

Artificial Intelligence

academic year 2024/2025

Giorgio Fumera

Pattern Recognition and Applications Lab

Department of Electrical and Electronic Engineering

University of Cagliari (Italy)



Introduction to Artificial Intelligence

historical notes and current research landscape

What is “intelligence”?


Broad definition: a set of capabilities that allow humans to learn, think, understand, communicate, be self-conscious, build abstract models of the world, plan, adapt to novel external conditions, etc. (some of these capabilities are exhibited also by animals, e.g., associative memory, reacting to stimuli, communicating).

Different aspects of intelligence have been studied since a long time by several disciplines: **logic**, **psychology**, **neurophysiology**, etc.

“Intelligent” machines

The idea of building “intelligent” machines has also been envisaged in the past, e.g.:

- ▶ Leonardo da Vinci’s robot knight (about 1495)
- ▶ “Mechanical Turk” automaton chess player (late 18th cent.)
- ▶ science fiction, e.g.:
 - the HAL 9000 computer of A.C. Clarke’s novel and S. Kubrick’s “2001: A Space Odyssey” movie (1968)
 - the “replicants” of R. Scott’s “Blade runner” movie (1982)



“Artificial intelligence”

The field named **Artificial Intelligence** (AI) was born in the 1950s from the confluence of two broad and articulated earlier research efforts:

- ▶ **understanding** human intelligence
- ▶ **building machines** capable of autonomously performing complex tasks that are deemed to require “intelligence”

Early investigations on human intelligence

Goal: **understanding** different aspects of **human** intelligence

▶ **high-level** manifestations:

– **rationality**

logic: Aristotle (4th cent. BC), G.W. Leibniz (17th-18th cent.), G. Boole (19th cent.), etc.

– **behavior** and **mind**

psychology and **cognitive science** (since 19th cent.)

▶ **low-level** biological support: the **brain**

neuroanatomy and neurophysiology (since 19th cent.):

McCulloch and Pitt's model of neuron (1943), D.O. Hebb's theory on neurons as basic units of thought, etc.

Early investigations on technology



Goal: **building machines** capable of autonomously performing complex tasks:

- ▶ **automata** (e.g., jacquard loom, 1804) and **cybernetics** (feedback and control):
N. Wiener, W.R. Ashby (1940s)
- ▶ **statistics** and **probability theory** as tools to deal with uncertainty in reasoning and decision-making: T. Bayes (18th cent.);
K. Pearson, R.A. Fisher, A. Wald,
J. Neyman (late 19th – 20th cent.)



Early investigations on technology: the *computer*

- ▶ Precursors: **mechanical** devices: B. Pascal and G.W. Leibniz (17th cent.), C. Babbage's analytical engine (19th cent.)



- ▶ Contributions from **engineering**: electromechanical and **electronic** devices (1940s): K. Zuse, J.P. Eckert, J.W. Mauchly, J. von Neumann
- ▶ Contributions from **logic** and **mathematics**: computational theory, the foundation of computer science: A.M. Turing, A. Church (1930s)

Early investigations on “thinking” computers

Alan M. Turing (1912-1954):

- ▶ first investigations into the nature of computing
- ▶ the “logical computing machine” (**Turing machine**): a universal computer
- ▶ envisioning intelligent computers:
Computing Machinery and Intelligence,
Mind, Vol. LIX, No. 263, 433–460, 1950;
operational definition of intelligence:
the **Turing Test**



Are electronic computers the right tool for building “intelligent” machines?

The birth of AI

In the 1950s many researchers from different disciplines were **investigating human intelligence**, and others were **building machines** capable of performing complex tasks.

In the summer of 1956 some of them met at a workshop at the Dartmouth College (USA), and founded a new discipline named “**artificial intelligence**”, whose aim was **to build “intelligent” machines**.

The founders: J. McCarthy, M. Minsky, A. Newell, H. Simon, C. Shannon, O. Selfridge, R. Solomonoff, and others.



AI early explorations: 1950s and 1960s

Goals: identifying **specific tasks** that require intelligence, and figuring out **how to get machines to do them**.

Great interest in mimicking **high-level** human thought and mental abilities, e.g.:

- ▶ logical reasoning and problem solving, game playing
- ▶ natural language processing (understanding, translation)

Some investigations also on **low-level** abilities:

- ▶ speech recognition
- ▶ vision (reading cursive script, object recognition, etc.)

Main problem: how do **humans** do that?

AI early explorations: 1950s and 1960s

Starting point: **games** and **toy problems** (easy to formalize and investigate), and some real-world ones

- ▶ game playing: 15-puzzle, checkers, chess, etc.
- ▶ theorem proving
- ▶ natural language processing (NLP)
- ▶ recognizing objects in images

AI early explorations: 1950s and 1960s

Dominant viewpoint: the essence of intelligence is deemed to be **symbol processing**.

Early AI research focused therefore on a **symbolic** approach, aimed at simulating **high-level** manifestations of human intelligence.

Main tools:


- heuristic search
- syntax analysis/generation
- symbolic knowledge representation (symbols, lists, graphs)
- symbolic knowledge processing: new programming languages (LISP, etc.)

Heuristic search methods


Symbol processing approach, applied to problems like:

- ▶ game playing: 15-puzzle, checkers, chess (the “Drosophila of AI”)

13	10	11	6
5	7	4	8
1	12	14	9
3	15	2	



- ▶ geometric analogy problems




- ▶ theorem proving
- ▶ mechanizing problem solving: A. Newell and H. Simon's *General Problem Solver* (1959)

Heuristic search methods

Common approach:

- ▶ knowledge representation: **lists of symbols**
(main feature of the LISP language: 1958, J. McCarthy)
- ▶ search methods: **search tree, heuristics**;
an example: search tree for the 8-puzzle problem



Natural language processing

Aim: understanding, generating and translating natural language.

A difficult problem, due also to different **linguistic levels**:



- ▶ **morphology**: word parts (e.g.: walking = walk + -ing)
- ▶ **syntax** (grammar): rules that define well-formed sentences (e.g.: *John hit the ball*: Yes; *ball the hit John*: No)
- ▶ **semantics**: meaning of a sentence
- ▶ **pragmatics**: context and background knowledge, e.g.:
*John went to the **bank***
*John threw the ball to the window and broke **it***
*John threw the glass to the wall and broke **it***

Natural language processing

Focus of earlier research: **syntactic** level (**symbol processing**).

Seminal work: N. Chomsky, *Syntactic Structures*, 1957.

Grammar definition: syntax rules for analyzing/generating sentences; main tool: **parse tree**.



Applications:

- ▶ question answering (original goal: computer interfaces)
- ▶ machine translation: early optimism, but it turned out to be a very difficult task

Non-symbolic approach

A secondary (by then) approach was a **non-symbolic** one, aimed at simulating **low-level** manifestations/capabilities of intelligence like perception (mainly, visual perception).

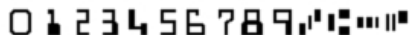
This approach gave rise to

- ▶ the **pattern recognition** discipline, which later emerged as a relevant branch of AI
- ▶ **artificial neural networks**, that became one of the main AI tools (now re-flourished as *deep learning*)

Pattern recognition

Goal: classifying different kinds of **signals** (images, sounds, electronic signals, etc.) into one of several categories

- ▶ first problem addressed: **image classification**
- ▶ first application: optical character recognition (OCR)



0 1 2 3 4 5 6 7 8 9 , : ; |

Main approaches:

- ▶ **template matching**
- ▶ **learning**: image pre-processing (noise filtering, line thickening, edge enhancement, ...), **feature** extraction (e.g., shape), classification “rules” **learnt from examples**

Artificial neural networks (ANNs)

Non-symbolic (**low-level**), **connectionist** approach.

The origins:

- ▶ McCulloch and Pitts' mathematical model of neuron (1943)
- ▶ the **perceptron** by F. Rosenblatt (1957): a potential model of human learning, cognition and memory
 - network of McCulloch-Pitts' neural elements
 - **learning** algorithm for adjusting connection weights **from examples**

First application: pattern recognition (mainly from images)

- ▶ OCR
- ▶ aerial images

Great expansion: mid-1960s to early 1980s

From toy/lab problems to real-world and commercial applications:

- ▶ computer vision
- ▶ mobile robots
- ▶ game playing
- ▶ speech recognition, NLP
- ▶ knowledge representation and reasoning



Relevant public funding:

- ▶ DARPA's Strategic Computing Program (USA)
- ▶ Fifth Generation Computer Systems (Japan)
- ▶ ESPRIT (Europe)

Great expansion: mid-1960s to early 1980s

Computer vision:

- ▶ MIT *Summer Vision project* (1966)
- ▶ low-level, hierarchical image processing (hints from biology); image filters; line, corner, surface detection; 3D reconstruction
- ▶ early application: guiding a robot arm to manipulate blocks





- ▶ high-level vision: finding objects in scenes (templates, parts)
- ▶ two main approaches emerge:
 - whole scene reconstruction: **difficult**
 - perceiving to guide robot action (purposive vision): **easier**

Great expansion: mid-1960s to early 1980s

Mobile robots:

- ▶ sensors, actuators, computer vision, environment modeling, planning
- ▶ route finding: heuristic search, **A* algorithm**
- ▶ first autonomous vehicles



Great expansion: mid-1960s to early 1980s


Game playing: progress in **chess**

- ▶ programs that attain human-level (not master) capability
- ▶ investigations into human ability: accumulated knowledge vs **massive search**

Great expansion: mid-1960s to early 1980s

Progress in NLP, on **less ambitious** goals than in the 1950s:

- ▶ improvements in grammars
- ▶ machine translation with humans in the loop
- ▶ speech recognition (**easy**), and understanding (**difficult**)
- ▶ **declarative** (logical languages, inference algorithms) vs **procedural** (“hard-wired”) knowledge
- ▶ dialog systems (1971, T. Winograd’s SHRDLU: blocks’ world, **procedural** knowledge):



Great expansion: mid-1960s to early 1980s

Knowledge representation and reasoning:

- ▶ development of consulting systems, decision support systems, **expert systems**;
main idea: solving **domain-specific** problems by embedding **expert** knowledge in the form of IF-THEN rules
- ▶ applications: chemistry, medical diagnosis, geology, military;
since the 1990s: business

Great expansion: mid–1960s to early 1980s

Until the 1970s AI research is mainly based on the **symbol processing** conception of human intelligence

- ▶ main approach: mimicking **high-level** human abilities through heuristic search and symbolic processing (“good, old-fashioned AI”, GOFAI)
- ▶ many successful applications through a pragmatic approach in **specific** tasks. . .
- ▶ . . . but very limited achievements with respect to early expectations for a **general** AI

Mid 1980s: the “AI winter”

Real-world tasks turned out to require much more “intelligence” than that achievable by heuristic search and symbolic processing (GOFAI). Two main issues emerge:

- ▶ **computational complexity**: **combinatorial** explosion
- ▶ human problem-solving relies on a large body of implicit **background knowledge** (including common sense)

The non-symbolic, connectionist approach (artificial neural networks) exhibits limitations as well.

Main consequences (“**AI winter**”):

- ▶ drop of interest in AI
- ▶ scaling back AI’s goals
- ▶ reduction of research funding

Mid-1980s to 1990s: technical and theoretical advances

The AI winter was overcome thanks to new results in several fields, based on solid theoretical foundations from:

- ▶ mathematics
- ▶ statistics and probability theory
- ▶ control engineering

This enabled concrete progress in real-world tasks, albeit still **far** from early expectations:


- ▶ knowledge representation and reasoning
- ▶ machine learning
- ▶ computer vision
- ▶ Intelligent Agent architectures

Mid-1980s to 1990s: technical and theoretical advances

Advances in **search algorithms**: evolutionary approach, genetic algorithms (inspired by evolution).

Advances in **knowledge representation and reasoning**:

- ▶ new paradigms, e.g.; fuzzy logic, soft computing (inspired by human mind)
- ▶ semantic networks, ontologies; e.g.:
 - WordNet, <https://wordnet.princeton.edu>
 - BabelNet, <https://babelnet.org>
- ▶ probabilistic reasoning to overcome the limits of logic (probabilistic graphical models, Bayesian networks, learning)



Mid-1980s to 1990s: technical and theoretical advances

The rise of **machine learning**:

- ▶ availability of huge **amount** of data in digital form, often **manually annotated** by users (e.g., user preferences, text translated in different languages by humans)
- ▶ main idea: automatically inferring knowledge (patterns, rules, etc.) **from data** instead of eliciting it from domain experts
- ▶ data analysis methods: data mining, etc.
- ▶ theoretical foundations: statistics
- ▶ novel techniques: inductive logic programming, decision trees, resurgence of ANNs (1986: back-propagation algorithm), support vector machines, ensemble methods, etc.
- ▶ many application fields: computer vision, natural language processing, etc.

Mid-1980s to 1990s: technical and theoretical advances

Computer vision:


- ▶ two main approaches persist: scene analysis, purposive vision
- ▶ main achievements: surface, depth; tracking, object recognition
- ▶ fruitful exchanges with research on animal/human vision
- ▶ novel techniques: hierarchical models, ANNs, **deep neural networks**
- ▶ extensive application of **machine learning** techniques

Mid-1980s to 1990s: technical and theoretical advances

Intelligent Agent architectures:

- ▶ sensor networks
- ▶ autonomous, cooperating robots; emergent behavior
- ▶ the **intelligent agent** paradigm

An example: soccer-playing robots



Mid 1990s – today: main achievements

The **original** goal of **building “intelligent” machines** is still **far-reaching**.

Nevertheless

- ▶ several real-world problems can be successfully addressed
- ▶ many commercial applications have been developed
- ▶ many start-up companies are exploiting AI techniques

Mid 1990s – today: main achievements

Some examples:

- ▶ **games**: master level has been achieved in checkers, chess and (very recently) Go
- ▶ **computer vision**: object recognition, scene understanding, etc.
- ▶ **autonomous vehicles**: driverless cars, space vehicles, etc.
- ▶ automatic **language translation**
- ▶ **pervasive** applications: home automation, route finding in maps (search algorithms), recommender systems (machine learning, social/collaborative filtering), characters in video games, etc.
- ▶ **medicine**: diagnosis, etc.
- ▶ business rule management systems
- ▶ automated (high-frequency) trading

Summary of the main approaches to AI

The approaches pursued so far to build intelligent machines can be categorized along two main dimensions:

	Human performance	Rationality
Mind/ thinking	Systems that think like humans (cognitive modeling approach)	Systems that think rationally ("law of thought" approach)
Behavior	Systems that act like humans (Turing test approach)	Systems that act rationally (rational agent approach)

The **rational agent** approach is the most general one, and is amenable to scientific/technological development, although it may be not useful enough for understanding **human** intelligence.

Philosophical issues

A long-standing question: **Can machines be “intelligent”?**

Two main hypotheses:

- ▶ **Weak AI**: machines can **emulate** intelligence (act intelligently)
- ▶ **Strong AI**: machines can **be** intelligent
(if they act intelligently, they are intelligent, e.g.: Turing test)

Another long-standing question: **Is human mind a machine?**

Philosophical issues

Some of the arguments raised against the **Weak AI** hypothesis:

- ▶ **machines can never do X** (X = make mistakes, learn from experience, have a sense of humor, enjoy ice creams, etc.)
- ▶ **machines are formal systems**, and formal systems cannot establish the truth of every mathematical sentence (Gödel incompleteness theorem), whereas humans (in principle) can
- ▶ **human behavior can not be captured by a set of rules**

Alan Turing's viewpoint (Mind, 1950):

The question “Can a machine **think?**” is **ill-posed**.

Consider this one: “Can a machine **fly/swim?**”

- ▶ airplanes “fly”, but not as birds “fly”
- ▶ ships “swim” in Russian, but not in English or in Italian...

Philosophical issues

Some arguments against the **Strong AI** hypothesis:

- ▶ even if machines can **emulate** intelligence, they cannot be **self-conscious**
- ▶ relationship between mental states and body (brain) states (free will, consciousness, intentions): **dualism** (R. Descartes, 17th cent.) vs **materialism** (“brains cause mind”)
- ▶ a machine running the “right” program (e.g., for natural language understanding) does not necessarily have a mind (the Chinese room thought experiment, J. Searle, 1980)
- ▶ intelligence is an **emerging** behavior that can be only supported by **biological brains**

Ethical issues

Some **ethical** issues against AI:

- ▶ even if we could build intelligent machines, **should we?**
- ▶ consequences on humans: loss of jobs, loss of the sense of being unique, end of human race, etc.
- ▶ accountability (e.g., driverless cars)
- ▶ ...

Ethical concerns are currently re-flourishing, as many believe that human-level AI is now in reach (e.g., this is the current focus of the Future of Life Institute, <https://futureoflife.org>).

A snapshot of current AI research landscape

Research topics:

- ▶ machine learning
- ▶ knowledge representation and reasoning
- ▶ reasoning under uncertainty or imprecision
- ▶ natural language understanding / translation
- ▶ multi-agent systems
- ▶ planning
- ▶ heuristic search
- ▶ robotics
- ▶ vision
- ▶ pattern recognition
- ▶ ...

A snapshot of current AI research landscape

Associations:

- ▶ Association for the Advancement of Artificial Intelligence (AAAI), <https://www.aaai.org>
- ▶ Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE)
<https://claire-ai.org/>
- ▶ Italian Association for Artificial Intelligence (AIxIA)
<https://aixia.it>

A snapshot of current AI research landscape

Conferences:

- ▶ International Joint Conference on Artificial Intelligence (IJCAI), <https://www.ijcai.org>
- ▶ The AAAI Conference on Artificial Intelligence
<https://aaai.org/Conferences/AAAI/aaai.php>

Scientific journals:

- ▶ Journal of Artificial Intelligence Research
<https://www.jair.org>
- ▶ Artificial Intelligence
<https://www.journals.elsevier.com/artificial-intelligence>

Some recent initiatives

Human Brain Project (co-funded by the EU)

<https://www.humanbrainproject.eu>

Overall goal: understanding the human brain and its diseases, and **emulating** its computational capabilities



RoboLaw (co-funded by the EU)

<http://www.robolaw.eu>

Regulating Emerging Robotic Technologies in Europe:
Robotics facing Law and Ethics



The Artificial Intelligence Act

<https://artificialintelligenceact.eu/>

A proposal for a European law on Artificial Intelligence
of the European Parliament and of the Council



Suggested reading

D.R. Hofstadter, *Gödel, Escher, Bach: an eternal golden braid*, Basic Books, 1979

S. Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, Profile Books Ltd, 2019

K. Crawford, *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*, Yale University Press, 2021

N. Cristianini, *The Shortcut – Why intelligent machines do not think like us*, CRC Press, 2023