

Artificial Intelligence (AI)

Machine Learning (ML)

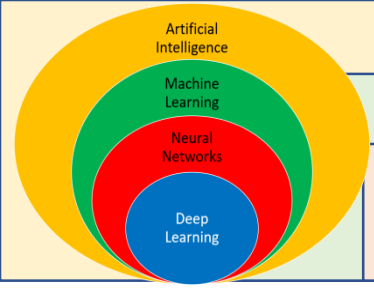
Neural Networks (NNs)

Deep Learning (DL)

Advanced Artificial Intelligence

Dr. Rastgoo





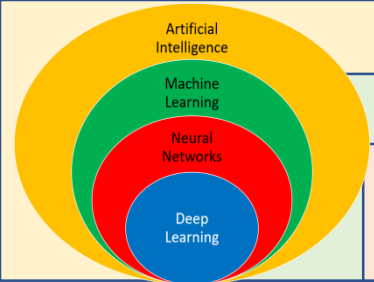
Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Part 1: Convolutional Neural Network (CNN)



Artificial Intelligence (AI)

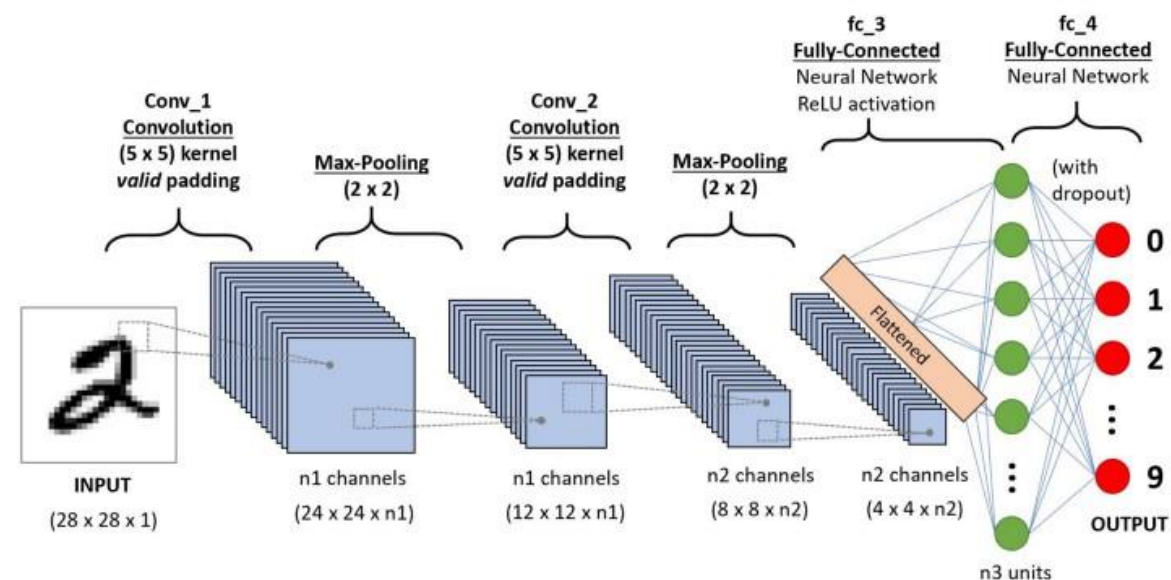
Machine Learning (ML)

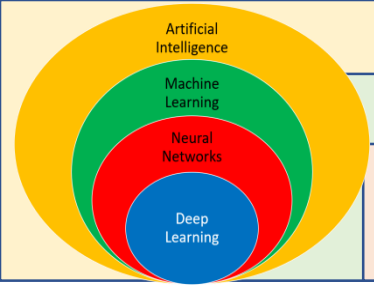
Neural Networks (NNs)

Deep Learning (DL)

Convolutional Neural Network (CNN)

- A convolutional neural network (CNN) is a type of artificial neural network used primarily for image recognition and processing, due to its ability to **recognize patterns** in images.
- A Convolutional Neural Network (ConvNet/CNN) is a **Deep Learning algorithm** which can take in an input image, assign importance (**learnable weights and biases**) to various aspects/objects in the image and be able to differentiate one from the other.





Artificial Intelligence (AI)

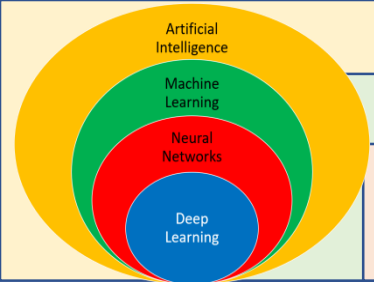
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Convolutional Neural Network (ConvNets/CNN)

- The pre-processing required in a ConvNet is much **lower** as compared to other classification algorithms.
- While in primitive methods filters are hand-engineered, with enough training, ConvNets have the ability to learn these filters/characteristics.



Artificial Intelligence (AI)

Machine Learning (ML)


Neural Networks (NNs)

Deep Learning (DL)

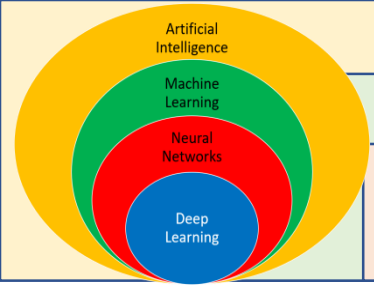
Why ConvNets over Feed-Forward Neural Nets?

- An image is nothing but a **matrix of pixel values**.
- So why not just flatten the image (e.g. 3x3 image matrix into a 9x1 vector) and feed it to a Multi-Level Perceptron for classification purposes?

1	1	0
4	2	1
0	2	1



1
1
0
4
2
1
0
2
1



Artificial Intelligence (AI)

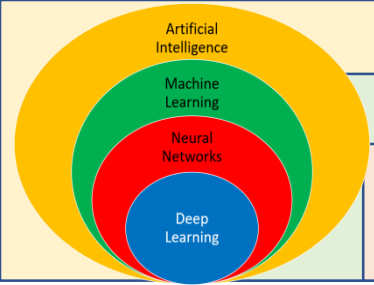
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Why ConvNets over Feed-Forward Neural Nets?

- A ConvNet is able to successfully capture the **Spatial** and **Temporal** dependencies in an image through the application of **relevant filters**.
- The architecture performs a better fitting to the image dataset due to the reduction in the number of parameters involved and **reusability** of weights.
- In other words, the network can be **trained** to understand the sophistication (complexity) of the image better.



Artificial Intelligence (AI)


Machine Learning (ML)

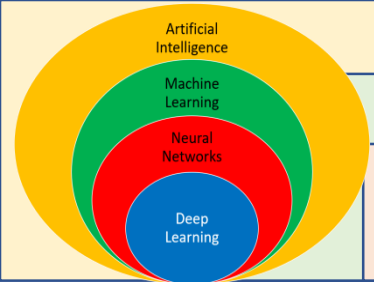
Neural Networks (NNs)

Deep Learning (DL)

Input image

- We have an RGB image which has been separated by its three color planes – Red, Green, and Blue.
- There are a number of such color spaces in which images exist – Grayscale, RGB, HSV, CMYK, etc.





Artificial Intelligence (AI)

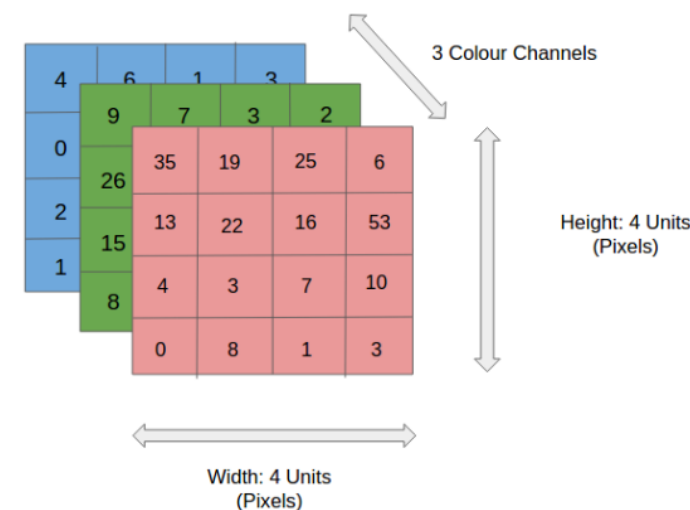
Machine Learning (ML)

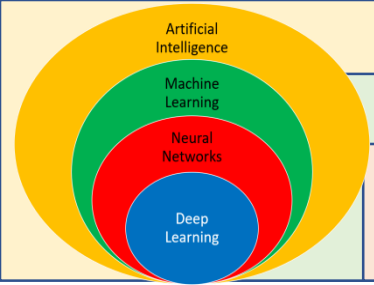
Neural Networks (NNs)

Deep Learning (DL)

Input image

- You can imagine how **computationally intensive** things would get once the images reach dimensions, say 8K (7680×4320).
- The role of the ConvNet is to **reduce** the images into a form which is easier to process, without losing features which are critical for getting a good prediction.
- This is important when we are to design an architecture which is not only good at **learning features** but also is **scalable** to massive datasets.





Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel


1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image

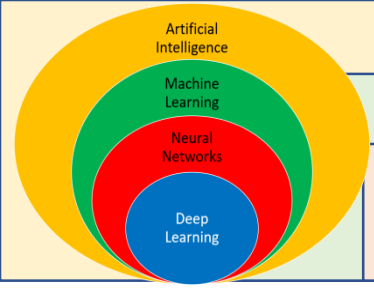
4		

Convolved
Feature

Convolving a 5x5x1 image with a 3x3x1 kernel to get a 3x3x1 convolved feature



Convolution Operation with Stride Length = 2




Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel



0	0	0	0	0	0	...
0	156	155	156	158	158	...
0	153	154	157	159	159	...
0	149	151	155	158	159	...
0	146	146	149	153	158	...
0	145	143	143	148	158	...
...

Input Channel #1 (Red)

0	0	0	0	0	0	...
0	167	166	167	169	169	...
0	164	165	168	170	170	...
0	160	162	166	169	170	...
0	156	156	159	163	168	...
0	155	153	153	158	168	...
...

Input Channel #2 (Green)

0	0	0	0	0	0	...
0	163	162	163	165	165	...
0	160	161	164	166	166	...
0	156	158	162	165	166	...
0	155	155	158	162	167	...
0	154	152	152	157	167	...
...

Input Channel #3 (Blue)

-1	-1	1
0	1	-1
0	1	1

Kernel Channel #1

↓
308

1	0	0
1	-1	-1
1	0	-1

Kernel Channel #2

↓
-498

0	1	1
0	1	0
1	-1	1

Kernel Channel #3

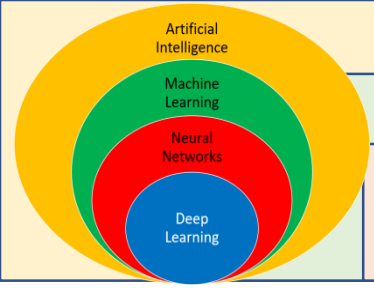
↓
164

+

+

+ 1 = -25
↑
Bias = 1

-25				...
				...
				...
				...
...



Artificial Intelligence (AI)

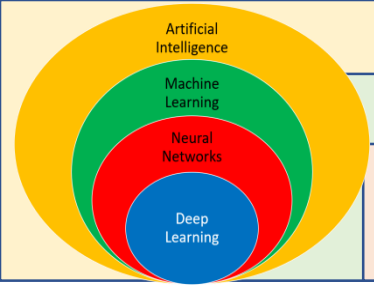
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel

- In the case of images with **multiple channels** (e.g. RGB), the Kernel has the same depth as that of the input image.
- Matrix Multiplication is performed between K_n and I_n stack ($[K_1, I_1]; [K_2, I_2]; [K_3, I_3]$) and all the results are **summed** with the bias to give us a squashed one-depth channel Convolved Feature Output.



Artificial Intelligence (AI)


Machine Learning (ML)

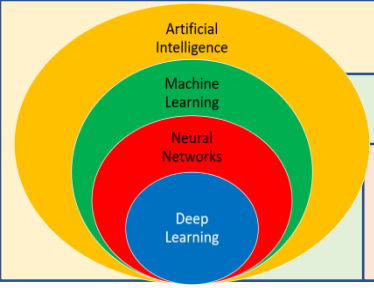
Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel

- The objective of the Convolution Operation is to extract the **high-level features** such as edges, from the input image.
- ConvNets need not be limited to only one Convolutional Layer.





Artificial Intelligence (AI)


Machine Learning (ML)

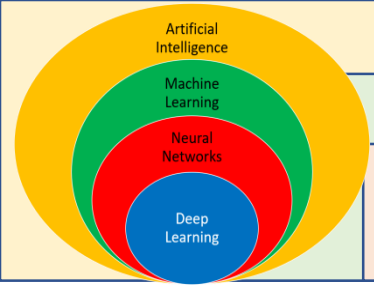
Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel

- Conventionally, the first ConvLayer is responsible for capturing the Low-Level features such as edges, color, gradient orientation, etc.
- With added layers, the architecture adapts to the **High-Level features** as well, giving us a network which has the wholesome understanding of images in the dataset, similar to how we would.





Artificial Intelligence (AI)


Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel

- There are two types of results to the operation: one in which the convolved feature is **reduced in dimensionality** as compared to the input, and the other in which the dimensionality is either **increased** or **remains the same**.
- This is done by applying **Valid Padding** in case of the former, or **Same Padding** in the case of the latter.



Artificial Intelligence (AI)


Machine Learning (ML)


Neural Networks (NNs)

Deep Learning (DL)

Convolution Layer: The Kernel

- **Same Padding:** When we augment the 5x5x1 image into a 6x6x1 image and then apply the 3x3x1 kernel over it, we find that the convolved matrix turns out to be of dimensions 5x5x1.





Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

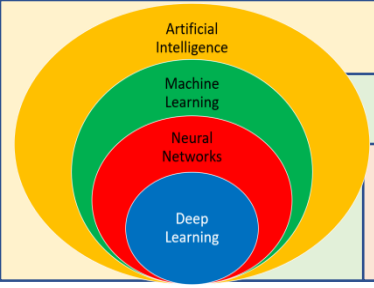
Deep Learning (DL)

Convolution Layer: The Kernel

- **Valid Padding:** On the other hand, if we perform the same operation without padding, we are presented with a matrix which has dimensions of the Kernel (3x3x1) itself.

3_0	3_1	2_2	1	0
0_2	0_2	1_0	3	1
3_0	1_1	2_2	2	3
2	0	0	2	2
2	0	0	0	1

12	12	17
10	17	19
9	6	14



Artificial Intelligence (AI)

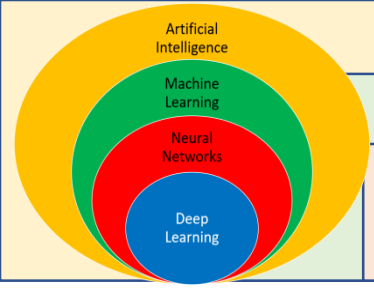
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Pooling Layer

- Similar to the Convolutional Layer, the Pooling layer is responsible for **reducing** the spatial size of the Convolved Feature.
- This is to decrease the **computational power** required to process the data through dimensionality reduction.



Artificial Intelligence (AI)

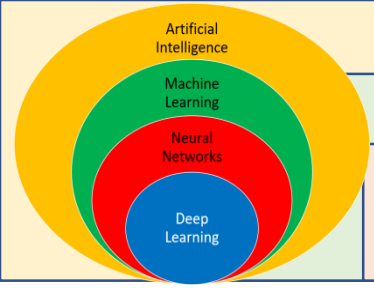
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Pooling Layer

- There are two types of Pooling: **Max Pooling** and **Average Pooling**.
- **Max Pooling** returns the maximum value from the portion of the image covered by the Kernel.
- On the other hand, **Average Pooling** returns the average of all the values from the portion of the image covered by the Kernel.



Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Pooling Layer

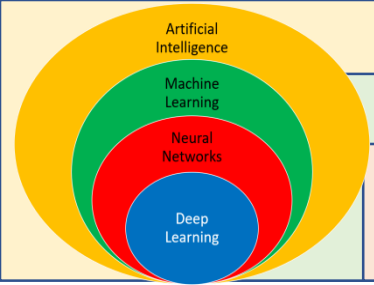
- **Max Pooling** also performs as a **Noise Suppressant**.
- It discards the noisy activations altogether and also performs **de-noising** along with dimensionality reduction.

Max Pooling

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45



Artificial Intelligence (AI)

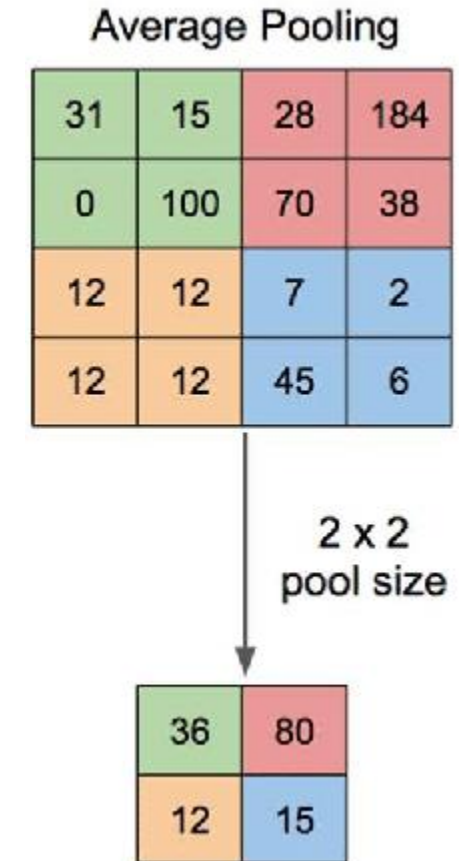
Machine Learning (ML)

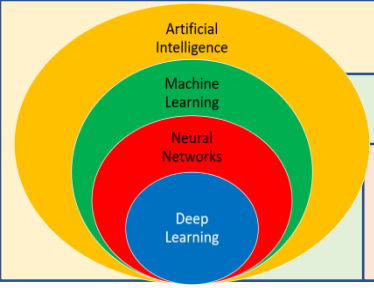
Neural Networks (NNs)

Deep Learning (DL)

Pooling Layer

- On the other hand, **Average Pooling** simply performs dimensionality reduction as a noise suppressing mechanism.
- Hence, we can say that **Max Pooling performs a lot better than Average Pooling**.





Artificial Intelligence (AI)

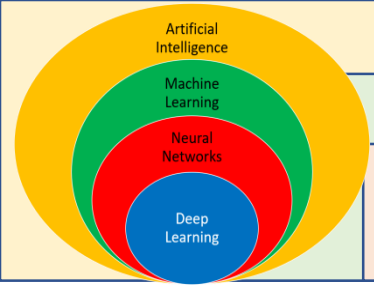
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Pooling Layer

- The Convolutional Layer and the Pooling Layer, together form the i -th layer of a **Convolutional** Neural Network.
- Depending on the complexities in the images, the **number** of such layers may be increased for capturing low-levels details even further, but at the cost of more computational power.



Artificial Intelligence (AI)


Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Pooling Layer

- After going through the above process, we have successfully enabled the model to **understand** the features.
- Moving on, we are going to flatten the final output and feed it to a regular Neural Network for **classification** purposes.



Artificial Intelligence (AI)


Machine Learning (ML)

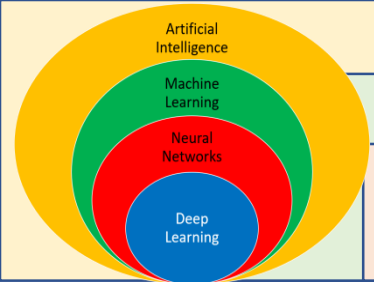
Neural Networks (NNs)

Deep Learning (DL)

Classification – Fully Connected Layer (FC Layer)

- Adding a Fully-Connected layer is a way of learning **non-linear** combinations of the high-level features as represented by the output of the convolutional layer.
- The Fully-Connected layer is learning a possibly non-linear function in that space.





Artificial Intelligence (AI)

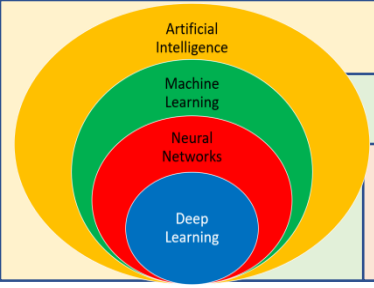
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Classification – Fully Connected Layer (FC Layer)

- Now that we have converted our input image into a suitable form for our Multi-Level Perceptron, we shall flatten the image into a column vector.
- The **flattened** output is fed to a feed-forward neural network and backpropagation applied to every iteration of training.
- Over a series of epochs, the model is able to distinguish between dominating and certain low-level features in images and classify them using the **Softmax Classification** technique.



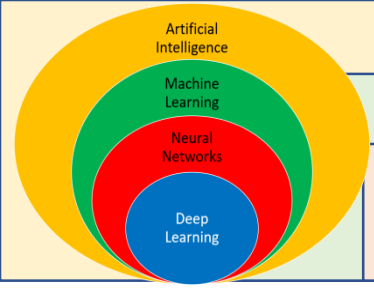
Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Part 2: Activation functions



Artificial Intelligence (AI)

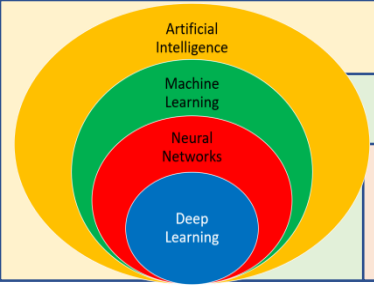
Machine Learning (ML)

Neural Networks (NNs)


Deep Learning (DL)

What is an activation function?

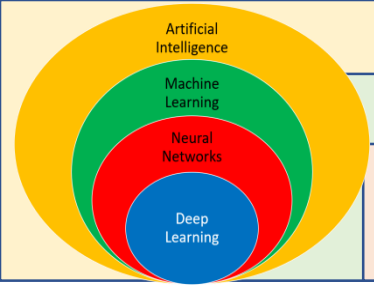
- An activation function is a function that is added into an artificial neural network in order to help the network **learn complex patterns** in the data.
- When comparing with a **neuron-based model** that is in our brains, the activation function is at the end deciding what is to be **fired** to the next neuron.
- That is exactly what an activation function does in an ANN as well. It takes in the **output signal** from the **previous cell** and converts it into some form that can be taken as input to the next cell.



What is an activation function?



A cartoon drawing of a biological neuron (left) and its mathematical model (right).



Artificial Intelligence (AI)

Machine Learning (ML)

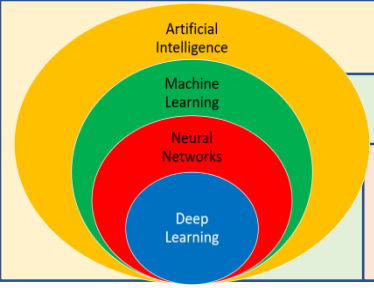
Neural Networks (NNs)

Deep Learning (DL)

Why is there a need for it?

There are multiple reasons for having non-linear activation functions in a network.

- Apart from the biological similarity, they also help in **keeping** the value of the output from the neuron **restricted** to a certain limit as per our requirement.
- This is important because input into the activation function is $W*x + b$ where W is the weights of the cell and the x is the inputs and then there is the bias b added to that.
- This value if not restricted to a certain limit **can go very high** in magnitude especially in case of very deep neural networks that have millions of parameters.



Artificial Intelligence (AI)

Machine Learning (ML)

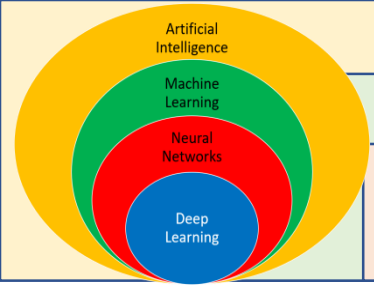
Neural Networks (NNs)

Deep Learning (DL)

Why is there a need for it?

There are multiple reasons for having non-linear activation functions in a network.

- This will lead to computational issues.
- For example, there are some activation functions (like **Softmax**) that output specific values for different values of input (0 or 1).
- The most important feature in an activation function is its ability to add **non-linearity** into a neural network.



Artificial Intelligence (AI)

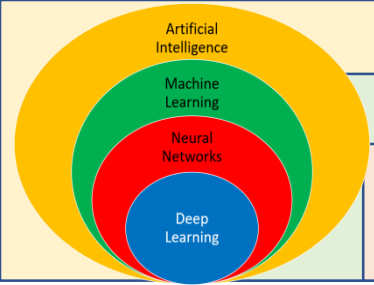
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Desirable features of an activation function

- **Vanishing Gradient problem:** The gradients tend to vanish because of the depth of the network and the activation shifting the value to zero. This is called the vanishing gradient problem. So we want our activation function to not shift the gradient towards zero.
- **Zero-Centered:** Output of the activation function should be symmetrical at zero so that the gradients do not shift to a particular direction.



Artificial Intelligence (AI)

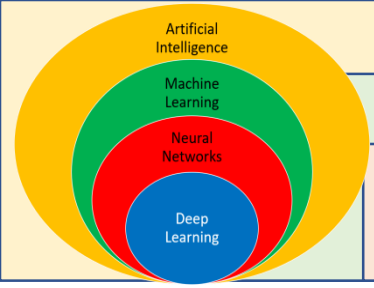
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Desirable features of an activation function

- **Computational Expense:** Activation functions are applied after every layer and need to be calculated millions of times in deep networks. Hence, they should be computationally inexpensive to calculate.
- **Differentiable:** Neural networks are trained using the gradient descent process, hence the layers in the model need to be differentiable or at least differentiable in parts. This is a necessary requirement for a function to work as an activation function layer.



Artificial Intelligence (AI)


Machine Learning (ML)

Neural Networks (NNs)

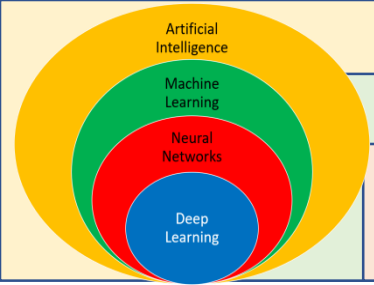
Deep Learning (DL)

Various non-linear activations in use

➤ Sigmoid



Sigmoid is generally used for **binary** classification problems.



Artificial Intelligence (AI)


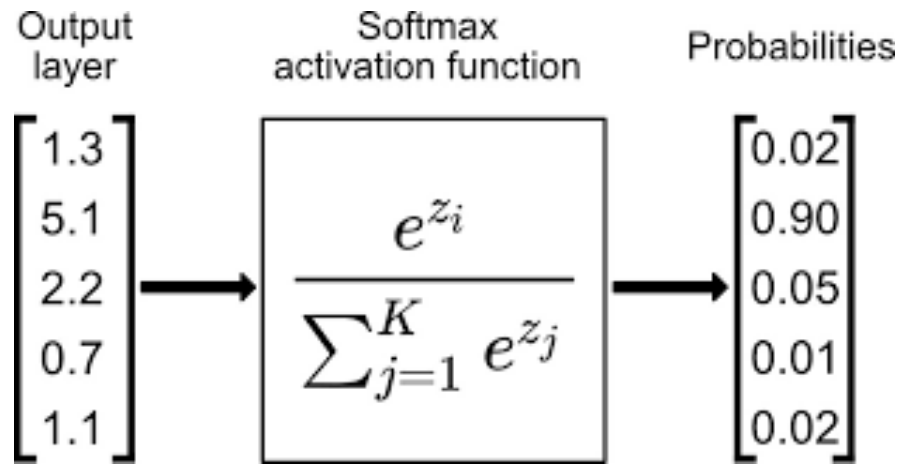
Machine Learning (ML)


Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

- **Softmax:** The Softmax is a more generalized form of the sigmoid. It is used in **multi-class classification** problems. Similar to sigmoid, it produces values in the range of 0–1 therefore it is used as the final layer in classification models.





Artificial Intelligence (AI)


Machine Learning (ML)

Neural Networks (NNs)

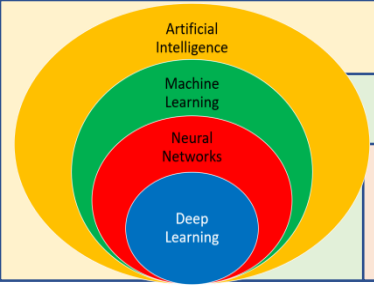
Deep Learning (DL)

Various non-linear activations in use

➤ Tanh



If you compare it to sigmoid, it solves just one problem of being zero-centered.



Artificial Intelligence (AI)


Machine Learning (ML)

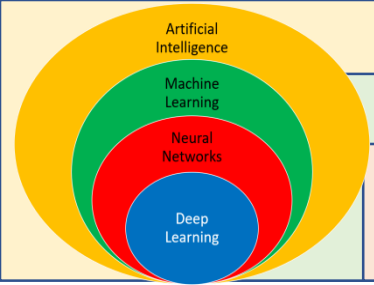
Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

➤ **ReLU:** ReLU (Rectified Linear Unit) is defined as $f(x) = \max(0, x)$





Artificial Intelligence (AI)


Machine Learning (ML)

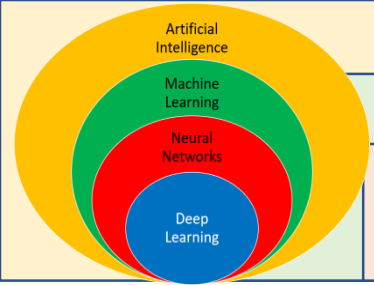
Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

- **ReLU:** ReLU (Rectified Linear Unit) is defined as $f(x) = \max(0, x)$
- This is a **widely** used activation function, especially with Convolutional Neural networks.
- It is easy to compute and **does not saturate** and **does not cause the Vanishing Gradient Problem**.





Artificial Intelligence (AI)


Machine Learning (ML)

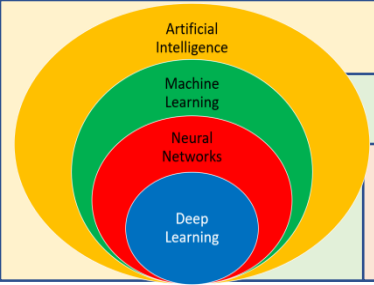
Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

- **ReLU:** ReLU (Rectified Linear Unit) is defined as $f(x) = \max(0, x)$
- It has just one issue of **not being zero centered**.
- It suffers from “**dying ReLU**” problem. Since the output is zero for all negative inputs. It causes some nodes to completely die and **not learn anything**.
- Another problem with ReLU is of **exploding** the activations since its higher limit is, well, inf. This sometimes leads to unusable nodes.





Artificial Intelligence (AI)


Machine Learning (ML)


Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

➤ **Leaky ReLU and Parametric ReLU:** It is defined as $f(x) = \max(\alpha x, x)$





Artificial Intelligence (AI)


Machine Learning (ML)


Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

- **Leaky ReLU and Parametric ReLU:** It is defined as $f(x) = \max(\alpha x, x)$
- Here α is a **hyperparameter** generally set to 0.01.
- Clearly, Leaky ReLU **solves** the “**dying ReLU**” problem to some extent.
- Note that, if we set α as 1 then Leaky ReLU will become a **linear function** $f(x) = x$ and will be of no use.





Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)


Deep Learning (DL)

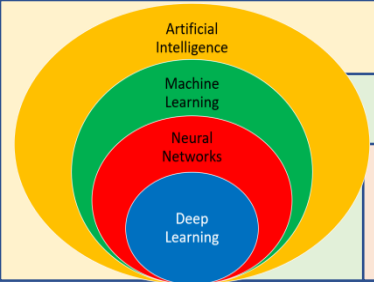
Various non-linear activations in use

➤ **Leaky ReLU and Parametric ReLU:** It is defined as $f(x) = \max(\alpha x, x)$

➤ Hence, the value of α is **never** set close to 1.

➤ If we set α as a hyperparameter for **each neuron separately**, we get **parametric** ReLU or PReLU.





Artificial Intelligence (AI)

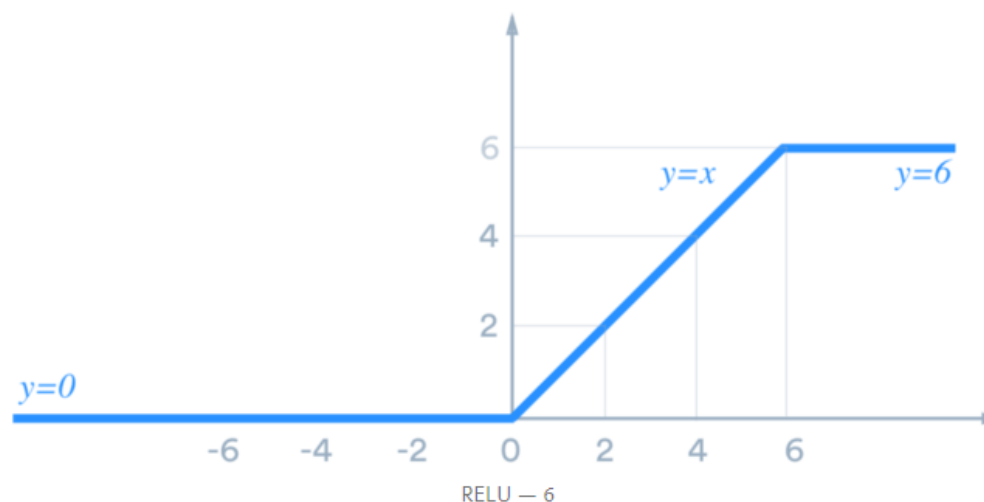
Machine Learning (ML)

Neural Networks (NNs)

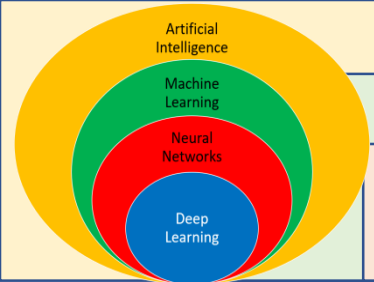
Deep Learning (DL)

Various non-linear activations in use

- **ReLU6**: It is basically ReLU restricted on the positive side and it is defined as $f(x) = \min(\max(0, x), 6)$



- This helps to **stop** blowing up the activation thereby stopping the gradients to **explode** (going to inf) as well another of the small issues that occur with normal ReLUs.



Artificial Intelligence (AI)


Machine Learning (ML)

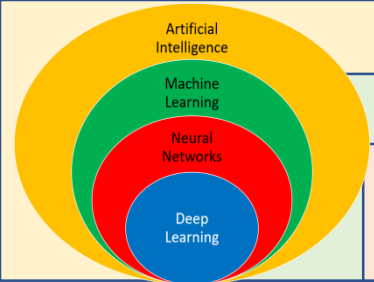
Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use

➤ **Swish:** It is defined as $f(x) = x * \text{sigmoid}(x)$.





Artificial Intelligence (AI)


Machine Learning (ML)

Neural Networks (NNs)

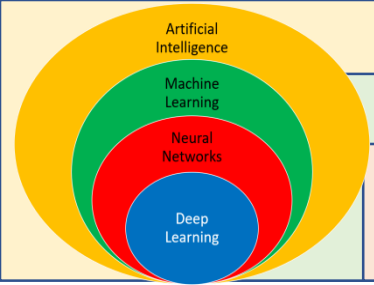
Deep Learning (DL)

Various non-linear activations in use

➤ **Swish:** Hard-Swish or H-Swish



➤ The best part is that it is almost similar to swish but it is **less expensive computationally** since it replaces **sigmoid** (exponential function) with a **ReLU** (linear type).




Artificial Intelligence (AI)

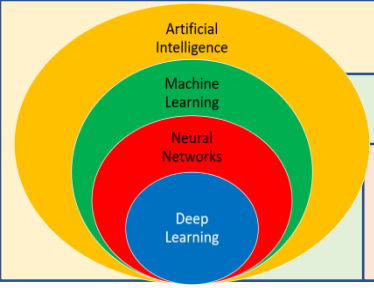
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Various non-linear activations in use





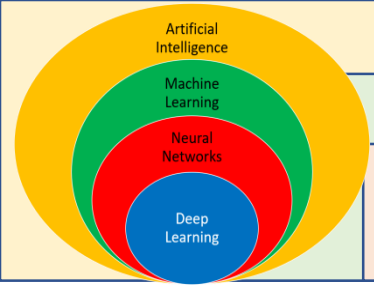
Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Part 3: Normalization



Artificial Intelligence (AI)

Machine Learning (ML)

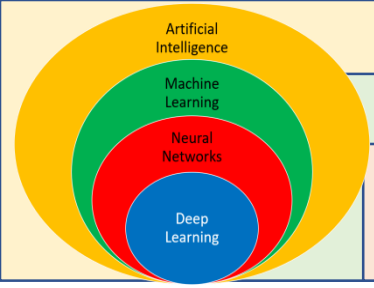
Neural Networks (NNs)

Deep Learning (DL)

Normalization

➤ Imagine that we have **two features** and a simple neural network. One is **age** with a range between 0 and 65, and another is **salary** ranging from 0 to 10 000. We feed those features to the model and calculate gradients.

- LeNet
- AlexNet
- VGGNet
- GoogLeNet
- ResNet
- ZFNet



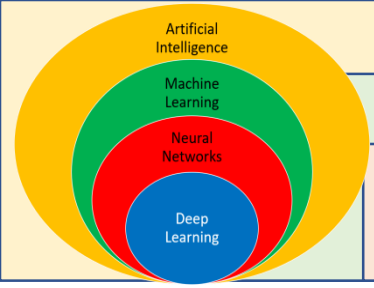
Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Part 4: Pre-trained Models



Artificial Intelligence (AI)

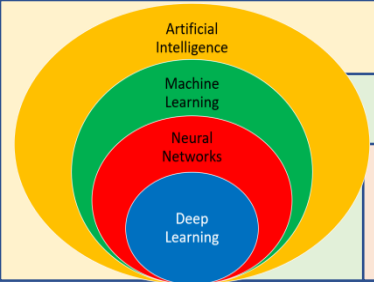
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

Pre-Trained CNNs

- There are various architectures of CNNs available which have been key in building algorithms which power and shall power AI as a whole in the foreseeable future. Some of them have been listed below:
 - LeNet
 - AlexNet
 - VGGNet
 - GoogLeNet
 - ResNet
 - ZFNet





Artificial Intelligence (AI)

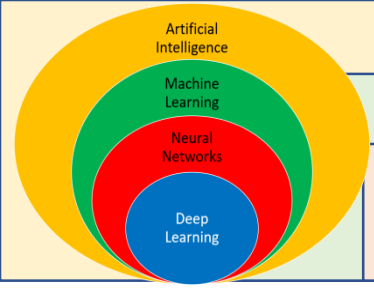
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

LeNet






Artificial Intelligence (AI)

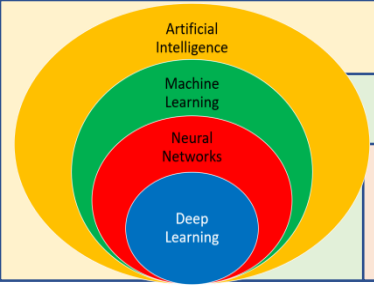
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

AlexNet






Artificial Intelligence (AI)

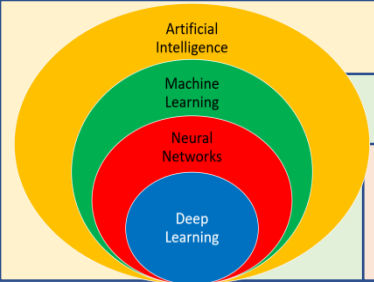
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

VGG16






Artificial Intelligence (AI)

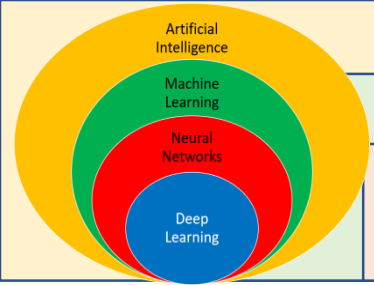
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

VGG19






Artificial Intelligence (AI)

Machine Learning (ML)

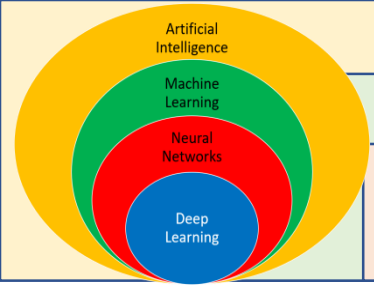
Neural Networks (NNs)

Deep Learning (DL)

GoogleNet (Inception v1)



22 layers deep




Artificial Intelligence (AI)

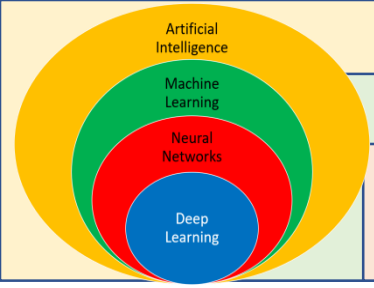
Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

ResNet50





Artificial Intelligence (AI)

Machine Learning (ML)

Neural Networks (NNs)

Deep Learning (DL)

