

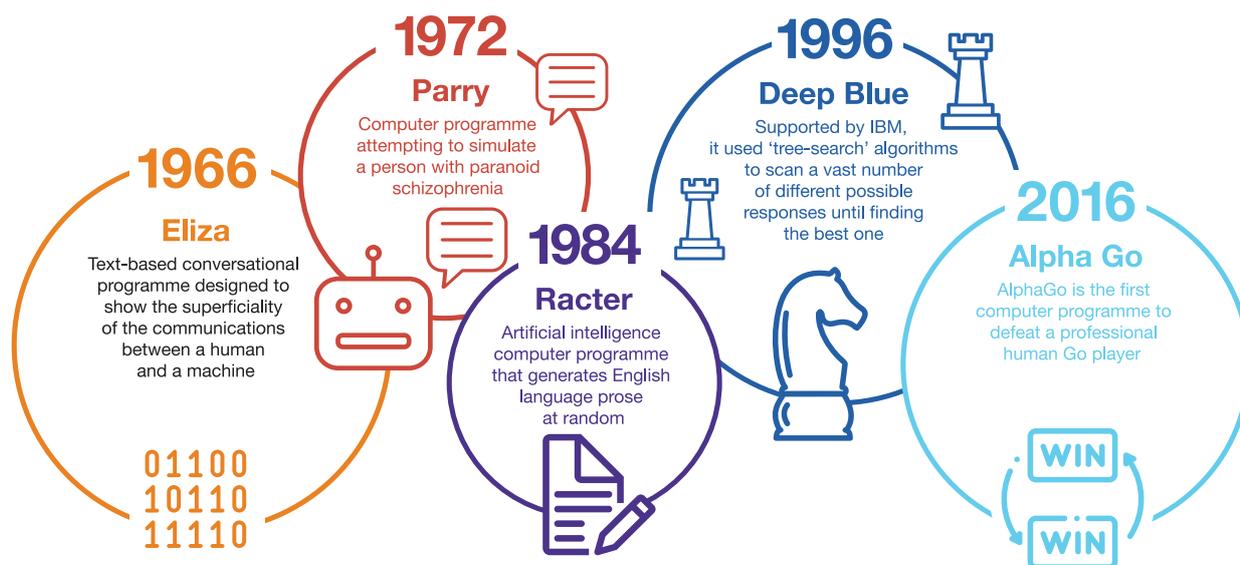
Byte-sized edtech research



AI Part 1 2019

Is Artificial Intelligence 'intelligent'?

Timeline of the development of AI



Logical Learning

Machine Learning

- In 1956, John McCarthy defined Artificial Intelligence (AI) as 'machines that have the ability to achieve goals like humans do'.
- Machines were said to achieve this by 'thinking humanly', 'acting humanly', 'thinking rationally' or 'acting rationally' Russel and Norvig (2016).
- Early AI systems were designed as a set of pre-programmed rules. It is extremely challenging to imitate human behaviour using this approach, because humans are not simply banks of knowledge that produce behaviour by following step-by-step, hard-wired rules. Humans have the ability to learn, adapt, feel, socialise, and respond unexpectedly (even irrationally).
- Newer AI systems are often based on Machine Learning (ML). ML techniques are based on finding statistical patterns in large collections of historical data.
- ML can only replicate human intelligence if it can access high-quality data that is representative of the specific problem, algorithms that can empower the ML to learn from this data, and affordable processing power and memory. Even with all this, ML is limited and can only 'imitate' the behavioural patterns that are shown in the data.
- Humans often make decisions based on intuition, instinct, and wrong assumptions. Machines use statistics, so are potentially more consistent.
- AI's capability to outperform human intelligence is limited to specific domains or tasks, such as playing chess. AI cannot yet transfer skills, respond creatively, or feel emotions.

“[AI is] machines that perform functions that require intelligence when performed by people.”

Kurzweil, 1990

“[AI is] the exciting new effort to make computers think... machines with minds, in the full and literal sense.”

Haugeland, 1965

“[Artificial General Intelligence is] the ability to accomplish any cognitive task at least as well as humans.”

Tegmark, 2017



AI Part 1 2019

01 Defining ‘Artificial Intelligence’ (AI)

In 1956, John McCarthy coined the term Artificial Intelligence (AI) and defined it as ‘the science and engineering of making intelligent machines that have the ability to achieve goals like humans do’.

In his proposal to a research workshop about AI, McCarthy stressed the assumption 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’. But we are still far from completely understanding human intelligence. What exactly is it that machines should imitate?

Conclusions: If AI is evaluated by achieving goals ‘like humans do’, we must first fully understand human intelligence to understand and design AI.

02 The rule-based early AI system (1960s – 1980s)

In the 30 years after the workshop that McCarthy attended in 1956, the first AI prototypes and systems were designed with the principal goal of imitating humans.

Three conversational programmes developed in this period are ELIZA (who imitated a Rogerian therapist), PARRY (imitating a person with paranoid schizophrenia) and RACTER (the amusing ‘artificially insane’ raconteur).

Those three simplistic bots emulated entire conversations using pre-programmed, rule-based logic. If the human’s statement does not fit the rough logic (always in English) that ELIZA has been programmed to respond to, ‘she’ can choose from a set of fixed phrases such as “Very interesting. Please go on.” or “Can you elaborate on that?”

Conclusions: Modelling real-world complex problems using a set of logic rules does not enable a machine to fully imitate human intelligence.

03 Machine Learning systems (1990s until today)

The dramatic increase in computer processing speed and capacity led IBM to develop Deep Blue in the 1990s. Deep Blue was an AI chess-playing programme, based on pre-programmed logic. Thanks to IBM’s massive-scaled hardware, it could defeat the world’s chess champion.

Years later, when developing its Go AI systems, Google’s engineers used ML techniques to defeat the (human) world Go champion in 2016. Unlike rule-based systems, ML techniques are based on statistical patterns found in large collections of historical data.

ML can sometimes perform specific tasks more effectively than humans. However, the extent to which this is possible depends on the availability of high-quality data, the

availability of sophisticated algorithms that learn, and a clear and well-articulated understanding of the problem to be solved. Even when all this is possible, ML is still limited to ‘imitating’ only those behavioural patterns that are shown in the data.

Both of these game-playing AI systems master a very specific task. They cannot play even a slightly different game, or show any intellectual ability. Thus, according to McCarthy’s definition, these two systems achieved their goal ‘like humans do’. However, having merely this single, non-transferrable skill cannot make them ‘intelligent’ in human intelligence terms.

Conclusions: Deep Blue and Go achieve their goal ‘like humans do’ and even outperform humans on a specific task within a specific knowledge domain. However, unlike human intelligence, their ‘intelligence’ cannot be transferred to other (even similar) tasks.