AI: Aiming for Intelligence

Paul S. Wang, Sofpower.com

July 19, 2020

Introduction

The term *Artificial Intelligence* or AI has seen wider and wider use, certainly in the past five years or so. Proponents say AI is the future of information technology (IT) and companies should get on the AI bandwagon. Tractica, a market research firm, forecasted huge market growth globally for AI software, expecting revenue to increase from about 10 billion U.S. dollars in 2018 to 126 billion by 2025.

That is certainly significant. But what is AI? What kind of intelligence is *artificial*? In what ways will AI help? What capabilities can it bring? Could AI become smarter than humans? Or, can it potentially take over the world and replace human beings all together?

We'll try to answer such questions in ways simple to understand. This article is part of my ongoing *Computational Thinking* blog (`computize.org/ctblog`).

What is Artificial Intelligence?

The word *intelligence* here refers to the ability to acquire and apply knowledge and skills. Intelligence is usually ascribed to human beings exclusively. But some animals are also intelligent. Dogs, for example, can learn to perform important tasks that humans can't even come close. Basically, the goal of artificial intelligence is to create man-made intelligent machines.

As the field of AI research and development advances, the meaning of the term AI has evolved. In practice, today AI usually refers to "*a computer system's ability to exhibit some kind of smartness*." A little history can shed more light on AI.

A Brief History of Artificial Intelligence

Of course, when we talk about intelligence we mean mental capacity and acumen, not information gathered by security agencies. As such, intelligence is not absolute but relative to some chosen baseline for comparison. Thus, natural intelligence is not easy to define and subject to different interpretations. Defining *artificial*, man-made or manufactured, intelligence then would be an even harder problem.

Alan Turing, the father of modern computer science, published a landmark paper in 1950 where he speculated about the possibility of creating machines that think. Because "thinking" is hard to define, he devised a test known as the *Turing Test*. Basically, if a machine can interact with a human via a teleprinter and make the person think another person is at the other end then it was reasonable to say that the machine was "thinking". Turing went on to argue convincingly that a "thinking machine" was at least plausible. The Turing Test was the first serious proposal in the philosophy of artificial intelligence.

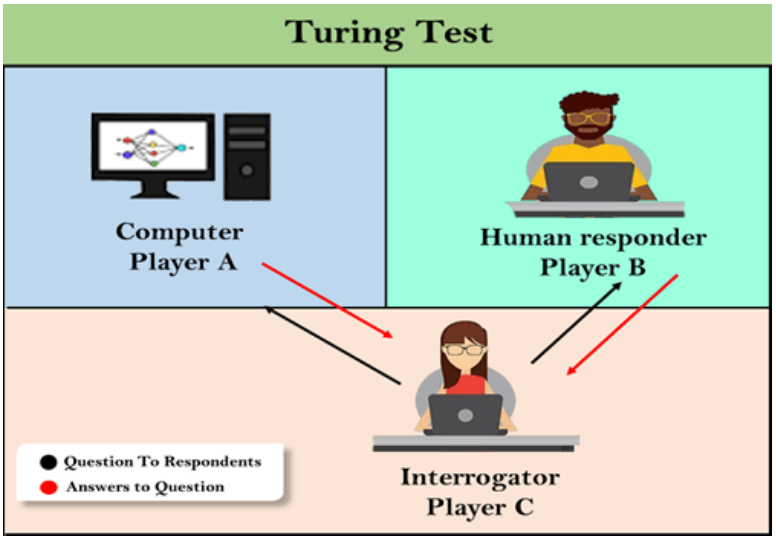

John McCarthy, one of its founding fathers, coined the term "Artificial Intelligence" in 1955, and organized the famous *Dartmouth Summer Research Project on Artificial Intelligence* in 1956. This meeting started AI as a field of research. McCarthy defines AI as "the science and engineering of making intelligent machines, especially intelligent computer programs." And in this context, he defines ingelligence as "the computational part of the ability to achieve goals in the world."

Henceforth, AI R&D endeavored to create systems that demonstrate their intelligence.

An early system *ELIZA* was created by Joseph Weizenbaum at the MIT AI Lab in 1966. ELIZA is a program that can carry on a conversation with a person by using simple rules and patterns. For example, if a user typed "I am not feeling well." ELIZA might display "Why are you not feeling well?" Or if the user cursed, ELIZA might respond "Did your parents teach you such talk?" It is an attempt to pass the Turing Test without the computer really understanding what's being said. The question has become "If passing (or nearly passing) the Turing Test is not enough, what does a program have to do to be considered intelligent?"

Next, programs that can play games were developed. Afterall, if a machine can play a game with a person and come close or even win then that machine is surely thinking or has intelligence. Such efforts include MIT's Greenblatt Chess playing system (1969), IBM's Jeopardy! (TV game show) winning system Deep Blue (2011), and DeepMind Technologies' AlphaGo (2016) which triumphed over the best human Go players in the world.

Yet, people ascribe these impressive performances to the raw speed and

data processing powers of modern computers rather than intelligence. Take action video games for example, does anyone really believe s/he can win against the computer? But that's not because it is smarter. Certainly, no one thinks Blackjack machines intelligent even when it can replace real dealers.

The more capable and wonderful computer system become, the more people realize that they are still just great programs, instead of possessing intelligence.

AI Systems in Practice

Proving man-made intelligence aside, attempts and efforts to build "intelligent machines" have led to many useful programming principles, techniques, paradigms as well as hardware devices. Their use resulted in many practical applications that enrich our lives as well as the field of AI.

In the early days, the MIT AI Lab (now MIT CSAIL) experimented with automatically balancing an inverted pendulum to simulate a child balancing a broomstick on an open palm. Much later, in 2001, the Segway HT, a two-wheeled, self-balancing personal transporter was invented and brought to market. It was a revolutionary riding device after the bicycle. Unfortunately, for various financial and safety reasons, it was not a commercial success.

Personal assistant apps, such as Google Assistant, Alexa, and Siri, use speech technologies to hear and execute user commands and to give verbal answers. They understand a given set of commands and can't really converse with the user. Pandora, Netflix, and Amazon Video can recommend songs or shows based on user experience and preference. Recommendations become better with time due to data accumulation. Google's *Nest* thermostat "learns" from your temperature preferences, activity schedule, and changes you make to the thermostat settings to adjust temperature automatically in order to save energy and keep your smart home comfortable.

Google Translate can translate words, phrases, sentences and even paragraphs from one language to another, among around 100 languages. It can even

make pronunciations. You can use it from any Web browser for free. Even though some translations are inaccurate, mistaken, or even ridiculous, the tool can be very useful when dealing with foreign languages or as an aid for translation related work. It also gathers user feedback and corrections to improve its database.

These are just a few examples of many such AI programs.

Learning and Data Mining

Even though the path to AI is still long and uncertain, many useful programming techniques have been devised along the way that greatly enhance the capabilities and smartness of programs.

One such technique is *learning*—An application is so organized that it can receive training, automatically customize itself, or make improvements by analyzing its past successes and failures. The AI term "learning" is used even though such program behaviors are a far cry from human or even animal learning. In other words, learning in AI programs refers to "making self improvements in performing predefined tasks through experience". It has little to do with the general ability of learning in creatures.

Another technique is *data mining*—An application will gather or be fed large amounts of data, such as "who visits what websites" or "scenes and situations encountered by an automobile driver". Statistical analyses are performed on the data to find insights or hidden correlations. Results mined from the mountains of data can often be surprising or revealing.

An interesting story can illustrate the basic idea. In 1992, a well-known retailer had its sales data analyzed and a correlation between beer and diapers was found, not by sophisticated data mining but just simple data queries. This means placing beer and diapers together on shelves can increase sales of both items.

Extracting useful information from *big data*, such as all phone calls intercepted in thirty days, has been a well-known technique for intelligence gathering by security agencies. Today, collecting and quickly processing large amounts of data to provide insight or gain effectiveness has become a major trend in AI.

These are just two examples of many AI programming techniques in use today. AI is being applied in all kinds of areas including marketing, manufacturing, banking/finance, agriculture, health care, image processing, security, self-driving vehicles, and many more.

Changing Context

The meaning of the term AI becomes difficult to pinpoint because it has a very complicated philosophical context that has never been generally agreed upon and that has been shifting with time and technology. It is almost as if when an AI program has clearly succeed in a task, such as winning the Go game, we see how it's done and say "that's not intelligent, its just following rules this way or that way." It is against this shifting context we interpret the success/failure of AI projects.

In *Computational Thinking* (CT) an important concept is *context* because in computing the ability to switch context is paramount. For example, when a computer pauses one program to resume another the CPU undergoes a *context switch* to restore the saved state of the target program in order to resume running it. This happens all the time when we use computers, moving from editing a document to surfing the Web for example.

Even within the same program, the same bit pattern is subject to different interpretations depending on its *data type*, i.e. its context.

```
'0'  00110000 (48)          '9'  00111001 (57)
'A'  01000001 (65)          'Z'  01011010 (90)
'a'  01100001 (97)          'z'  01111010 (122)
```

In the preceding table, you see that the bit pattern representing 65 as an *integer type* will represent the letter 'A' as a *character type*.

Thus, we have the CT rule:

"**The meaning of any symbol depends on its context.**"

We all know taking words out of their original context is a common problem. The verb "**computize**" means putting CT to use. For example, we should computize and prevent context related pitfalls by avoiding words such as "last night" and "next week" in our emails.

Applying this CT principle, we naturally realize that the meaning of the symbol AI shifts along with its context, ranging from the lofty "man-made intelligence" to the merely "added smarts in a program".

Finally

Many somewhat clever computer systems are promoted and marketed under the AI label. AI vacuum cleaners are a vivid example of such promotion or hype. The exaggeration of AI technologies has negative effects. According to a *Bloomberg Opinion* article, "Executives who say they've adopted the technology may be responding to hype and the fear of missing out." Politicians, funding agencies, and business managers may be similarly influenced.

Still, many AI systems do offer value, smartness, and even amazing capabilities. Those based on big data, such as Google Translate or face recognition systems, are in vogue. Limitations of these AI techniques include accuracy and bias in the data used as well as insufficient background knowledge and expertise in the target areas.

Computer systems can be very impressive and powerful without relation to AI. Examples are everywhere: the Internet, the Web, MS Word, LaTeX, MAXIMA, Skype, Zoom, turn-by-turn navigation, encryption/decryption, and so on.

The original goal of AI, that of creating man-made thinking, learning, intelligent and even conscious machines, is difficult to achieve. There is yet no generally agreed concrete criteria to replace the Turing Test for AI.

Will AI systems lead to intelligence? Who knows. Yet, who cares? If a program is sufficiently smart, useful in many applications, why not label it AI? It may be even more attractive in the marketplace. By aiming for intelligence, the industry may have hit a pot of gold, that is not such a disaster!

As to AI wiping out humans, not any time soon.