



# ARTIFICIAL NEURAL NETWORKS & DEEP LEARNING


AI for Engineers



# WHAT ARE NEURONS?

**Neurons** (also called neurones or nerve cells) are the fundamental units of the brain and nervous system, the cells responsible for receiving sensory input from the external world, for sending motor commands to our muscles, and for transforming and relaying the electrical signals at every step in between.


An artificial neuron is a mathematical function conceived as a model of biological neurons, a neural network.



# WHAT IS ARTIFICIAL NEURAL NETWORKS (ANN)?

An Artificial Neural Network is an information processing model that is inspired by the way biological nervous systems, such as the brain, process information.

In simpler terms it is a mathematical model of the brain which is used to process nonlinear relationships between inputs and outputs in parallel like a human brain does every second.




Artificial Neural Networks are used for a variety of tasks, a popular use is for classification. You can collect datasets of images for example of different breeds of dogs and then train a neural network on the images, then if you supply a new image of a dog it will give a statistical score on how closely the new image matches the model and then will output what breed of dog the image is.

Neural Networks are also used in Self Driving cars, Character Recognition, Image Compression, Stock Market Prediction, and lots of other interesting applications.

## *WHAT ARE ARTIFICIAL NEURAL NETWORKS USED FOR?*



# HOW DOES A NEURAL NETWORK LEARN?



Information flows through a neural network in two different ways. When the model is learning (being trained) or operating normally (after being trained either being used or tested), patterns of information from the dataset are being fed into the network via the input neurons, which trigger the layers of hidden neurons, and these in turn arrive at the output neurons. This is called a feedforward network.

Not all neurons “fire” all the time. Each neuron receives inputs from the neurons to its left, and the inputs are multiplied by the weights of the connections they travel along. Every neuron adds up all the inputs it receives in this way and (this is the simplest neural network) if the sum is more than a certain threshold value, the neuron “fires” and triggers the neurons it’s connected to (the neurons on its right).



# *HOW DOES A NEURAL NETWORK LEARN? (CONTD.)*

For an artificial neural network to learn, it has to learn what it has done wrong and is doing right, this is called feedback. Feedback is how we learn what is wrong and right and this is also what an artificial neural network needs for it to learn.

This is where you start to see similarities to the human brain. If you are learning to play a game like tennis you learn that if you hit the ball too hard it will go out of the court and you will lose the point, or if you don't hit the ball hard enough it won't go over the net but if you hit it perfectly it will go onto the other side in the court and if could win a point, this is a classic example of feedback where you lose the point or potentially gain a point.

This is how we learn what we are doing correct or wrong and this is what a neural network needs to learn.






# HOW DOES A NEURAL NETWORK LEARN? (CONTD.)

Neural networks learn things in exactly the same way as the brain, typically by a feedback process called back-propagation (this is sometimes shortened to “backprop”). This is where you compare the output of the network with the output it was meant to produce, and using the difference between the outputs to modify the weights of the connections between the neurons in the network, working from the output units through the hidden neurons to the input neurons going backward.

Over time, back-propagation causes the network to learn by making the gap between the output and the intended output smaller to the point where the two exactly match, so the neural network learns the correct output.



# WHAT IS DEEP LEARNING?



Deep learning is a subset of machine learning where artificial neural networks, algorithms inspired by the human brain, learn from large amounts of data. Similarly, to how we learn from experience, the deep learning algorithm would perform a task repeatedly, each time tweaking it a little to improve the outcome.

We refer to ‘deep learning’ because the neural networks have various (deep) layers that enable learning. Just about any problem that requires “thought” to figure out is a problem deep learning can learn to solve.

Deep learning allows machines to solve complex problems even when using a data set that is very diverse, unstructured and inter-connected. The more deep learning algorithms learn, the better they perform.





# REAL LIFE USES OF DEEP LEARNING

1. Virtual assistants- Whether it's Alexa or Siri or Cortana, the virtual assistants of online service providers use deep learning to help understand your speech and the language humans use when they interact with them.
2. Translations- In a similar way, deep learning algorithms can automatically translate between languages. This can be powerful for travelers, business people and those in government.
3. Artificial Vision- for driverless delivery trucks, drones and autonomous cars. The way an autonomous vehicle understands the realities of the road and how to respond to them whether it's a stop sign, a ball in the street or another vehicle is through deep learning algorithms. The more data the algorithms receive, the better they are able to act human-like in their information processing—knowing a stop sign covered with snow is still a stop sign.
4. Chatbots and service bots -Chatbots and service bots that provide customer service for a lot of companies are able to respond in an intelligent and helpful way to an increasing amount of auditory and text questions thanks to deep learning.



# REAL LIFE USES OF DEEP LEARNING (CONTD)

5. Image colorization -Transforming black-and-white images into colour was formerly a task done meticulously by human hand. Today, deep learning algorithms are able to use the context and objects in the images to colour them to basically recreate the black-and-white image in colour. The results are impressive and accurate.

6. Facial recognition- Deep learning is being used for facial recognition not only for security purposes but for tagging people on Facebook posts and we might be able to [pay for items in a store just by using our faces](#) in the near future. The challenges for deep-learning algorithms for facial recognition is knowing it's the same person even when they have changed hairstyles, grown or shaved off a beard or if the image taken is poor due to bad lighting or an obstruction.

7. Medicine and pharmaceuticals - From disease and tumour diagnoses to personalised medicines created specifically for an individual's genome, deep learning in the medical field has the attention of many of the largest pharmaceutical and medical companies.

8. Personalised shopping and entertainment -Ever wonder how Netflix comes up with suggestions for what you should watch next? Or where Amazon comes up with ideas for what you should buy next and those suggestions are exactly what you need but just never knew it before? Yep, it's deep-learning algorithms at work.

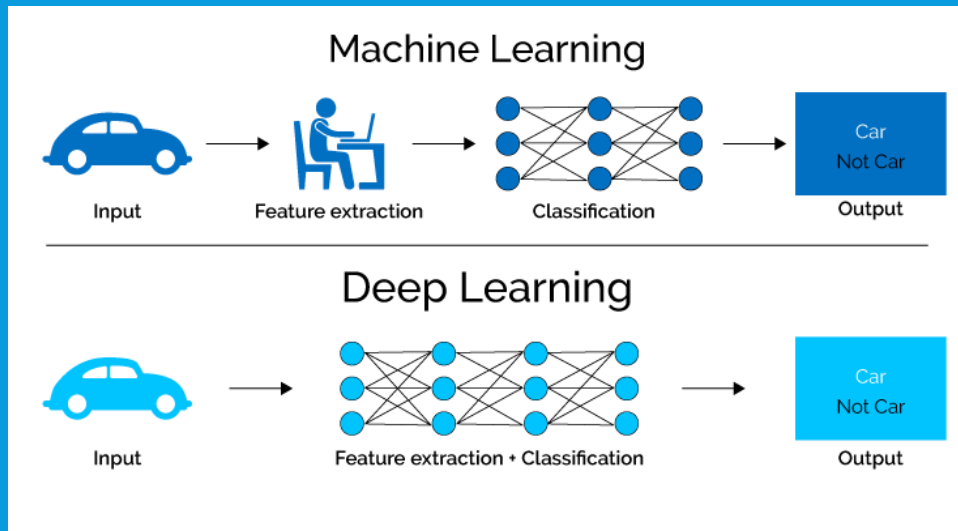


# DIFFERENCE BETWEEN ANN AND DNN

- ANN in its simplest form has only three layers.
- In real use an ANN that is made up of more than three layers – i.e. an input layer, an output layer and multiple hidden layers – is called a ‘deep neural network’, and this is what underpins deep learning. A deep learning system is self-teaching, learning as it goes by filtering information through multiple hidden layers, in a similar way to humans.
- The two are closely connected in that one relies on the other to function. Without neural networks, there would be no deep learning.




# HOW IS DEEP LEARNING DIFFERENT FROM MACHINE LEARNING?




- Classical, or "non-deep", machine learning is dependent on human intervention to learn, requiring labeled datasets to understand the differences between data inputs. For example, if I were to show you a series of images of different types of fast food, I would label each picture with a fast food type, such as "pizza," "burger," or "taco." The machine learning model would train and learn based on the labelled data fed into it, which is also known as supervised learning.
- "Deep" machine learning can leverage labeled datasets to inform its algorithm, but it doesn't necessarily require a labeled dataset; instead it can also leverage unsupervised learning to train itself. While supervised learning leverages labeled data, unsupervised learning uses unstructured, or unlabeled, data.
- By observing patterns in the data, a machine learning model can cluster and classify inputs. Taking the same example from earlier, we could group pictures of pizzas, burgers, and tacos into their respective categories based on the similarities identified in the images. With that said, a deep learning model would require more data points to improve its accuracy, whereas a machine learning model relies on less data given the underlying data structure. Deep learning is primarily leveraged for more complex use cases, like virtual assistants or fraud detection.

# TYPES OF ARTIFICIAL NEURAL NETWORKS

- The perceptron is the oldest neural network, created by Frank Rosenblatt in 1958. It has a single neuron and is the simplest form of a neural network.
- Feedforward neural networks, or multi-layer perceptrons (MLPs). The feedforward neural network was the first and simplest type of artificial neural network devised. They comprise of an input layer, a hidden layer or layers, and an output layer.




Perceptron Model (Minsky-Papert in 1969)



An example of a Feed-forward Neural Network with one hidden layer ( with 3 neurons )

# CONVOLUTIONAL NEURAL NETWORK



A Convolutional neural network (CNN) is a neural network that has one or more convolutional layers and are used mainly for image processing, classification, segmentation and also for other auto correlated data.

A convolution is essentially sliding a filter over the input.. Rather than looking at an entire image at once to find certain features it can be more effective to look at smaller portions of the image. Each convolutional layer contains a series of filters known as convolutional kernels. The filter is a matrix of integers that are used on a subset of the input pixel values, the same size as the kernel. Each pixel is multiplied by the corresponding value in the kernel, then the result is summed up for a single value for simplicity representing a grid cell, like a pixel, in the output channel/feature map.



# UNIVERSAL APPROXIMATION THEOREM

- The **Universal Approximation Theorem** states that a neural network with one hidden layer can approximate any continuous function for inputs within a specific range.
- Universal approximation theorems imply that neural networks can *represent* a wide variety of interesting functions when given appropriate weights. On the other hand, they typically do not provide a construction for the weights, but merely state that such a construction is possible.

# GENERATIVE ADVERSARIAL NETWORKS (GANS)

- Generative Adversarial Networks, or GANs for short, are an approach to generative modeling using deep learning methods, such as convolutional neural networks.
- 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.