

AI

WHAT

Broadly speaking, anything that lets computers mimic human behaviour can be termed as artificial intelligence.

SUB-CATEGORIES

Machine learning is a sub-set of artificial intelligence. Broadly speaking, it can be:

1. Supervised learning – learn from data given by humans (input and output mapping).
 - a. Examples include linear regression, overfitting, gradient descent, logistic regression, SVMs, k-nearest neighbours, decision trees, random forests, etc.
2. Unsupervised learning – No output given by humans, computers expected to find their own patterns.
 - a. Clustering, k-means, dimensionality reduction, etc.
3. Neural networks or deep learning – a collection of connected nodes is an artificial neural network (ANN), inspired by a human brain, and a huge ANN is commonly referred to as “deep” learning.
 - a. Chess, image classification (famous cat experiment by Andrew Ng), adding sound to silent movies, etc.

IMPACT OF GPU's

Huge amounts of matrix multiplications, a key element in deep learning, can be done in parallel on GPUs that can have hundreds of cores while typically, CPUs usually don't have more than 12 cores. GPUs also have a higher bandwidth to retrieve from memory.

BIG DATA and DEEP LEARNING

The best performing nets are deep and deepness of a net is directly proportional to the amount of training data fed to it. Bigger the dataset, better the net in terms of accuracy. A general rule of thumb is “at-least 10 times the degree of freedom”. So, for example, if you have 3 weights, you should at least have 30 data points. Not all data can be stored in memory and thus comes the need to store data on disk.

HOW IS DATA STORED?

Flatfiles, HDFS, RDBMS, NoSQL, or custom data structure.

MAJOR FRAMEWORKS USED

TensorFlow, Torch, Theano, Caffe, Deeplearning4j, Keras, Microsoft Cognitive Toolkit, NVIDIA Deep Learning SDK, etc.

WHO ARE THE BIG PLAYERS?

Dell – Offers 5 different types of boxes specifically designed for Machine Learning that has SSDs and some with NVIDIA GPUs. Have also partnered with [Bright Computing Software](#) and ships their boxes with it. Bright includes most of the famous Machine Learning softwares we discussed above. Also have a [HPC cluster management service](#).

NVIDIA – Offers [DGX](#) systems built on NVIDIA's GPUs. [HERE](#) is a nice 1.5 minutes video on DGX systems. Also offer a [GPU cloud](#) on Amazon EC2. They have also forked some of the open source ML frameworks and created their own [SDK](#).

IBM – They call it [Power AI](#) that is built on their S822LC systems that includes NVIDIA GPUs. PowerAI includes most of the popular open source frameworks. They also offer this solution in their cloud built on [Nimbix](#).

OTHER SMALLER PLAYERS

Exact – Offers servers with NVIDIA GPUs.

Amax – Offers servers with NVIDIA GPUs. Storage is based on their [STORMAX](#) product.

HPE – Collaborated with NVIDIA and provides HPE series for Machine Learning solutions. Their portfolio is [HERE](#).

CRAY – Provides supercomputers accelerated with NVIDIA GPUs specifically designed for Deep Learning.

REFERENCES

1. Najafabadi, Maryam M., et al. "Deep learning applications and challenges in big data analytics." *Journal of Big Data* 2.1 (2015): 1.
2. Chen, Xue-Wen, and Xiaotong Lin. "Big data deep learning: challenges and perspectives." *IEEE access* 2 (2014): 514-525.
3. Ingo Mierswa. What is Artificial Intelligence, Machine Learning, and Deep Learning? <https://rapidminer.com/artificial-intelligence-machine-learning-deep-learning/>, 2017. [Online; accessed 02-Nov-2017].
4. Michael Copeland. What's the Difference Between Artificial Intelligence, Machine Learning, and Deep Learning? <https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/>, 2016. [Online; accessed 02-Nov-2017].
5. Larry Hardesty. Explained: Neural networks. <http://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>, 2017. [Online; accessed 03-Nov-2017].
6. Jason Brownlee. 8 Inspirational Applications of Deep Learning. <https://machinelearningmastery.com/inspirational-applications-deep-learning/>, 2016. [Online; accessed 03-Nov-2017].
7. Ottokar Tilk, Carlos E. Perez, Tim Dettmers. Why are GPUs well-suited to deep learning? <https://www.quora.com/Why-are-GPUs-well-suited-to-deep-learning>, 2014. [Online; accessed 07-Nov-2017].
8. Geoffrey Craig. Deep Learning To Enhance Big Data Analytics, 2016. [Online; accessed 07-Nov-2017].
9. Zygmunt Z. How much data is enough? <http://fastml.com/how-much-data-is-enough/>, 2013. [Online; accessed 07-Nov-2017].
10. Dell EMC. Dell EMC Machine Learning and Deep Learning. <http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Machine-Learning-Solution-Overview.pdf>, 2017. [Online; accessed 08-Nov-2017].