
Chapter 1

<i>Title</i>	General Introduction
<i>Subject</i>	Principals and Approaches Of Artificial Intelligence
<i>Author</i>	Chourouk Guettas
<i>Grade level</i>	Master 1
<i>Objective</i>	Learn the basics and the approaches of AI
<i>Materials</i>	Book; Artificial Intelligence: A modern Approach
<i>Activities and procedures</i>	Presentations and exercises
<i>Lecture's Plan</i>	Chapter 1 General Introduction 1- What's AI a. The cognitive modelling Approach : Thinking Humanly b. The logistic Approach : laws of thoughts 2- The Foundations of AI 3- History of AI 4- AI Today 5- Applications of Artificial Intelligence

We call ourselves Homo-sapiens -man of wisdom- because our mental capacities are so important to us. For thousands of years, we have tried to understand *how we think*; that is, how a mere of handful stuff can perceive, understand, predict, and manipulate a world far larger and more complicated than itself. The field of artificial intelligence, or AI, goes further still: it attempts not just to understand but also to build intelligent entities.

1.1 What is AI?

Since the beginning of AI, there were a lot of attempts to find a unified definition for Artificial Intelligence that all went to vain. It's hard to claim that one definition is the right one but according to different perspectives, here are eight definitions of what can be called AI:

Systems that think like humans	Systems that think rationally
<p>“The exciting new effort to make computers think ... <i>machines with mind</i>, in the full and literal sense.” (Haugeland, 1985)</p> <p>“[the automation of] activities that we associate with human thinking, activities such as decision-making, problem-solving, learning ...” (Bellman 1978)</p>	<p>“The study of mental faculties through the use of computational model.” (Charniak and McDermott, 1985)</p> <p>“The study of the computational that make it possible to perceive, reason, and act” (Winston, 1992)</p>
Systems that act like humans	Systems that act rationally
<p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<p>“Computational intelligence is the study of the design of intelligent agents” (Poole et al. 1998)</p> <p>“AI ... is concerned with intelligent behavior in artifacts.” (Nilsen, 1998)</p>
Figure 1.1 Some definitions of artificial intelligence.	

Historically, all four approaches to AI have been followed. A human-centered approach must be an empirical science, involving hypothesis and experimental confirmation. A rationalist approach involves a combination of mathematics and engineering. Let's take a look to each one of these approaches.

- **Acting humanly: The Turing Test approach**

The **Turing Test**, proposed by Alan Turing (1950), was designed to provide an operational definition of intelligence. Rather than a list of qualifications, a test was proposed to decide whether a machine is undeniably intelligence entity. The computer passes the test interrogator, after posing some written questions, cannot tell whether the written responses come from a person or not. But for a computer to pass the test, it needs to possess the following capabilities:

- **Natural language processing:** to enable it to communicate successfully.
- **Knowledge representation:** to store what it knows or hears.
- **Automated reasoning:** to use the stored information to answer questions and to draw new conclusions.
- **Machine learning:** to adapt to new circumstances and to detect and extrapolate patterns.

Turing's test deliberately avoided direct physical interaction between the interrogator and the computer, because physical simulation of a person is unnecessary for intelligence. However the so-called *total Turing test* includes a video signal so that the interrogator can test the subject's perceptual abilities, and he can pass physical objects "through the hatch". To achieve this the computer will need also:

- **Computer vision:** to perceive objects, and
- **Robotics:** to manipulate objects and move about.

These six disciplines compose most of AI, and Turing deserves credit for designing a test that remains relevant for over 70 years later. Yet it was important to surpass Turing test because it is more interesting to understand intelligence rather than just duplicate an exemplar. Exactly like the quest for "artificial flight" which only succeeded when the Wright brothers stopped imitating birds and used aeronautical engineering. As the goal of their field wasn't making "machines that fly exactly like pigeons that they can fool even other pigeons".

- **Thinking humanly: The cognitive modeling approach**

If you are going to say that a given program thinks like a human, we must have some way of determining how humans think. There are two ways to know how the human mind work actually; through introspection (catching our own thoughts) and through psychological experiments. Once we have a sufficiently precise theory of the mind, it becomes possible to express the theory as a computer program.

The interdisciplinary field of **cognitive science** brings together computer models from AI and experimental techniques from psychology to try to construct precise and testable theories of the workings of the human mind.

In the early days of AI there was often confusion between the approaches: an author would argue that an algorithm performs well on a task and that it is therefore a good model of human performance, or vice versa. Modern authors separate the two kinds of claims; this distinction has allowed both AI and cognitive science to develop more rapidly. The two fields continue to lean on each other, especially in the areas of vision and natural language.

- **Thinking rationally: The “laws of thought” approach**

Aristotle was one of the first to attempt to codify “right thinking”, that is, irrefutable reasoning processes. His *syllogism* provided patterns for argument structures that always yielded correct conclusions when given correct premises. His study initiated the field called *logic*.

Logicians in the 19th century developed a precise notation for statements about all kinds of things in the world and about the relations among them. By 1965, programs existed that could, in principle, solve any solvable problems described in logical notation.

There are two main obstacles to this approach. First, it is not easy to take informal knowledge and state it in the formal terms required by logical notations. Second, there is a big difference between being able to solve a problem “in principle” and doing so in practice. Even problems with just a few dozen facts can exhaust the computational resources of any computer. Although both of these obstacles apply to any attempt to build computational reasoning systems, they appeared first in the logicist tradition.

Acting rationally: The rational agent approach

An **agent** is just something that acts (agent comes from the Latin *agere*, to do). But computer agents are expected to have other attributes that distinguish them from mere “program”, such as operating under autonomous control, perceiving their environment, persisting over a prolonged time period, adapting to change, and being capable of taking on another’s goals. A **rational agent** is one that acts so as to achieve the **best outcome** or, when there is uncertainty, the **best expected outcome**.

In the “law of thought” approach to AI, the emphasis was on correct inferences. Making correct inferences in order to achieve one’s goals is sometimes part of being a rational agent. On the other hand, there are often situations where there is no provably correct thing to do, yet something must still be done. A reflex action may also be considered a successful rational action without involving inference.

The study of rational agent design has at least two advantages for the AI discipline. First, it’s more general than the “laws of thought” approach, because correct inference is just one of several possible mechanisms for achieving rationality. Second, it is more amenable to scientific development because human behaviour is more than rational thoughts. Human behaviour, being well-adapted for one specific environment, it needs

the skills to represent knowledge, generate comprehensive sentences in natural language, learning and visual perception to have a better idea of how the world works.

1.2 Foundations of AI

In this section, a brief history of the disciplines that contribute ideas and techniques to AI will be presented in chronological order:

Philosophy (428 B.C. ~ present)

- Can formal rules be used to draw valid conclusions?
- How does the mental mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?

Aristotle was a pioneer in formulating a law of thoughts to govern the rational thinking. Around 1500 Leonardo da Vinci designed but didn't build a mechanical calculator that proved later to be functional.

René Descartes (1596 – 1650) gave the first clear discussion between mind and matter and the free will. He was also a proponent of dualism discussing the soul/spirit part of human.

The next problem is to establish the source of knowledge. The empiricism movement started with the principle of induction, confirmation theory (to understand how knowledge can be acquired from experience) and by 1928 an explicit computational procedure for extracting knowledge from elementary experiences.

The final element is the connection between knowledge and action, which is vital to AI. Intelligence requires actions as well as reasoning and only by understanding how actions are justified can we understand how to build an agent whose actions are justifiable (or rational).

Mathematics (c. 800 ~ present)

- What are the formal rules to draw valid conclusion?
- What can be computed?
- How do we reason with uncertain information?

Economics (1776 ~ present)

- How should we make decisions so as to maximize payoff?
- How should we do when others may not go along?
- How should we do this when the payoff may be far in the future?

Neuroscience (1861 ~present)

-
- How do brains process information?

Psychology (1897 ~present)

- How do humans and animals think and act?

Computer engineering (1940 ~ present)

- How can we build an efficient computer?

Control Theory and Cybernetics (1948 ~present)

- How can artifacts operate under their own control?

Linguistics (1957 ~ present)

- How does language relate to thoughts?

1.3 The History of Artificial Intelligence

▪ The gestation of AI (1943 ~ 1955)

The first work that is recognized now as AI was done by Warren McCulloch and Walter Pitts (1943). Based on physiology and functions of neurons, propositional logic and Turing's theory of computation; they proposed a model of artificial neurons in which each neuron is characterized as being "on" or "off". They showed that any computable function could be computed by some network of connected neurons while implementing all the logical connectives (and, or, not, etc.).

Donald Hebb introduced Hebbian learning rule in 1949 demonstrating the learning ability in neural networks but changing the connection strength. Marvin Minsky and Dean Edmond, students at Harvard, built the first neural network computer in 1950. But it was Alan Turing who first articulated a complete vision of AI in his 1950 article "Computing Machinery and Intelligence". Therein he introduced the Turing test, machine learning, genetic algorithms, and reinforcement learning.

▪ The birth of Artificial Intelligence (1956)

After graduating from Princeton, John McCarthy moved to Dartmouth College where the birth of the field officially announced. McCarthy convinced Minsky, Claude Shannon, and Nathaniel Rochester to help him bring together U.S researchers interested in automata theory, neural nets, and the study of Intelligence. They organized a two-month workshop at Dartmouth in the summer of 1956. There were 10 attendees in all (MIT, Stanford, CMU, and IBM), and although the workshop didn't lead to any breakthroughs, it did introduce all the major figures who will dominate the field for the next 20 years and an agreement to adopt McCarthy's new name for the field "*Artificial Intelligence*".

▪ Early enthusiasm, great expectations (1956 ~ 1969)

The early years of AI were full of successes, in a limited way given the primitive computers and programming tools of the times. It was astonishing whenever computer did anything remotely clever.

- *General Problem Solver*: Program designed to imitate human problem-solving protocols..
- *Geometry Theorem Prover*: Prove mathematical theorem
- *Lisp*: By McCarthy, AI programming language, and later on *Advice Taker*, to solve general knowledge problems of the world, the program was able to accept new axioms therefore allowing it to achieve new competence in new areas and by that embodying the principals of knowledge representation and reasoning.
- *Analogy program*: solved the geometric analogy problems that appear in IQ test.
- *Blocks world*: solid blocks placed on a tabletop, and a typical task was to rearrange the blocks in a certain way using a robot hand. (Project vision included: learning theory, natural language processing, and the planner).
- *Perceptrons*: neural networks and learning algorithms to adjust the connection strengths to match any input data.

▪ A dose of reality (1966 ~ 1973)

The visible future where Herbert expected computer to become chess champion and prove significant mathematical theorem didn't come only after 40 years along with a lot of failures in a huge selection of real problems.

- *Syntactic manipulation failure*: The first attempt to speed up translate from Russian to English (Sputnik launch). The famous retranslation of “the spirit is willing but the flesh is weak” as “the vodka is good but the meat is rotten” illustrates the difficulties encountered. After that all the US Academic funding for translation projects were all cancelled.
- The failure to solve the problem when scaling up and countless of useless hours of calculations in machine evolution.
- Some fundamental limitations on the basic structures being used to generate intelligent behaviour; one perceptron could not be trained to recognize when its two inputs were different.

▪ Knowledge-based systems: The key to power? (1969–1979)

General-purpose search methods that string together elementary reasoning steps to find complete solutions were labelled as **weak method** because, although general, they do not scale up to large or difficult problem instances. The alternative is to use more powerful, domain-specific knowledge that allows larger reasoning steps and can handle occurring cases in narrow areas of expertise.

- *Expert Systems*: DENDRAL was that it was the first successful *knowledge-intensive* system (molecular structure): its expertise derived from large numbers of special-purpose rules. After that MYCIN to diagnose blood infections was introduced with about 450 rules and incorporated a calculus of uncertainty called the **certainty factor** associated with medical knowledge.

-
- *Understanding natural language*: Winograd's SHRDLU system for understanding natural language was presented and it was able to overcome ambiguity and understand pronoun references.
 - *Frames*: The widespread growth of applications to real-world problems caused a concurrent increase in the demands for workable knowledge representation schemes. A large number of different representation and reasoning languages were developed; Prolog language became popular in Europe, PLANNER family in the United States and FRAMES Minsky's idea of **frames** (1975), adopted a more structured approach, assembling facts about particular object and event types and arranging the types into a large taxonomic hierarchy analogous to a biological taxonomy.

- **AI becomes an industry (1980–present)**

The first successful commercial expert system, R1, began operation at the Digital Equipment Corporation (McDermott, 1982). The program helped configure orders for new computer systems; by 1986, it was saving the company an estimated \$40 million a year. By 1988, DEC's AI group had 40 expert systems deployed, with more on the way. DuPont had 100 in use and 500 in development, saving an estimated \$10 million a year. Nearly every major U.S. Corporation had its own AI group and was either using or investigating expert systems.

Overall, the AI industry boomed from a few million dollars in 1980 to billions of dollars in 1988, including hundreds of companies building expert systems, vision systems, robots, and software and hardware specialized for these purposes. Soon after that came a period called the *“AI Winter,”* in which many companies fell by the wayside as they failed to deliver on extravagant promises.

- **The return of neural networks (1986–present)**

In the mid-1980s at least four different groups reinvented the **back-propagation** learning algorithm first found in 1969 by Bryson and Ho. The algorithm was applied to many learning problems in computer science and psychology. As occurred with the separation of AI and cognitive science, modern neural network research has bifurcated into two fields, one concerned with creating effective network architectures and algorithms and understanding their mathematical properties, the other concerned with careful modelling of the empirical properties of actual neurons and ensembles of neurons.

- **AI adopts the scientific method (1987–present)**

Recent years have seen a revolution in both the content and the methodology of work in artificial intelligence. It is now more common to build on existing theories than to propose brand-new ones, to base claims on rigorous theorems or hard experimental evidence rather than on intuition, and to show relevance to real-world applications rather than toy examples. AI was founded in part as a rebellion against the limitations of existing fields like control theory and statistics, but now it is embracing those fields.

- *Hidden Markov models (HMMs)*: provides a mathematical framework for understanding and support the engineering to solve problems of speech and character recognition.

-
- *Data Mining*: Using improved methodology and theoretical frameworks, the neural network field arrived at an understanding in which neural nets can now be compared with corresponding techniques from statistics, pattern recognition, and machine learning.
 - *Bayesian networks*: was invented to allow efficient representation of, and rigorous reasoning with, uncertain knowledge. This approach largely overcomes many problems of the probabilistic reasoning systems of the 1960s and 1970s. The approach allows for learning from experience, and it combines the best of classical AI and neural nets.

- **The emergence of intelligent agents (1995–present)**

Researchers started to look at the “whole agent” problem again and one of the most important environments for intelligent agents is the Internet. AI systems have become so common in Web-based applications that the “-bot” suffix has entered everyday language. It was then widely appreciated that sensory systems (vision, sonar, speech recognition, etc.) cannot deliver perfectly reliable information about the environment. Hence, reasoning and planning systems must be able to handle uncertainty.

- *human-level AI*: Despite these successes, some influential founders of AI believe AI should return to its roots of striving for, in Simon’s words, “machines that think, that learn and that create.” Instead of creating ever-improved versions of applications that are good at a specific task.
- *Artificial General Intelligence*: held its first conference and organized the Journal of Artificial General Intelligence in 2008. AGI looks for a universal algorithm for learning and acting in any environment.

- **The availability of very large data sets (2001–present)**

Throughout the 60-year history of computer science, the emphasis has been on the *algorithm* as the main subject of study. But some recent work in AI suggests that for many problems, it makes more sense to worry about the *data* and be less picky about what algorithm to apply.

Word-sense disambiguation and filling in holes in a photograph were ones of the first attempts to solve problems with huge amount of data instead of algorithms. Works like these suggest that the “*knowledge bottleneck*” in AI—the problem of how to express all the knowledge that a system needs—may be solved in many applications by learning methods rather than hand-coded knowledge engineering, and provide the learning algorithms with enough data to go on.

1.4 AI Today

What can AI do today? A concise answer is difficult because there are so many activities in so many subfields. Here we sample a few applications;

- *Robotic vehicles and self-driving cars*: safely driving in traffic through the streets, obeying traffic rules and avoiding pedestrians and other vehicles.
- *Speech recognition*: A traveller calling United Airlines to book a flight can have the entire conversation guided by an automated speech recognition and dialog management system.

-
- *Autonomous planning and scheduling:* Autonomous planning program to control the scheduling of operations for a spacecraft; REMOTE AGENT generated plans from high-level goals specified from the ground and monitored the execution of those plans—detecting, diagnosing, and recovering from problems as they occurred.
 - *Game playing:* IBM's DEEP BLUE became the first computer program to defeat the world champion in a chess match when it bested Garry Kasparov.
 - *Spam fighting:* Each day, learning algorithms classify over a billion messages as spam, saving the recipient from having to waste time deleting messages.
 - *Logistics planning:* During the Persian Gulf crisis of 1991, U.S. forces deployed a Dynamic Analysis and Replanning Tool, DART (Cross and Walker, 1994), to do automated logistics planning and scheduling for transportation. This involved up to 50,000 vehicles, cargo, and people at a time, and had to account for starting points, destinations, routes, and conflict resolution among all parameters. The AI planning techniques generated in hours a plan that would have taken weeks with older methods.
 - *Robotics:* The iRobot Corporation has sold over two million Roomba robotic vacuum cleaners for home use. The company also deploys the more rugged PackBot to Iraq and Afghanistan, where it is used to handle hazardous materials, clear explosives, and identify the location of snipers.
 - *Machine Translation:* computer program automatically translates from Arabic to English, allowing an English speaker to see the headline “Ardogan Confirms That Turkey Would Not Accept Any Pressure, Urging Them to Recognize Cyprus.” The program uses a statistical model built from examples of Arabic-to-English translations and from examples of English text totalling two trillion words (Brants et al., 2007). None of the computer scientists on the team speak Arabic, but they do understand statistics and machine learning algorithms.