

Introductory talk on Artificial Intelligence, Machine Learning & Deep Learning

M. Phiri¹

`mitchell.phiri@cern.ch`

The Second African Conference on Fundamental and Applied Physics, 7-11 March 2022

¹Faculty of Engineering and the Built Environment
University of Johannesburg



Introduction

What is AI, ML and DL and what are the differences?

Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) are three different things.

- **AI** studies ways to build intelligent programs and machines that can creatively solve problems
- **ML** is a subset of AI that provides systems the ability to automatically learn and improve from experience without being explicitly programmed
- **DL** is a subset of machine learning, which uses the neural networks to analyse different factors with a structure that is similar to the human neural system.

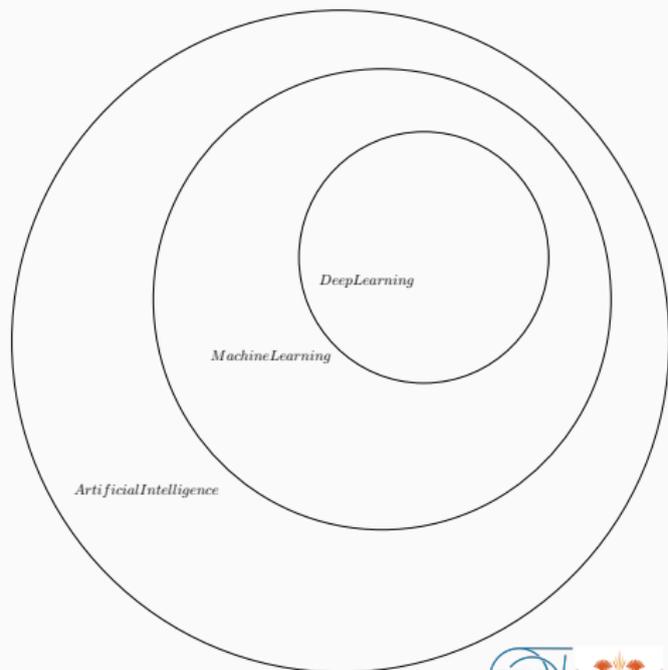


Figure 1: AI Hierarchy

Machine Learning Algorithms can be subdivided into four groups, based on how they learn namely:

1. **Supervised Learning:** a teacher helps the program throughout the training process: there is a training set with labelled data. Examples: Support Vector Machines, Decision Trees, K-Nearest Neighbours. Used for: spam filtering, language detection, computer vision
2. **Semi-supervised Learning:** the input data is a mixture of labeled and unlabelled samples
3. **Reinforcement Learning:** very similar to how humans learn: through trial. Examples: Q-Learning, Genetic algorithm, SARSA, DQN, A3C. Used for self-driving cars, games, robots
4. **Unsupervised Learning:** you do not provide any features to the program allowing it to search for patterns independently. Examples: Latent Dirichlet allocation (LDA), Singular Value Decomposition (SVD), Principal Component Analysis (PCA). Used for data segmentation, anomaly detection, recommendation systems



Deep Learning

- In a neural network, the information is transferred from one layer to another over connecting channels
- All neurons have a unique number called bias. This bias is added to the weighted sum of inputs reaching the neuron, to which then an activation function is applied
- The result of the function determines if the neuron gets activated. Every activated neuron passes on information to the following layers
- This continues up to the second last layer. The output layer in an artificial neural network is the last layer that produces outputs for the program

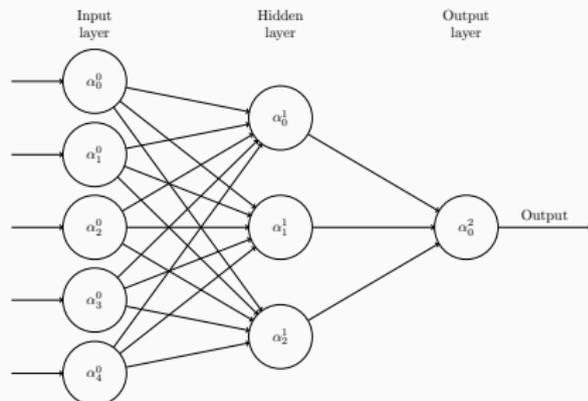


Figure 1: Neural network layers

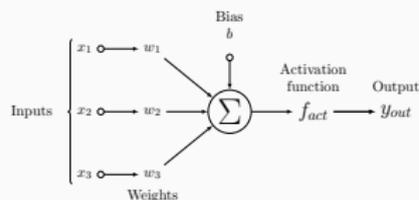
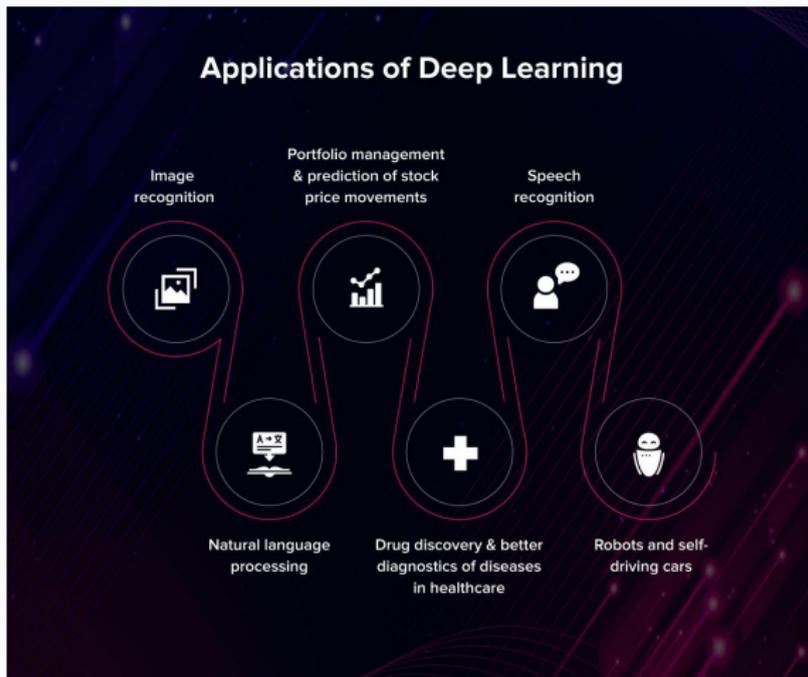


Figure 2: Neural network pipeline



Example & Resources

Where to start?

- For anyone interested in pursuing this avenue of analysis there's an overwhelming amount of free resources available on the internet and they have all been listed in the notebook that is accompanying these slides
- This list is not exhaustive but possesses more than enough resources to get very deep into AI starting with learning the basics such as Linear Algebra, Statistics and Python, to learning about different Deep Learning/ Machine Learning frameworks to seeing the latest State-of-the-Art papers along with discussions from the authors as well as places to discuss your own personal research with other practitioners
- My recommendation is to use Python as it has the most resources for ML



Thank You!

