Artificial Intelligence and Machine Learning

1. What is Machine Learning (ML)?

The general idea about Machine Learning (ML) can be traced back to 1959 with the approach proposed by Arthur Samuel, one of its pioneers. He defined ML[1] *as the field of study that gave computers the ability to learn without being explicitly programmed*.[2] In his IBM days, Arthur Samuel developed a computer programme that could play checkers game better than him. Machine Learning was developed as a subset of Artificial Intelligence that enables computer programmers to develop applications more lucidly. In 1998, Tom Mitchell, another well regarded machine learning researcher, added to the definition, and stated that *a computer program is said to learn from experience E with respect to some task T and some performance measure P, if its performance on T, as measured by P, improves with experience E.* Figure 1 represents this concept graphically.

For example, a financial organization wants to detect credit card frauds for its existing customers. Based on the classical definition of ML presented in Figure 1, the primary task T of the system will be to identify whether a given transaction is good or bad successfully,

Figure 1. – A Computer Program Gains Experience (E) to Perform A Machine-learning Task (T) At A Specific Performance-level (P)

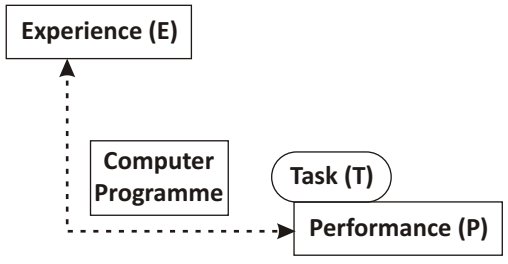

Source: Author's Compilation

and if possible, count the number of fraudulent transactions. The performance measure P for the ML task T is provided by the proportion of F within the total transactions. The fraud detection system learns through experience E, which has been built with the knowledge set of previous transactions. The experience was gained with time and through numerous iterations of the task, where transactions were successfully identified as good or bad. Therefore, the choice of P is crucial for an ML task.

2. What is Artificial Intelligence (AI)?

Artificial Intelligence involves the development of machines which can

[1]Machine Learning is referred as ML in this article hereafter

[2]Wiederhold G & McCarthy J (1992), Arthur Samuel: Pioneer in Machine Learning, *IBM Journal of Research and Development*, Vol 36 (3), pp.329-331

accomplish tasks that would otherwise need human aptitude. Typical examples of AI[3] tasks include natural language processing, voice command recognition, and image identification. AI has a presence in driverless cars, Google Voice commands, auto-complete action in emails, and many more real-life activities. AI can be broadly categorized into two parts, *weak AI* and *strong AI*. Weak AI can respond to situations but cannot think for themselves whereas machines with strong AI can think and act just like human beings. In reality, most instances of AI are weak AI, and rarely any strong AI can be found. In 1965, Herbert Simon the founding father of AI, stated that *machines would be capable of doing any work a man can do.* This event was the stepping stone towards developing the concept of AI[3]. Some relevant examples of AI are visual perception, speech recognition, decision-making and translation between languages.

3. What is the Difference between Artificial Intelligence and Machine Learning?

In today's world, AI and ML are the buzzwords, and often they are used interchangeably. However, these two concepts are not the same. Machine Learning can be defined as a subset of Artificial Intelligence. AI is a broader concept than ML which signifies that machines can carry out tasks that typically require human knowledge. ML states that machines are given access of data, and they learn the techniques on their own. The emergence of ML has come after two significant breakthroughs - one is the idea that machines can learn by themselves (Samuel, 1959[4]), and the other is the evolution of the internet and the wave of *big data.* Big data usually refers to large and complex data sets which the traditional data processing techniques find quite difficult to deal with. However these data sets being exceptionally valuable from the point of view of the vast information hidden within, it is creating a new generation of decision support data management system. Grover et al. (2018) put forth that the market of big data technology and services is expected to grow at a 23.1 per cent compound annual rate, reaching $48.6 billion in 2019[5].

Figure 2. – Artificial Intelligence and Machine-Learning Disciplines

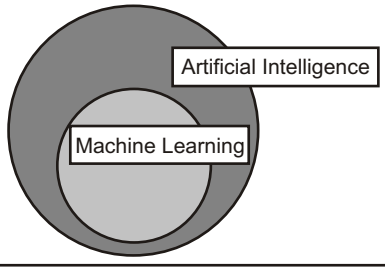

[3] Artificial Intelligence is referred as AI in this article hereafter

[4] Arthur Samuel is most known within the AI community for his groundbreaking work in computer checkers in 1959, and seminal research on machine learning

[5] Grover, V., Chiang, R.H.L., Liang, T-P & Zhang, D. (2018). Creating Strategic Business Value from Big Data Analytics: A Research Framework, *Journal of Management Information Systems,* Vol. 35 (2), pp. 388–423.

Further, the development of *neural networks* has propelled the faster emergence of AI. *Neural networks* were built to emulate human brains and neurons, hence the name. E.g. ML applications can read the text, and also can understand whether a statement is a complaint or a compliment.

4. How can Artificial Intelligence be useful for real world applications?

With increasing data availability, AI is now being used to solve real-life business problems and will be used to drive tangible business value for a wide range of industries. Often it will be useful for firms to look at AI through the lens of business capabilities rather than technology alone. Examples of AI include visual perception, speech recognition, decision-making, and translation between languages, self-driving cars, and humanoid robots (such as ASIMO and Boston Dynamics) among many others. AI can support three critical business needs: automation of business processes, gaining insight through data analysis and engaging with customers and employees.

5. What are the primary types of Machine Learning?

Machine Learning tasks are typically classified into three broad categories:

a. Supervised learning: Here, the machine learns by the experience gathered from the input data that is appropriately labelled (i.e., predictors and target/outcome variables are known apriori). The algorithm iterates until it achieves a desired level of accuracy.

When to use: Abundant labelled data is available.

Common Algorithms: Simple Regression, CART (Classification and Regression Tree), Random Forests, Naïve Bayesian Classifiers, Support Vector Machines.

b. Semi-supervised learning: The machine learns through an iterative training exercise, due to the scarcity of labelled data records.

When to use: Labelled data is hard to get in the public domain, Natural Language Processing.

Algorithms: Self-Training Algorithms; Transductive Support Vector Machine (SVM), Graph-based Algorithms

c. Unsupervised learning: Here, the machine cannot learn from the data because there are no target/outcome variables at all. This scenario is also known as non-labelled data. This type of techniques are particularly used for clustering, customer segmentation into different groups. Often researchers apply unsupervised learning to identify the best predictor/variables(s)

When to use: Labelled data is hard to get in the public domain, Sentiment Analysis of Social Network Messages.

Algorithms: Clustering, Association Rule Mining.

d. Reinforcement learning: Machine is exposed to an unknown environment and learns by trial and error method. Data (in the form of rewards and punishments) are given only as feedback to the programme's actions, such as driving a vehicle or playing a game against an opponent.

When to use: Dynamic or approximate dynamic programming; the only way to collect information about the environment is to interact with it.

Algorithms: Markov Decision Process, Monte Carlo Methods, State–action–reward–state–action (SARSA).

6. What are the risks associated with Artificial Intelligence?

While Artificial Intelligence has come a long way in accomplishing many technological goals, the risks associated with it are many and may turn evil. Researchers are not sure about the timeline by which AI may become super-intelligent and overpower humans. Most of them are assuming that human-level AI will happen before 2060 and if the projection is correct, it is high time we are aware of the probable malicious effects of AI on humankind. The worst fear of robots taking control over human by sheer use of Artificial Intelligence was accepted by Stephen Hawkins who claimed that "the development of full artificial intelligence could spell the end of the human race." The narrow AI system that is operative across the board is designed to outperform human intelligence, but researchers are striving to develop general AI that will be powered to beat every human cognitive task. If that is being achieved, AI may develop a destructive method of making its goals as human's desire and machine's target is almost impossible to align fully. Further, as intelligence enables controls, it is also a matter of debate whether in the long run machine will control over a human.

7. What are the Ethical Challenges associated with Artificial Intelligence?

As pointed out by Peter Norvig, Director of Research, Google, and one of the pioneers of Machine Learning, "the challenge now is to make sure everyone benefits from this technology.[6]" Researchers are trying to enhance the reach of AI to every one of the society including the bottom of the pyramid while trying to ensure that its advantages will not remain restricted only to the niche section of society. The software on which AI is based is often complex and difficult to understand, and hence the scrutiny needed for ethical acceptance of any software remains undone most of the time. This step reduces the moral dimension of a machine learning process or artificial intelligence, in due course of time which may pose a grave risk in the future. Further, without any standardized safety test,

[6]"Why the Biggest Challenge Facing AI is an Ethical One", BBC Future Now, 2017, http://www.bbc.com/future/story/20170307-the-ethical-challenge-facing-artificial-intelligence

policymakers remain suspicious to implement this technique on a mass scale.

The biggest ethical challenge, however, lies in the issue of machine replacing workforce and the question of tackling jobless economic growth. Whether increasing automation will lead to a spurt in unemployment is questionable. But at the same time, there are also evidence from countries like Germany or China where robots are working alongside human and even have already replaced large-scale workforce.

8. What are the open questions that Artificial Intelligence poses today?

a. Machines will rise to the extent that it will be smarter than a human being with too much power and control. The machine will dominate the planet with benevolence or malevolence.

b. Robots, which run on digital computers are assumed to have no subjective awareness and thus cannot understand qualitative perspectives. But, if general AI is developed, it may happen that the robots will "behave" like human without having any "sense."

c. Understanding ourselves well enough to build intelligent machines is an essential but highly debated criterion to achieve the next level of Artificial Intelligence. Replicating ourselves need a perfect understanding of the human brain's functioning and delivering Artificial Intelligence up to that mark is especially a tricky task faced by researchers.

d. As robots now get citizenship which ensures personhood, it is also questioned whether they have human rights.