Inteligencia Artificial: pasado, presente y futuro

Artificial Intelligence: past, present and future

Gilles Paul Pieffet

Facultad de Ciencias
Universidad Antonio Nariño
gp.pieffet@uan.edu.co

Fotografía de Gilles Pieffet.

Keywords:
Artificial intelligence, Data science, Machine learning, Algorithms.

Palabras clave:
Inteligencia Artificial, Ciencia de datos, Aprendizaje automático, Algoritmos.

Artificial Intelligence can be defined as the automatic (i.e. without human intervention) making of a decision to solve a specific problem within a specific given situation. This work looks at the evolution of artificial intelligence from its early beginning until present day with a glimpse of what is yet to come.

Gilles Pieffet es investigador de la Universidad Antonio Nariño y trabaja en el estudio de sistemas biológicos a partir de simulaciones computacionales. Graduado de la Universidad Pierre et Marie Curie en Física-Química y Doctor en Matemáticas y Ciencias Naturales de la Universidad de Groningen.

Gilles Pieffet is a researcher at the Universidad Antonio Nariño working on the study of biological systems using computer simulations. He has a Msc. from the University Pierre et Marie Curie and a Phd. from the University of Groningen.

Abstract

Artificial Intelligence (AI) can be defined as the automatic (i.e. without human intervention) making of decisions to solve a specific problem within a specific given situation. This work looks at the evolution of artificial intelligence from the advances it brought and the struggles it faced in its early beginnings to the most recent successes and the renewed interest they ignited in present days. The scientific and technological foundation of the field will be overviewed. The influence of the dramatic increase in computational power will be discussed, together with the effect of the availability of cheap and accurate data at an unprecedented level and how it fueled the latest AI revolution. We will conclude with a glimpse of what is yet to come or is already there but still unnoticed: the very serious implications this technology will have at a societal level and the real threat it already poses to privacy

Resumen

La Inteligencia Artificial (AI, por sus siglas en inglés) se puede definir como la toma automática (es decir, sin intervención humana) de una decisión para resolver un problema específico dentro de una situación específica. Este trabajo analiza la evolución de la inteligencia artificial a partir de los avances que trajo y las luchas que enfrentó desde sus inicios hasta los éxitos más recientes que encendieron el renovado interés en tiempos mas recientes. Se analizarán los fundamentos científicos y tecnológicos del campo. Se discutirá la influencia del aumento dramático en el poder computacional, junto con el efecto de la disponibilidad de datos baratos y precisos a un nivel sin precedentes y cómo este alimentó la última revolución de la AI. Concluiremos con una rápida ojeada de lo que está por venir: las serias implicaciones que tendrá esta tecnología a nivel social y las amenazas reales que ya representa para la privacidad.

ciencia hoy

Introduction

Artificial intelligence (AI) can be defined as the automatic decision, making process to solve a specific problem based on a given input meant to describe a specific situation. It can also be more simply described as the science of having machines performing human tasks. In this work, we will look at the various phases AI went through from its early beginning more than half a century ago to its latest incarnation and the new ethical challenges it poses at a societal level.

History

The field of artificial intelligence (AI) was first developed in the 1950s, when scientist started wondering how computers could solve problems on their own. Among the most prominent pioneers of AI were Alan Turing, John McCarthy (who coined the term) and Marvin Minsky (Figure 1), whose work (Turing, 1950; McCarthy, 1958; Minsky, 1961) set the foundations of AI.

In one seminal publication in the field, Alan Turing asks if machine can think (i.e. if they are intelligent) and to answer this question went on to devise the Turing test, where an interrogator is to determine
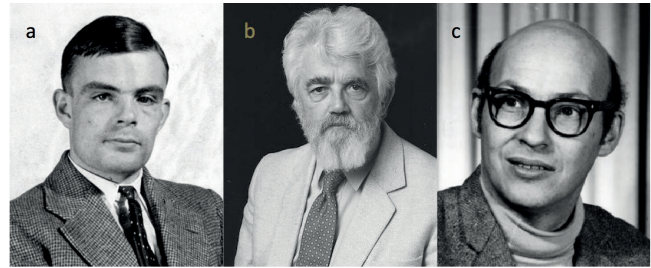

Figure 1: The pioneers of Artificial Intelligence (a) Alan Turing, (b) John McCarthy and (c) Marvin Minsky.

between two entities which one is a human and which one is a machine (Figure 2).
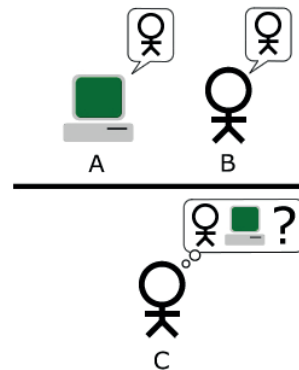

Figure 2: Representation of the Turing test. An interrogator C is set to determine which of entity A or B is human and which is a machine, solely from their written answers to his questions. If the interrogator cannot tell the difference between them then A is intelligent

The AI developed during this period would later be known as the classical AI. Classical AI was based on symbolic calculations, abstract symbols (words, numbers and mathematical operator) and on the belief that human intelligence could be fully represented using them. Classical AI led to advances in computational search and computational logic and enjoyed early successes such as the General Problem Solver, a computer program capable of solving

basic word problems and puzzles. However, the complexity of AI problems (increasing exponentially with the size of the problem) coupled with the limited computational power available at the time meant that it could only be practically applied to toy problems in very constrained environment. Public failures and clear deficiencies in applications such as machine translation that turned out to be slower, less accurate and more expensive than human translations led to a brutal loss of interest and lack of funding for research in AI that would be retrospectively known as the first AI winter (1974-1980).

Starting from the 1980s a new knowledge-based AI appeared that focused on how to capture, represent and infer knowledge. This led to the creation of so-called expert systems that were computer program designed to answer question or solve problems in a limited domain of knowledge. In these systems, knowledge was represented as a set of fact with rules used to connect them, forming a knowledge base. They also used an inference engine that deduced new facts from known facts contained in the knowledge base. These expert systems could perform tasks like diagnose diseases, detect chemical compounds from sensor reading or manage inventory parts, like the XCON system from DEC that could process orders with a 95% accuracy. However, these systems only proved useful in a few special cases throughout the 1980s because of their limitations.

This led to the creation of so-called expert systems that were computer program designed to answer question or solve problems in a limited domain of knowledge.

Ultimately, they would turn out to be very expensive to maintain (the system needing constant updating with new rules), besides being rigid and inflexible (producing vastly incorrect answers when given unusual input). They were also very slow, which limited their use to solve simple problems. The termination of the 5th generation computer project in Japan, designed to enable the future of AI technology, without reaching any of its objectives marked the beginning of the 2nd AI winter that would last from 1988 until the year 2000.

### Modern Artificial intelligence: data driven AI

The new AI paradigm came with the realization that knowledge could be learnt from data and a critical part of it is machine learning. Machine learning is such a important part of modern AI that it has now to a large extent become synonymous with it. However, while AI is the broad science of mimicking human abilities, machine learning is a specific subset of AI that teaches a machine how to learn using data and that is based on statistics. Machine learning builds analytical models automatically. To do this it uses methods from neural networks, statistics and physics to find hidden

patterns in data without being explicitly programmed where or what to look for or what to conclude.

The resurgence of AI is in no small parts due to advances in machine learning with the apparition of deep learning, which is the technology behind all the recent successful applications of AI, some of the most visible ones being the virtual assistants developed by the major Hi-Tech companies and now present everywhere: Siri for Apple, Alexa for Amazon, Google assistant for Google and Cortana for Microsoft just to name the most famous ones.

If machine learning is a subset of AI, deep learning is a subset of machine learning where neural networks are expanded into huge networks with a large number of layers trained using large amount of data. This was only made possible because of the availability of large amount of data, but also more than anything because of the explosion in recent years of parallel computing power coming from the development of the GPU architecture in graphics cards driven by the video game industry. As a consequence, milestones in the development of AI have been occurring increasingly rapidly. In 2009, Google self-driving Toyota Prius executes several trips of 100 miles. In 2011, IBM Watson won the US Quiz show Jeopardy. In 2016, Google DeepMind AlphaGo won against a human grand master in Go and since then more have occurred indicating that the interest isn't waning (see Figure 3 for a timeline of AI).

Figure 3. AI Timeline.
Source: Image courtesy of Dr. Paul Marsden (2017).
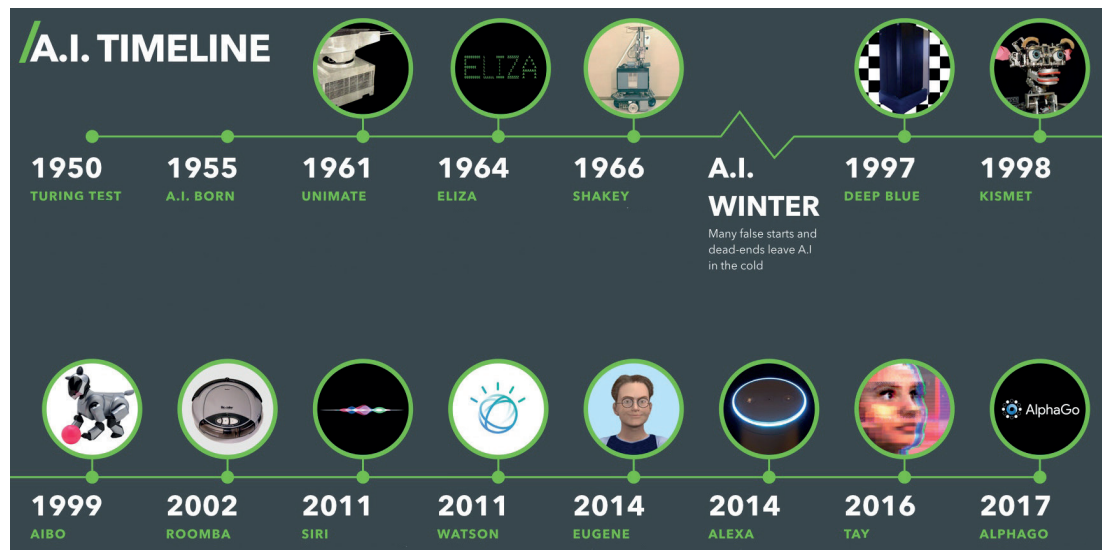
Conclusion

While it is clear that AI models will continue to improve in accuracy with time, what is less clear is what the implications will be in the long term and the challenges it will place on society. Some aspect will be largely beneficial especially in the healthcare domain where it will be of increasing help with result analysis, for example to diagnose difficult cases of tumor in x-rays. Other aspects will be more controversial such as the issue of privacy with the development of public video surveillance and facial recognition. Other implications will be clearly negative with the automation of an ever increasing of amount of jobs and it is unclear whether new jobs will be created at the same pace or in the same number than they will be destroyed. The future of Humanity Institute from Oxford University surveyed several hundreds of machine learning experts and asked them how long they thought it would take for AI to reach specific human capabilities (Grace *et al.*, 2018, p. 729). According to them, essays written by machines that could pass as written by human could happen in 2026 and truck driver could be replaced as soon as 2027. Other very interesting yet frightening predictions were that there is a high (50%) probability that AI will outperform humans in all tasks in the next 45 years and will automate most human jobs within 120 years. The time to start discussing these matters and to decide whether or not to place boundaries on the technology is now.

References

Grace, K., Salvatier, J., Dafoe, A., Zhang, B. and Evans, O. (2018). When Will AI Exceed Human Performance? Evidence from AI Experts. Journal of Artificial Intelligence Research, 62, pp. 729-754.

Marsden, P. (2017). Artificial Intelligence Timeline Infographic – From Eliza to Tay and beyond. Retrieved from: [https://digitalwellbeing.org/artificial-intelligence-timeline-infographic-from-eliza-to-tay-and-beyond/].

McCarthy, J. (1958). Programs with common sense, in National Physical Laboratory proceedings in Teddington, England, November 1958.

Minsky, M. (1961). Steps toward Artificial Intelligence, in Proceedings of the IRE, 49, 1, pp. 8-30.

Turing, A. (1950). I.—Computing machinery and intelligence, Mind, LIX (236), pp. 433-460

ciencia hoy