IOT ONLINE COURSE

Fundamentals of Artificial Intelligence

F-AI-3: Deep Learning

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham
Université de Pau, France

Slides realized by J. Mantilla

IoT – from idea to reality

What is Deep Learning?

- Learning from examples.
- A Machine Learning (**ML**) technique that teaches a computer to filter inputs (observations in the form of images, text, or sound) through layers in order to learn how to predict and classify information.
- Inspired by the way that the human brain filters information!
- Involves mostly using Artificial Neural Networks (**ANN** or **NN**)
- Deep Learning (**DL**) models can be supervised or unsupervised.

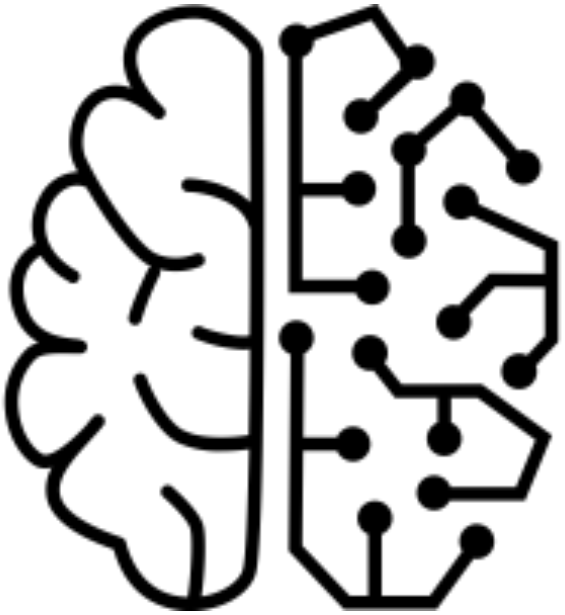

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

- Combines advances in computing power and special types of Neural Networks to learn complicated patterns in large amounts of data.
- State of the art for identifying objects in images and words in sounds.
- Applied successes in pattern recognition to more complex tasks such as automatic language translation, medical diagnoses and numerous other important social and business problems.

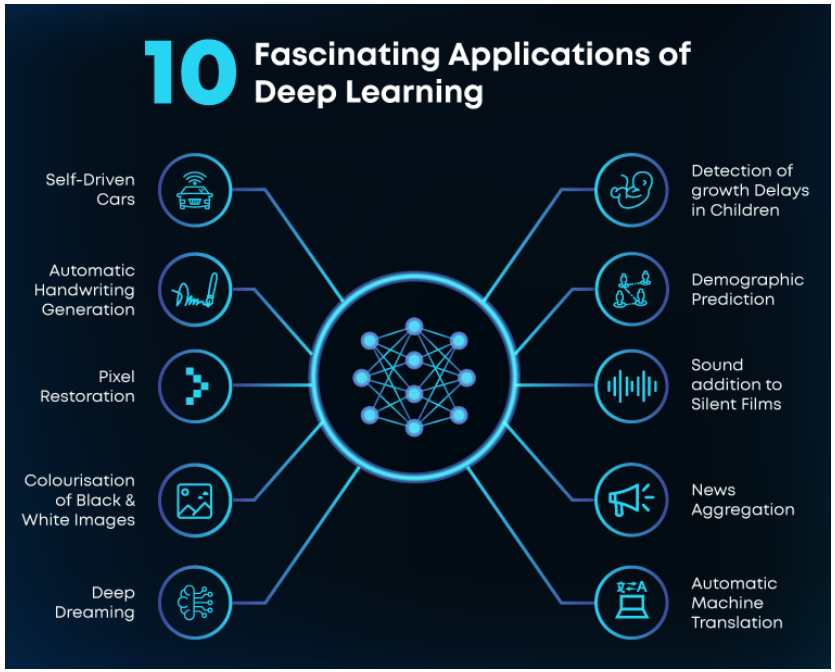

Perceptron

- Mathematical representation of a biological neuron
- First implementation by Frank Rosenblatt in the 1950s
- Rosenblatt's perceptron is activated when there is sufficient stimuli or input. (Neurons have been found to perform a similar process, in which experience strengthens or weakens dendrites' connections)

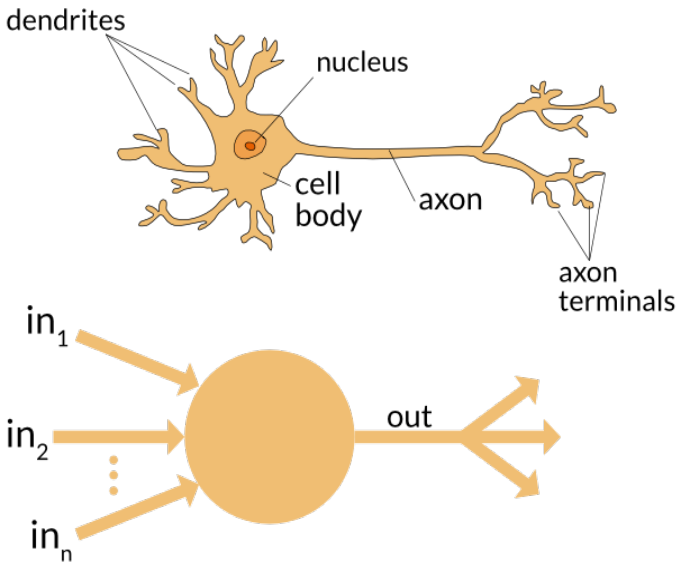

How does a Perceptron Work?

- Perceptron receives the value of the attributes of an input, just as dendrites do in a neuron.
- Each attribute has a **weight** that measures its *contribution* to the final result, which is the sum of the multiplications of inputs of each attribute by its corresponding weight.
- If the sum is greater than zero Perceptron returns a value of 1, otherwise it yields 0.

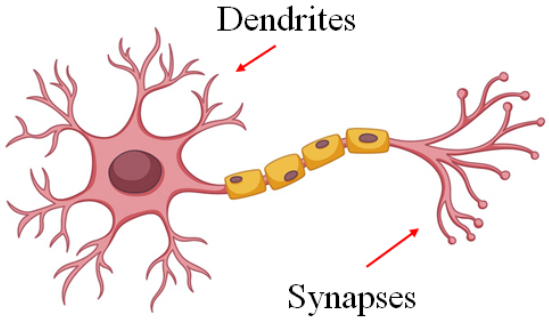

NEURON

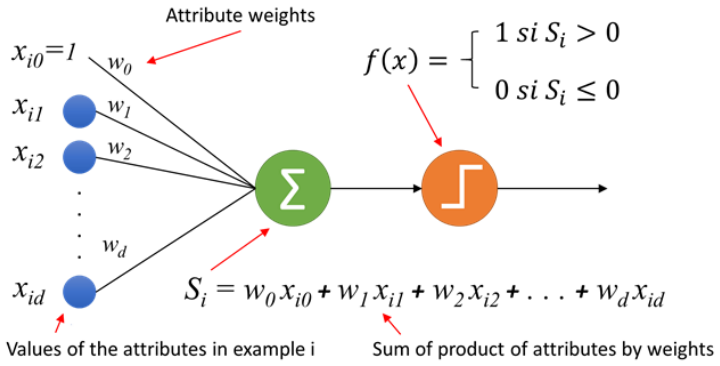

PERCEPTRON

Neural Networks

- Neurons by themselves are kind of useless, in large groups, they work together to create some serious magic!
- Neural Networks are no more than a **stacking** of multiple *perceptrons* in layers to produce an output.
- Input into one layer that creates an output which in turn becomes the input for the next layer, and so on. This happens until the final output signal.

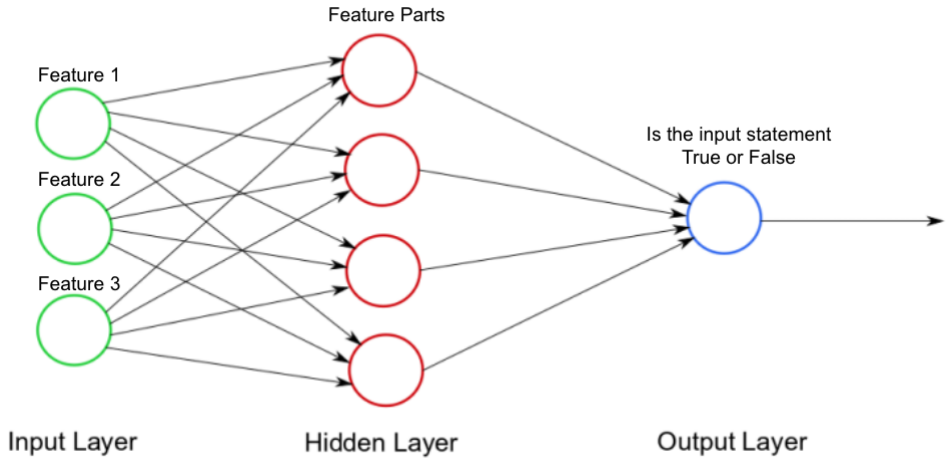

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

Deep Neural Networks

- In the 1980s, most Artificial Neural Networks (**ANN**) were single-layered due to the cost of computation and availability of data.
- Nowadays is possible to afford more hidden layers in ANN, hence the moniker “Deep Neural Networks” (**DNN**).
- Regained popularity since ~2006.
- Rebranded field as Deep Learning (**DL**)

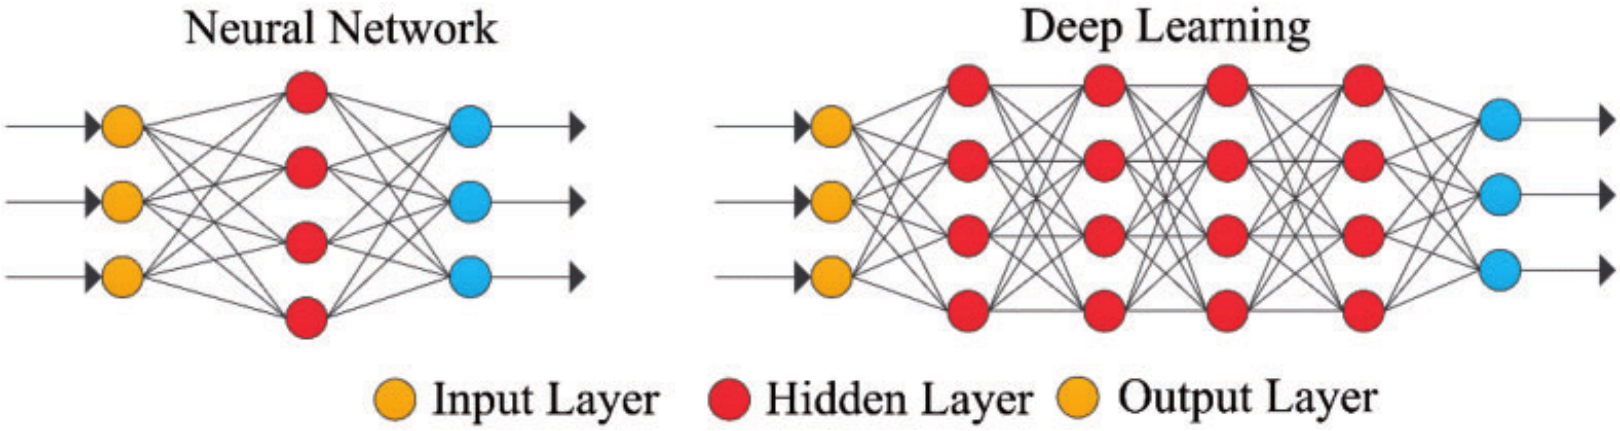

Loss Function

- Also known as Cost Function
- Weights must be adjusted (**learned** from the data)
- Idea: define a function that tells us **how "close"** the network is to generating the desired output
- **Minimize** the loss ➔ **optimization** problem
- With a **continuous** and **differentiable Loss Function**, we can apply *gradient descent*
- **Retro-propagation**. Iterative algorithm that, after each repetition, travels backwards a DNN adjusting the weights of the connections based on the contribution to the error.
- This is how DNN learn!

Gradient Descent

"We place an imaginary hiker at different points (on the Loss Function) with just one instruction: Walk only downhill until you can't walk down anymore." [Andrew Ng]

- The hiker doesn't *actually know* where she's going – she just feels around to find a path that might take her down the mountain, to the Global Minimum of the Loss Function.
- Minimize the Loss Function ➔ Better outputs/predictions.
- The DNN learns!

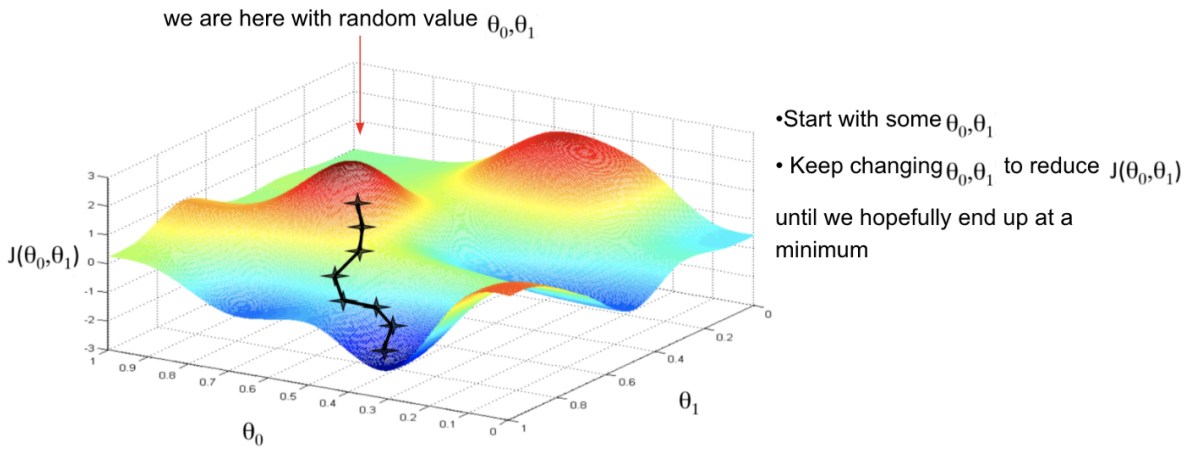

Types of Deep Neural Networks

- Feedforward Neural Networks (**FFNs, ANNs** or **NNs**)
- Recurrent Neural Networks (**RNNs**)
- Convolutional Neural Networks (**CNNs**)
- Autoencoder Neural Networks (**AEs**)

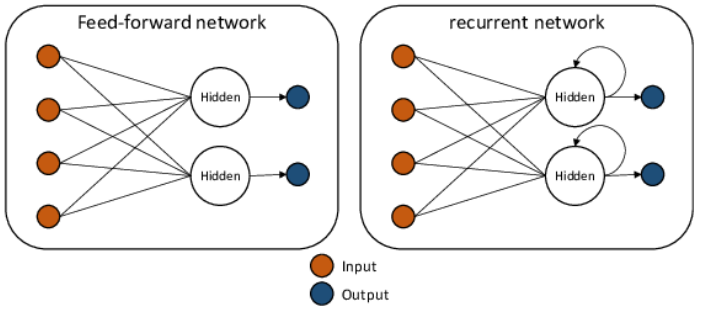

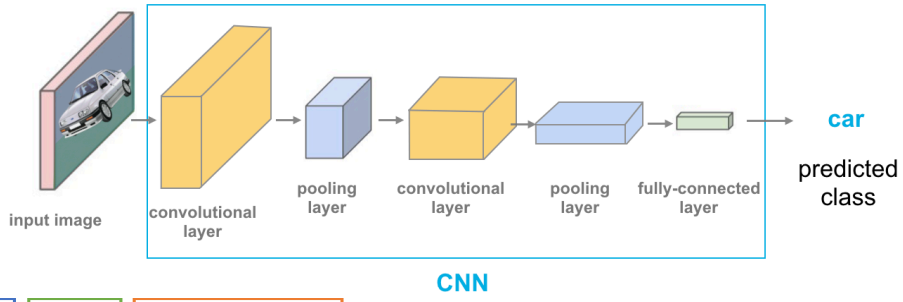

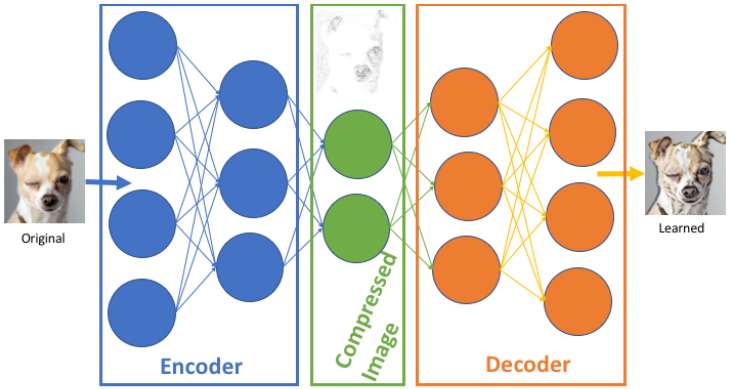

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

Feedforward Neural Networks

- In which each perceptron in one layer is connected to every perceptron from the next layer.
- Information is fed forward from one layer to the next in the forward direction only.
- There are no feedback loops.
- The classic ANN with several layers.

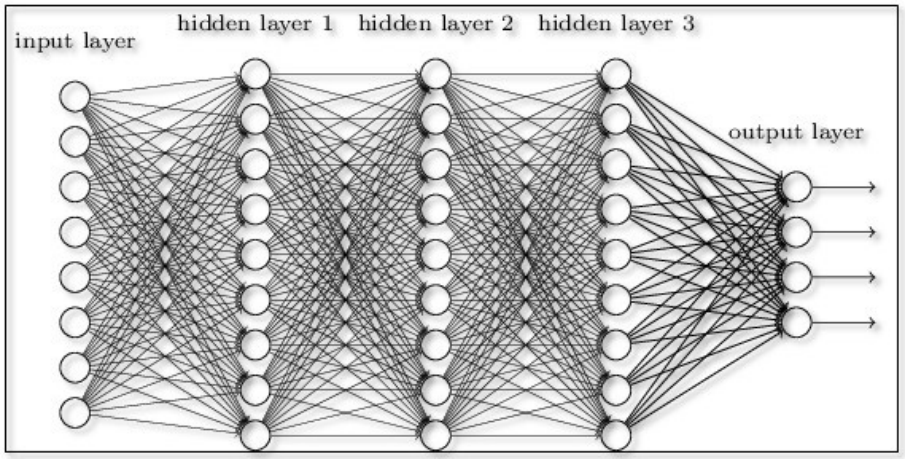

Recurrent Neural Networks

- Use sequential information such as time-stamped data from a sensor device or a spoken sentence, composed of a sequence of terms.
- Unlike traditional neural networks, all inputs to a recurrent neural network are not independent of each other, and the output for each element depends on the computations of its preceding elements.
- RNNs are used in forecasting and time series applications, sentiment analysis and other text applications.

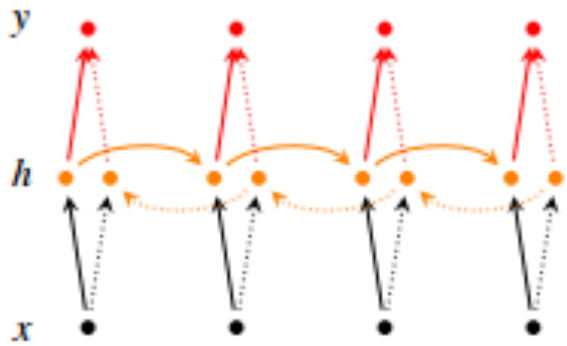

Figure 6: A bi-directional RNN model

Convolutional Neural Networks

- Contain five types of layers:
  - Input
  - Convolution
  - Pooling
  - Fully connected
  - Output.

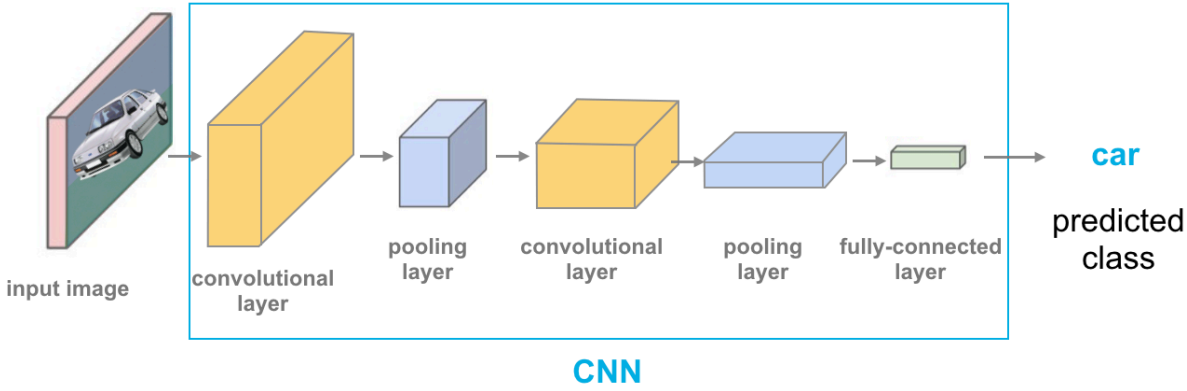

- Each layer has a specific purpose, like summarizing, connecting or activating.
- CNN have popularized image classification and object detection.
- Also applied to other areas, such as natural language processing and forecasting.

Autoencoder Neural Networks

- Used to create abstractions called encoders
- Created from a given set of inputs.
- Autoencoders seek to model the inputs themselves
- Unsupervised method.
- The premise is to desensitize the irrelevant and sensitize the relevant.
- As layers are added, further abstractions are formulated at higher layers (layers closest to the point at which a decoder layer is introduced).

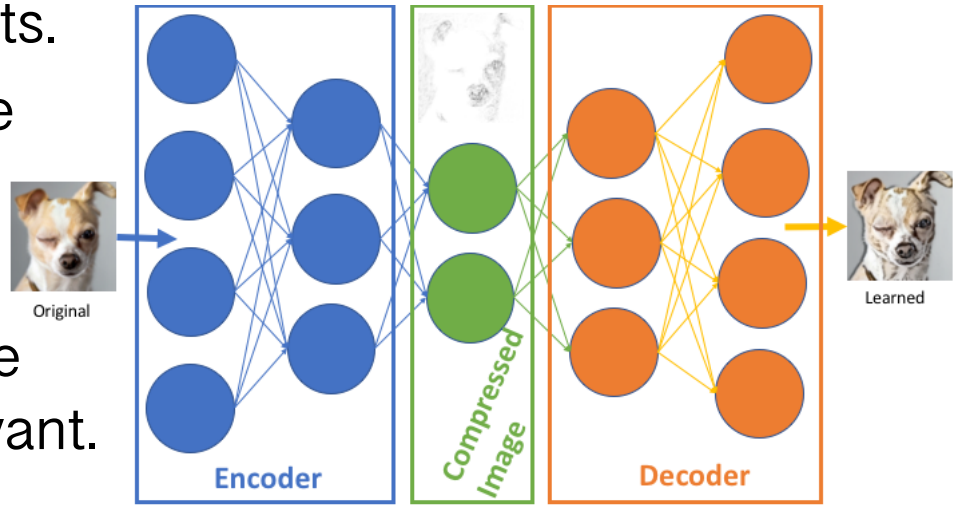

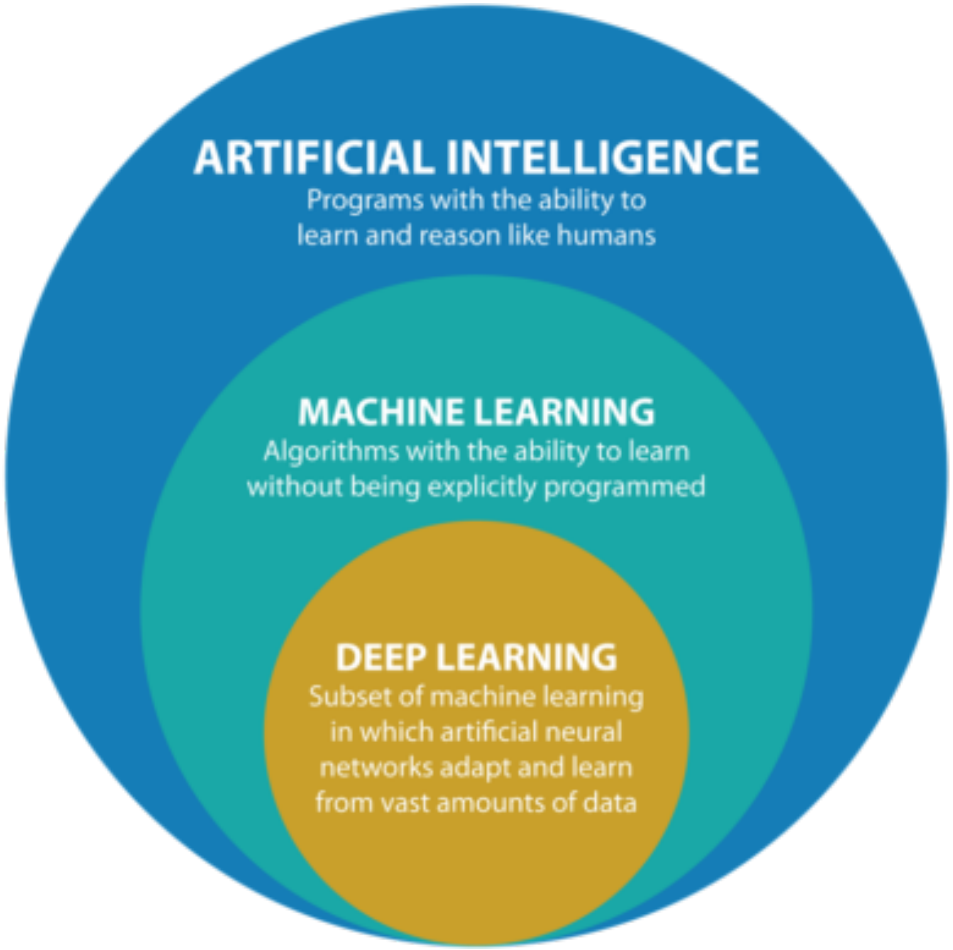

IOT ONLINE COURSE

Fundamentals of Artificial Intelligence

Continue with
F-AI-4: Regression

WAZIUP

IoT – from idea to reality

WAZIHUB