

THE NEW SPRING FOR ARTIFICIAL INTELLIGENCE

We interact with AI systems for a long time. Most of the time, the interaction was silent: we were not aware of the interaction. Nowadays, AI comes to the front page of journals: Watson won Jeopardy, AlphaGO won the world championship, an accident with an autonomous car, etc. In this article, we argue that the yeast of this vigorous AI Spring is machine learning and data sciences.

JOÃO GAMA^(1, 2),
ALÍPIO JORGE^(1, 3)

⁽¹⁾ INESC TEC;

⁽²⁾ FACULTY OF ECONOMICS OF THE UNIVERSITY OF PORTO;

⁽³⁾ FACULTY OF SCIENCES OF THE UNIVERSITY OF PORTO

joao.gama@inesctec.pt

alipio.jorge@inesctec.pt

The idea of Artificial Intelligence as we conceive it today has existed since the 1940s. Around 1940, Alan Turing, suggested that machines, like humans, could also think. In 1950, Turing wrote an article on the topic proposing to answer the question "Can machines think?"

The term Artificial Intelligence was first used in 1956, at a conference at Dartmouth College, organized by Marvin Minsky, John McCarthy, with the participation of Claude Shannon, Arthur Samuel, Allen Newell and Herbert A. Simon among others. All these scientists have played a very relevant role for decades in AI research.

As we conceive it today, artificial intelligence is a branch of computer science that aims to develop computational models that simulate the human ability to reason, make decisions and solve problems. These capabilities are defined by: ability to link reasoning, apply logical rules and derive conclusions; learning from facts and observations, making future actions more effective; recognizing patterns; applying reasoning to everyday situations. The 60s and 70s were years of inflated expectations. Or relative advances in areas such as automatic theorem demonstration, robotics, automatic translation, development of logic programming languages, etc. The initial success is well illustrated by the GPS - General Problem Solver, developed by Newell and Simon [1], capable of solving problems automatically. In the 1980s, Japan launched a 10-year project for the Fifth Generation of Computer Systems, to create computers using massive parallel computing and logical programming.

Since the 1970s, there has been an effort to use Artificial Intelligence to solve real problems. Initially, the problems were treated by AI through the acquisition of knowledge from specialists in any given domain. These systems, known as expert systems, were modular with the inference engine being independent of the knowledge base. For each specific domain, a knowledge base was built through interviews with experts in the field to discover the rules used by the expert to make decisions. In the late 1980s, early 1990s, expectations were followed by a period of marked disillusionment, leading to disinvestment in the area. AI had entered its winter. In 1997, IBM's DeepBlue computer won over world chess champion G. Kasparov. IBM stressed that the victory was due to the machine's processing capacity and not to the use of AI technologies!

In the 1980s, more sophisticated and autonomous computational tools for extracting knowledge from facts and data began to gain popularity. In the 90s, these tools gained maturity, beginning to be used in companies. These tools had great impetus with the development of computer networks and the WWW, associated with the ability to collect, store and process large amounts of digital information.

AI appeared on headlines of the newspapers when, in October 2005, the DARPA Grand Challenge took place in the Nevada desert. It was the first time that an autonomous car successfully completed the challenge. Sebastian Thrun said that "The robot's software system relied predominately on state-of-the-art AI technologies, such as machine learning and probabilistic reasoning" [2]. It is the beginning of a new spring of AI, which will be reinforced when machine learning is used in other AI areas such as Knowledge Representation, Computer Vision, Natural Language Processing. AI has developed algorithms and technologies capable of solving difficult problems, and which have come to be widely used in the most diverse sectors.

In 2011, the Mackinsey Institute published the report "Big data: The next frontier for innovation, competition, and productivity" [3] which relaunched public and private investment in AI technologies, machine learning and data science.

The greatest growth occurs in companies, where the use of AI is used as a business strategy, as is the case of Google and Facebook, even for the development of marginal business applications, such as the automatic assistants common in apps and websites of several banks. Netflix, for example, uses AI in its recommendation system and to identify the patterns of preferences of its users. Some examples in which Google employs AI, and which are already part of the daily lives of its users, we can mention: organization of photos in Google Photos, in which ML is used to identify the elements of the photos or to group the photos by standards among other uses. Automatic captions for videos on YouTube. Recommendation for quick responses to email messages in Gmail. Use of artificial neural networks, more specifically Deep Learning, to improve the effectiveness of translations in Google Translate.

Examples of successful applications of ML techniques to real problems include: interfaces that use natural language (written or spoken), facial recognition, spam filtering in emails, fraud detection by banks and credit card operators, decision support on medical diagnosis through analysis of clinical data, from images and/or genetic data, recommendation of products based on the consumer's profile and consumption history, intelligent behavior in game characters, playing Chess and Go with a performance comparable to that of human champions. An application with intensive use of AI and ML that is currently very popular is that of autonomous cars. Several manufacturers, such as Tesla, Volvo, BMW, Mercedes-Benz, Ford and Land Rover, have autonomous vehicle designs. Several commercial models already integrate technologies for partial autonomy. In these models, ML is used diversely in several tasks, such as the detection and recognition of objects and plates, object classification, location, prediction and tracking of moving objects.

The Internet of Things (IoT), counts on millions of devices sensing the environment, processing that information, and forwarding it to other machines. IoT is at the origin of technologies smart-cities, smart-grids, smart-farms, etc. The creation of a level of information about the production process gave rise to Industry 4.0, where decision processes involve people and machines. The machines' ability to explain how they came to a decision is critical to a climate of trust. On the other hand, a large part of the economy is developed in a virtual universe. Any company has a website, and some large companies, such as Facebook, Netflix, Airbnb, Uber, etc. only exist on the web. All companies are accessible 24/7 on our smartphones. Interactions with users are monitored in order to create user profiles: likes, preferences, fears, etc. These profiles can be used for marketing, recommendation or influencing choices. In the digital universe, where there are no borders, concepts such as privacy, the "public" and the "private", and ethics need to be rethought [4].

References

1. N Newell, A., Shaw, J.C., Simon, H.A. (1959). Report on a general problem-solving program. Proceedings of the International Conference on Information Processing, pp. 256-264.
2. Thrun S. et al, Stanley, The robot that won the DARPA Grand Challenge. J. Field Robotics 23(9): 661-692 (2006).
3. Manyika S., Chui M., Brown B., Bughin J., Dobbs R., Roxburgh C., and Byer A. (2011). Big data: The next frontier for innovation, competition, and productivity, Mackinsey Global Institute, Technical Report.
4. High-Level Expert Group on AI, Ethics Guidelines for Trustworthy Artificial Intelligence, European Commission, 2019.

Figure 1 - Autonomous Stanley Touareg.

