The Evolution of AI

White Paper Published By

krtrima AI™
Intelligence Delivered

Jan-2020

Table of Contents

The objective of this paper is to describe how the Artificial Intelligence (AI) domain has evolved over the past few decades.

Before looking at AI's history, let's first understand where AI fits in the world of automation and how it is related to Information Technology (IT).

1. Introduction

The original objective of AI was to develop machines that could emulate human intelligence. Given this ambitious goal, despite more than eight decades of research and development, the progress has been slow. However, over the years there have emerged several useful applications that have had great impact on humanity.

AI is a sub set of the broader IT domain that deals with the automation of business processes to achieve scale, efficiency and reach. As the diagram explains, Machine Learning is a subset of AI, and Deep Learning is a subset of Machine Learning.

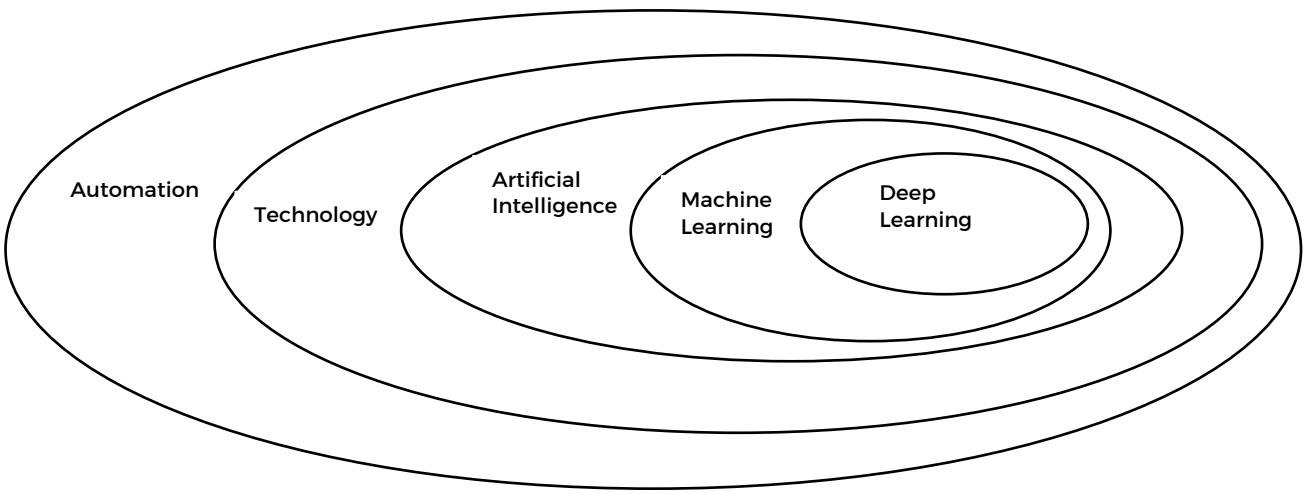

In traditional programming that is typically used in IT, we convert the existing process into a set of rules, apply them on the data to produce the desired output (answers). Classical AI (Business Intelligence) also works in similar fashion.

However, Machine Learning works on a different programming paradigm. It takes the data and answers as inputs and generates the rules as output. For example, in an image classification/ recognition problem, we present the image and contents of the image as inputs during the training, and the model learns rules that convert pixels in the image to what the image represents. After the training, if we present an image to the model, it will explain what's in the image using the rules it learnt during the training.

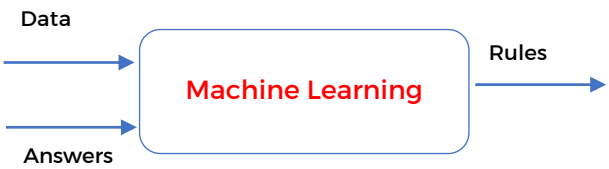

This approach has helped enterprises solve many complex problems, where it is nearly impossible to define a set of rules to describe a process, like while playing chess, or an object, such as photos, videos and speech.

Over past few decades, researchers have taken many different approaches to Machine Learning such as symbolic methods, probabilistic models, genetic algorithms, Deep Learning (Neural Networks) etc. All the recent success stories of AI such as digital assistants, autonomous cars, drones, language translation etc. use Deep Learning approach.

Let us now look at historic developments in the AI field.

2. The Gestation of AI (1940 to 1950)

The initial work on AI tried to model the human brain using the concept of neuron that goes "ON" or "OFF" based on input signal strength.

In 1943, Warren McCulloch (University of Illinois, Chicago) and Walter Pitts developed a neural network that could compute any computable function and implement all logical connectives (and, or, not, etc.). They also suggested that well-defined networks could learn with past history.

In 1949, Donald Hebb (McGill University, Canada) demonstrated a simple rule for updating connection strengths between neurons, which is now called Hebbian Learning.

In 1950, two undergraduate students at Harvard, Marvin Minsky and Dean Edmonds, built the first neural network computer, called SNARC, using 3000 vacuum tubes to simulate 40 neurons.

Between 1947 and 1950, Alan Turing (British scientist, Victoria University of Manchester) designed and promoted the Turing Test, machine learning, genetic algorithms, and reinforcement learning.

The Turing Test defined the criteria which, when met, enabled the machine to match human intelligence. It defined the rules for AI though it did not use the specific term.

The concept of autonomous vehicles/self-driven cars was first demonstrated by General Motors in 1939 at the Futurama exhibit New York world fair. The working model was released in 1958. However, it did not use any of AI techniques, but was an electric vehicle guided by radio-controlled electromagnetic fields generated with magnetized metal spikes embedded in the roadway.

3. The birth of AI (1956)

John McCarthy (Computer Science Professor, Dartmouth College, New Hampshire) convinced Marvin Minsky (MIT professor), Claude Shannon (MIT professor), and Nathaniel Rochester (Computer Scientist, IBM) to help him bring together US researchers interested in automata theory, neural nets, and the study of intelligence. These four together organized a two-month workshop at Dartmouth College in the summer of 1956.

The proposal for the workshop states:

"We propose that a two-month, 10-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve

themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer."

This conference is considered to be the birth of AI. It laid the foundation for AI, and set the goals and objectives for the field to evolve.

4. Early enthusiasm and great expectations (1950 to 1980)

Between 1952 to 1956, Arthur Samuel (MIT graduate, AI researcher at IBM) developed a series of programs for checkers that eventually learned to play at a strong amateur level.

The term Machine Learning was coined by Arthur Samuel in 1959.

In 1952, A team at Bell Labs designed the Audrey, a machine capable of understanding spoken digits.

In 1954, Georgetown-IBM experiment could successfully translate 60 Russian sentences into English. Researchers of this project claimed that machine translation would be a solved problem within 3 to 5 years.

In 1958, John McCarthy, after moving to MIT, developed the first AI programming language LISP. In the same year, he published a paper "Programs With Common Sense", in which he described "Advice Taker", a hypothetical program that can be considered as the first complete AI system.

In 1959, Herbert Gelernter (AI researcher at IBM) constructed the Geometry Theorem Prover, which is a tool to prove geometric theorems automatically.

In 1961, Newell and Simon from Carnegie Mellon University, developed the General Problem Solver (GPS), which mimicked human approach to problem solving. The success of GPS led Newell and Simon to formulate a hypothesis on physical symbol system.

Early work on neural networks was carried forward by Bernie Widrow (Stanford) and Frank Rosenblatt (Cornell) with their concepts of adalines and perceptrons in 1962. In the same year, Block et al published their perceptron convergence theorem.

In 1966, Joseph Weizenbaum, an MIT Professor, developed the first chatbot called Eliza. It worked on pattern matching techniques by dividing the user input into sentences and parsing it to look for key words and phrases. Then it responds by finding the closest match to the key words or phrases, in pre-programmed responses.

In 1969, Bryson and Ho developed the back-propagation learning algorithm for training neural networks in the name of "multi-stage dynamic system optimization method". For some reason this never got much attention from researchers and industry leaders. The one developed by Rumelhart (University of California, San Diego), Geoffrey Hinton (University of California, San Diego, later moved to Carnegie Mellon, then to University of Toronto) et al, later in mid 1980s became popular and is used even today in deep learning.

In the mid-1970s, knowledge-based systems, which came to be known as expert systems were developed. The first successful commercial expert system, R1, became operational at Digital Equipment Corporation in 1982. During the same period, initial work started in understanding natural language processing and machine translation using the same knowledge

representation schemes of expert systems. These initiatives led to the development of the Prolog language that became popular in Europe, and the PLANNER family in US.

In 1980s, Lidar and computer vision based autonomous cars and vehicles were introduced by Mercedes Benz and DARPA funded projects.

5. AI Winter (1970 to 2005)

Although initial attempts to create intelligent machines seemed successful, their weaknesses got exposed while trying to apply them in solving real life problems.

In 1969, Marvin Minsky & Seymour Papert published a book "Perceptrons", which proved how the perceptron algorithm was unable scale up to solve real life problems. With this all the optimism around neural networks initiatives died off.

Attempts to commercialize LISP machines (Hardware designed for LISP applications) failed.

Developments in Machine Translation, Speech Recognition, and chatbots could not scale up to solve real life problems, and constrained themselves to laboratories.

Expert Systems did not take off due to their high maintenance cost as their updates were expensive, and no learning mechanisms in built.

People soon realized that AI promises were hyped up, and the algorithms, computing power, knowledge repository available then were not sufficient to scale up to solve complex problems. Governments started cutting off the funds for research on AI.

https://en.wikipedia.org/wiki/AI_winter provides more details on this.

6. Renewed Interest in AI (2005 to 2019)

Despite serious setbacks, AI researchers continued to work on overcoming the limitations of algorithms.

In the mid-80s, several researchers improved the back-propagation learning algorithm which enabled the return of neural networks. They came on par with corresponding techniques from statistics, pattern recognition and machine learning. As a result of these developments a new field of data mining emerged.

In 1990, Yann LeCun, a leading AI researcher, developed an algorithm to recognize handwritten digits. He invented a concept called "convolution" which enabled neural network to recognize the object even if it was within the image. The US Postal Services adopted this to sort mails automatically based on zip codes.

IBM developed a chess playing machine, Deep Blue, which defeated, world champion, Gary Kasparov in May, 1997. It used Good Old Fashioned AI (GOFAI, classical AI) that uses rules based search and not a machine learning algorithm.

Over the years, machine translation and image recognition algorithms have improved by leaps and bounds. Advances in statistical models like the Hidden Markov Models helped address most of the problems associated with earlier solutions to speech recognition problem

After the arrival of Internet, e-commerce and social media platforms gained popularity, which generated a lot of unstructured data like text, images, videos, speech etc. At the same time, computing power kept increasing, while the cost of computing

power kept coming down, and enterprises started accumulating enormous amounts of data.

The high availability of computing power at lower cost, high volumes of data, coupled with advancements in algorithms and neural network architectures (Deep Learning) led to improvements in the performance of natural language processing, speech recognition and computer vision applications that led to following key developments.

In 2006, Netflix announced a competition to improve its recommendation system performance by at least 10%. A team of scientists from AT & T labs won this competition.

Google Translate was launched in 2006. It uses a proprietary statistical machine translation algorithm that reads vast amount of previously translated documents and finds patterns in them. Till 2016 it used classical machine learning techniques. From 2016 it started using deep learning approach, which improved the accuracy dramatically. It is now available in more than 100 languages with a fair degree of accuracy. However, in a few languages where previously translated documents are less, Google translator may not give good accuracy, as it could not learn the patterns adequately.

Google started self-driving car project in 2009, and it spun off as Waymo, a subsidiary of Alphabet Inc. (Google's parent company). Uber, Lyft, Tesla and many existing carmakers have also joined the race. A few of these cars are expected to be on road in early 2020 in the US.

Google voice enabled search came up in 2008, Apple announced Siri, a digital assistant, in 2011.

In 2011, IBM's Watson, a question-answering computer defeated Jeopardy champions Brad Rutter and Ken Jennings, and won the USD 1 million prize money

In 2014, Microsoft and Amazon developed their versions of digital assistants called, Cortana and Alexa respectively. In 2016, Google released its smart speaker and digital assistant Home and Assistant. In 2018, Google added duplex, a feature in its assistant, enabling it to take appointments/reservations for restaurants, salons etc. Duplex has a natural human voice, rather than robotic voice.

ImageNet is a large-scale visual recognition annual competition that was held between 2010 and 2017. Fei-Fei Li, a computer science professor at Princeton started the ImageNet project in 2007 to develop a very large data set of the world's objects organized into various categories with appropriate labels. The competition is to develop an algorithm that can recognize the objects accurately. In 2010, when it was launched winning team's accuracy was 71.8%. The algorithm identified the top five guesses and if the ground truth did not match with any of these, it was considered as an error, else accurate. Over a period of time, she and her team managed to collect 10 million images organized into 1000 categories.

In 2012, Geoff Hinton's team from University of Toronto, using a deep convolutional neural network architecture (called AlexNet), won the ImageNet competition with 84.7% accuracy (top-5 error rate of 15.3%). This was the turning point for deep learning approach to computer vision specifically, and AI in general. All the subsequent annual competitions, deep learning based algorithms won the competition, and further improved the accuracy to the level of 97.3% in 2017.

In 2015, Google launched Google Photos, a photo sharing and storage service. It recognizes and organizes photos based on the contents (people, places and things) of images, and enables search by content, and also organizes albums automatically.

DeepMind, a UK based AI company (now acquired by Alphabet Inc.), developed the AlphaGo program, which defeated Lee Sedol, world champion in Go game, in 2016. AlphaGo used two deep neural networks, data on past historical games and rules of the game.

Subsequently AlphaGo Zero and AlphaZero further improved the performance on multiple games like Go, Chess etc. AlphaZero plays itself to master the game without any historic games data. The only input required are the rules of the game. It also mastered multiple games, Go, Chess and Shogi (Japanese equivalent of Chess). It uses the combination of deep learning and reinforcement learning, now called as deep reinforcement learning. It has overcome two main limitations of the deep learning approach, namely requirement of large labeled data set, and solving narrowly defined problems (mastering only chess, Go or Shogi).

In 2018 Google launched its "Talk to Books" service. It uses natural language understanding to answer questions by reading 100,000 books that were pre indexed by Google.

7. The Road Ahead (2020 and beyond)

While we have seen great progress in Machine Translation, Natural Language Processing, Speech Recognition, and Computer Vision applications, none of them is fool proof. They all have severe limitations, and work in a closed environment within narrowly defined scope.

Research is going on to further improve these applications, as well as develop algorithms to expand the horizons of AI. A sample of these developments can be found here.

In 2006, Elon Musk founded Neuralink, a company focused on developing implantable brain computer interfaces (https://www.neuralink.com/).

Since 2017, Facebook has been working on developing non-invasive brain computer interface wearables, which can type without keyboard or mouse. By wearing the device overhead, one can transfer his/her thoughts from brain to computer, without typing it using a keyboard or a mouse.

There are many such research initiatives from Facebook, Google, Microsoft, OpenAI, deeplearning.ai and many innovative startups across the globe. Hence, we continue to see the advancement of these AI technologies, applications, and enterprise adaption.

8. Conclusion

Though AI has been evolving for years, several breakthrough products and solutions have come into the market using classical AI, Deep Learning, and reinforcement learning methods only in the past decade. Many of these solutions are being adopted by business enterprises and retail consumers, making them a huge commercial success.

Yet, the original goal of AI i.e. artificial general intelligence or human level intelligence in machines is still far from getting realized. Already many research initiatives are underway to further improve

the path to artificial general intelligence. We will continue to see new products and solutions every few months.