIVATION FUNCTIONS

- In artificial neural networks, the activation function of a node defines the output of that node given an input or set of inputs.
- So, what does an artificial neuron do? Simply put, it calculates a "weighted sum" of its input, adds a bias and then decides whether it should be "fired" or not.

$$Y = \sum (weight * input) + bias$$

- Now, the value of Y can be anything ranging from -inf to +inf. The neuron really doesn't know the bounds of the value. So how do we decide whether the neuron should fire or not (why this firing pattern? Because we learnt it from biology that's the way brain works, and brain is a working testimony of an awesome and intelligent system).

NEURAL NETWORK ACTIVATION FUNCTIONS

- We decided to add "activation functions" for this purpose. To check the Y value produced by a neuron and decide whether outside connections should consider this neuron as "fired" or not. Or rather let's say — "activated" or not.
- If the value of Y is above a certain value, declare it activated. If it's less than the threshold, then say it's not.

Activation function A = "activated" if Y > threshold else not

Alternatively, A = 1 if y> threshold, 0 otherwise

TYPES OF NEURAL NETWORKS

FEED-FORWARD NEURAL NETWORK

This type of neural network essentially consists of an input layer, multiple hidden layers and an output layer.

There is no loop and information only flows forward.

Feed-forward neural networks are generally suited for supervised learning with numerical data

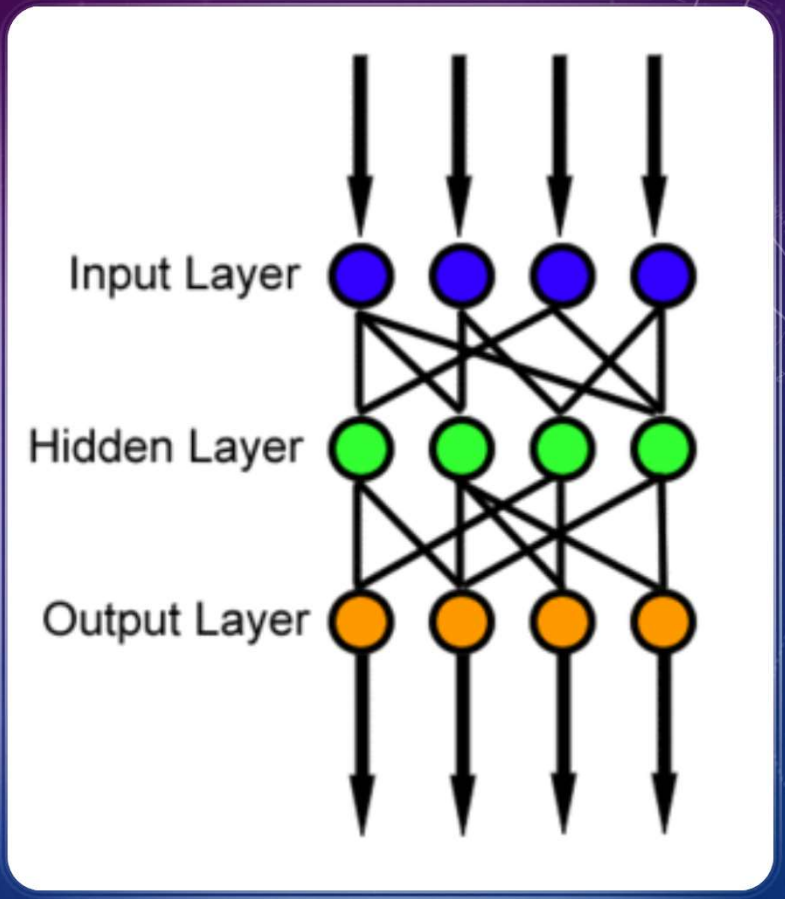

CONVOLUTIONAL NEURAL NETWORKS (CNN)

In deep learning, a convolutional neural network (CNN, or ConvNet) is a class of artificial neural network, most applied to analyze visual imagery.

Convolutional networks were inspired by biological processes in that the connectivity pattern between neurons resembles the organization of the animal visual cortex. Individual cortical neurons respond to stimuli only in a restricted region of the visual field known as the receptive field. The receptive fields of different neurons partially overlap such that they cover the entire visual field.

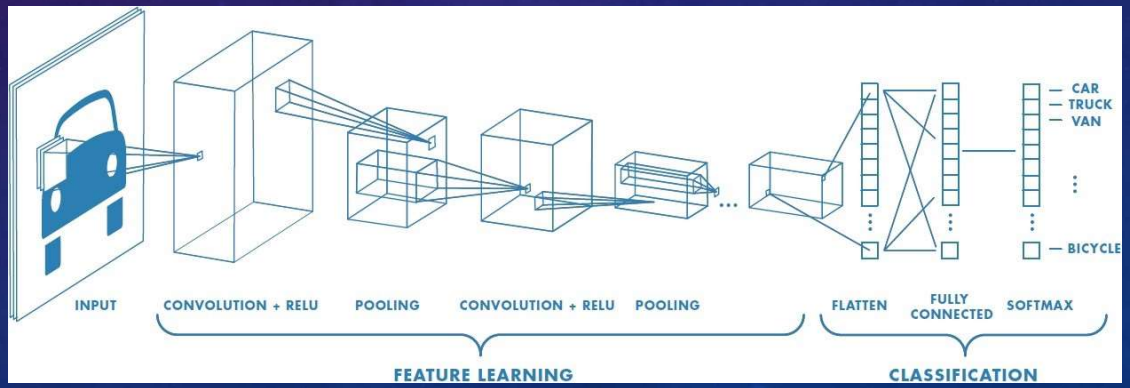

RECURRENT NEURAL NETWORKS

What CNN means for images; Recurrent Neural Networks are meant for text. RNNs can help us learn the sequential structure of text where each word is dependent on the previous work, or a word in the previous sentence.

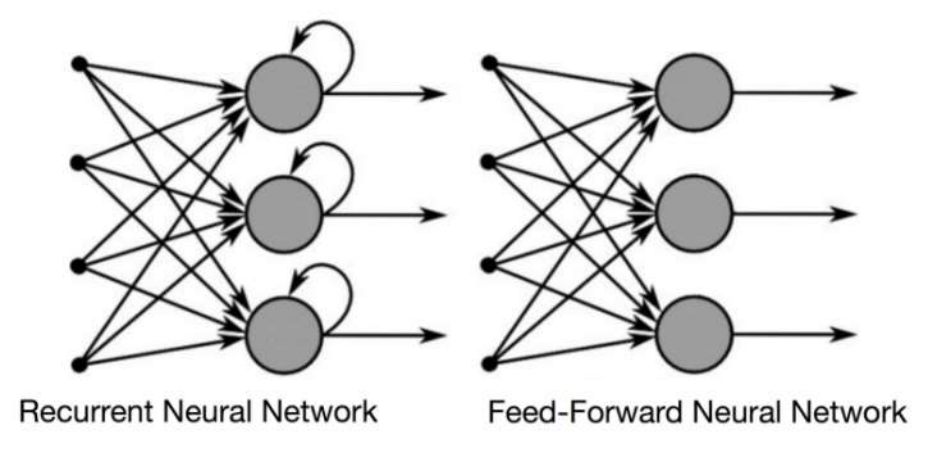

TRANSFORMERS

A transformer is a deep learning model that adopts the mechanism of attention, differentially weighing the significance of each part of the input data.

It is used primarily in the field of natural language processing (NLP) and in computer vision (CV).

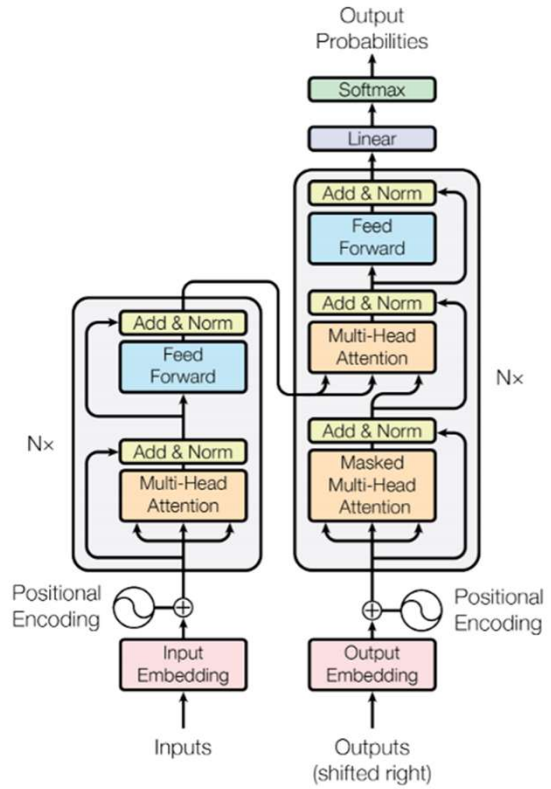

Figure 1: The Transformer - model architecture.

GENERATIVE ADVERSARIAL NETWORKS

A generative adversarial network (GAN) is a class of machine learning frameworks designed by Ian Goodfellow and his colleagues in 2014. Two neural networks contest with each other in a game (in the form of a zero-sum game, where one agent's gain is another agent's loss).

- GANs typically employ two dueling neural networks to train a computer to learn the nature of a dataset well enough to generate convincing fakes. One of these neural networks generates fakes (the generator), and the other tries to classify which images are fake (the discriminator). These networks improve over time by competing against each other.
- Perhaps it's best to imagine the generator as a robber and the discriminator as a police officer. The more the robber steals, the better he gets at stealing things. At the same time, the police officer also gets better at catching the thief.

Source: All of them are fake

THE END