

Living with AI

Alfio Quarteroni



Narrow and General Artificial Intelligence

First, let's clarify the meaning of Artificial Intelligence.

When we talk about AI today, we refer to what is known as narrow or weak intelligence: systems designed to perform specific, limited tasks for which they are trained. Different, at least in theory, is general Intelligence (AGI, Artificial General Intelligence), considered the “Holy Grail” of research, towards which enormous investments are being concentrated. It is the type of intelligence that should be able to match human intelligence in every way. But, at least for now, we are still a long way from this goal.

The AI in use today responds to concrete objectives: to provide solutions to well-defined problems using specific models and algorithms.

In robotics, for example, the goal is to build machines capable of moving autonomously and safely, and therefore navigation and control algorithms are needed. In the medical field, on the other hand, image and clinical data analysis systems are designed to support professionals in diagnosis and therapy. In both cases, narrow AI develops tools that serve a specific purpose.

Generative intelligence – which produces text, images or videos based on human instructions – is a particular variation of this same family of algorithms. Here too, as in the previous examples, the underlying logic is that of machine learning: the ability to learn autonomously through a training phase made possible by a large amount of available data.

However, defining what intelligence really is, whether human or artificial, remains complex. When the term Artificial Intelligence was coined, the aim was to provide machines with software capable of simulating human behaviour, enabling them to solve tasks like us. This is the principle behind the famous Turing test, designed to assess whether a machine can exhibit behaviour indistinguishable from that of a human being.

Machine Learning

Today, we see examples of ‘intelligent’ machine behaviour every day: expert systems, robots capable of learning, self-driving vehicles, or our smartphones, which recognise images, understand speech and translate text in real time.

Machine learning is the paradigm that makes all this possible. To enable a machine to learn, rather than programming it step by step according to rigid instructions, it is trained on large amounts of data, thanks to which it will learn to solve problems independently. This is a radical change: the transition from deterministic programming to learning by experience, i.e. through data.

An effective example is language learning. We can study its rules and grammar theoretically, or we can learn it by living among those who speak it, listening and experimenting. This is what happens to children with their mother tongue: they learn it through experience, without knowing the rules, but by using it. In the same way, machines “learn” by analysing huge sets of texts, images or numbers that represent their world of reference.

Data is the real fuel of artificial intelligence.

Without large amounts of data, or information, AI performance would remain modest. Advances in machine translation, for example, stem from the enormous availability of multilingual texts – books, articles, conversations – on which algorithms are trained.

Neural networks, the heart of many AI systems, are mathematical models inspired by the functioning of the human brain. Trained on huge volumes of data – images, texts, sounds – they learn to recognise patterns, identify correlations and draw inferences. This enables them to identify a face they have never seen before or translate an unfamiliar sentence, based on statistical and probabilistic models learned during training.



Human and Artificial Intelligence

Comparing human and artificial intelligence is useful but inevitably approximate and misleading.

Humans learn through theoretical knowledge, reason, emotions, intuition and memory, building new knowledge based on experience. Machines, on the other hand, are not aware of the meaning of the data they process. They have no emotions, intentions or awareness: they do not “know” that they are translating a text or analysing an image; they simply perform a task assigned to them.

This is why AI results can sometimes appear trivial or predictable. However, in many areas, machines already surpass humans in terms of efficiency and accuracy: they can compare thousands of images in a matter of seconds, recognise unknown voices or summarise complex texts with remarkable accuracy. It is a complementary collaboration: we offer context, judgement and meaning; machines bring computing power and the ability to identify hidden patterns in data, which are often surprising even to us.

However, consciousness, awareness and common sense remain the prerogative of human intelligence. AGI – future general intelligence – aims to bridge this gap, but for now, humans retain a decisive advantage. In summary, the comparison between neurons and neural networks is useful for educational purposes, but it should not be misleading: machines imitate and extrapolate, generalising, but they do not understand. They do not know that they are behaving like a human being, an ant or a plant.

A telling example: after a few attempts, a child learns how to load a dishwasher; no robot is yet capable of doing so. An 18-year-old, after a few dozen hours of driving, tackles city traffic with an ability to adapt that autonomous vehicles are still trying to achieve.

ancora cercando di raggiungere.

Even machines make mistakes: Bias and Hallucinations

AI learns from data, but if the data is incorrect, partial or biased, the results will inevitably be misleading. It is easy to “trick” a system with distorted or manipulated information. Algorithms, in fact, have no ethical or moral criteria, nor the ability to assess the veracity of what they process, unless such criteria are explicitly included by programmers.

The latest generative AI platforms, such as ChatGPT, include security filters that prevent responses on illegal or dangerous topics. However, their nature remains dialogical: they are designed to converse and make communication pleasant rather than rigorous. We could say that they tend to “please” the interlocutor. Talking to them is more like chatting in a bar than consulting the Encyclopaedia Britannica.

Under human control

The rapid development of artificial intelligence generates conflicting feelings: enthusiasm and hope on the one hand, fear and mistrust on the other. Alarmist statements, especially when made by authoritative experts, receive widespread media coverage. These concerns are not unfounded: they relate to privacy, work, personal freedom and the very future of the human species. The idea that, one day, machines will be able to communicate with each other without human control – a remote but not impossible scenario – contributes to amplifying these fears.

Yet AI has enabled extraordinary scientific achievements in recent years. As a personal testimony, I proudly recall that my research group at the Politecnico di Milano created the first complete mathematical model of the human heart, now used in numerous hospitals in Italy and abroad. It is an important achievement, the result of the synergy between science, mathematics, physics and artificial intelligence. This news has had little resonance in the media, where catastrophic narratives about AI, which fuel fears and mistrust, are given ample space.

Artificial intelligence has enormous potential, but it must remain a tool under human control. With the AI Act, the European Union has introduced rules to assess its risks and limit its abuse. However, it would also be desirable to have a preventive consensus, similar to that adopted for nuclear energy: international agreements that set limits and responsibilities before problems arise. The Manifesto for an International Consensus on AI, recently signed by over 200 experts, is a step in this direction.

Being an optimist by nature, I believe that the major players in the sector – Google, Microsoft, OpenAI, Meta, Amazon – will have to listen to users' concerns and reassure them about the effects of their technologies. They will not do so out of philanthropy, but because it will be in their own interest.

Alfio Quarteroni is an Emeritus Professor at Politecnico di Milano and at EPFL, Lausanne. He is the founder of MOX at Politecnico di Milano and the co-founder and president of MOXOFF.

Quarteroni is a member of several prestigious academies, including the Accademia Nazionale dei Lincei, the European Academy of Sciences, the Academy of Europe, the Lisbon Academy of Sciences, and the Italian Academy of Engineering and Technology. He has authored 25 books (translated into several languages) and more than 400 research papers.

Quarteroni has been honored with numerous awards, including the NASA – Group Achievement Awards (1992), the Galilean Chair from Scuola Normale Superiore in 2011, the International Galileo Galilei Prize for Sciences in 2015, the ECCOMAS Euler Medal in 2022, the ICIAM Lagrange Prize in 2023, the Blaise Pascal Prize for Mathematics in 2024, the ECCOMAS Ritz-Galerkin medal in 2024, the SIAM Ralph Kleinman Prize in 2025. His research spans applications in medicine, earthquake geophysics, environmental science, aeronautics, and the oil industry. He led the mathematical modeling for the design of Alinghi, the Swiss yacht that won the America's Cup in 2003 and 2007 and developed the first comprehensive mathematical model of the human heart.

According to Google Scholar, he is the Italian mathematician with the highest number of citations and the largest H-index."

<https://mox.polimi.it/people/alfio-quarteroni/>

