

Artificial Intelligence: Evolution & Perspectives

Digital Culture

Andrea Omicini
andrea.omicini@unibo.it

Master in Finanza, Controllo e Auditing + Marketing and Sales / Bologna Business School

Academic Year 2018/2019




- 1 The Idea of AI
- 2 The Raise of AI
- 3 Main Themes in Early AI
- 4 The AI Winter
- 5 The Modern Era of AI
- 6 AI Today & Tomorrow



Disclaimer

- a number of different sources were used to prepare these slides—including of course more than two decades of research in the field by myself
- however, the organisation of the material is largely inspired by [Wooldridge, 2018], which represents a fresh reading on the subject, indeed




Next in Line...

- 1 The Idea of AI
- 2 The Raise of AI
- 3 Main Themes in Early AI
- 4 The AI Winter
- 5 The Modern Era of AI
- 6 AI Today & Tomorrow



The Idea of Intelligence

- understanding how intelligence works is a persistent issue for humans
- Aristotle's logics is the most outstanding example of that
- “understanding”, for humans, typically means to be able to model and reproduce
- building machines that can reproduce intelligence is a way to measure how much we do understand the way in which intelligence works



The Idea of Intelligent Machines

- humans as makers of machines exhibiting intelligent behaviours
- machines vs. humans—from The Turk to AlphaGo
- the basic question: what is intelligence, which we aim at capture and reproduce?



The Turing Test I

- after providing the fundamental theoretical machinery for computation (the Turing Machine [Turing, 1937]), Turing started exploring the limits of automatic computation: could it be a path towards intelligence of machines?
- the Turing Test [Turing, 1950]

"You are interacting via a computer keyboard and screen with something that is either another person or a computer program. The interaction is in the form of text – questions and answers. Your task is to determine whether the thing being interrogated is in fact a person. Now suppose, after some time, you cannot tell whether the thing is a person or program. Then, Turing argued, you should accept that the thing being interrogated has human-like intelligence." [Wooldridge, 2018]

The Turing Test II

- *behavioural* interpretation of intelligence—independent from how it works, dependent on what it can achieve
- indistinguishable from human behaviour
- it has limitations, it has been refined, it is anyway meaningful



Against AI I

- many scientists have argued against the chance of machine intelligence
- mostly philosophers, like Searle [Searle, 1980]

The Chinese Room

"Imagine a room in which a man, who understands no Chinese, receives, through a slot in the door, questions written in Chinese. When he receives a question, the man carefully follows detailed instructions written in English to generate a response to the question, which he passes back out through the slot. Now suppose the questions and responses are part of a Chinese Turing test, and the test is passed." [Wooldridge, 2018]



Against AI II

- interesting argument
- so many assumptions and mistakes in it
- which does not mean its conclusions are wrong, of course



Next in Line...

- 1 The Idea of AI
- 2 The Raise of AI**
- 3 Main Themes in Early AI
- 4 The AI Winter
- 5 The Modern Era of AI
- 6 AI Today & Tomorrow



AI is Born in 1956

The Golden Age of AI

- Dartmouth College, New Hampshire, Summer School, 1956
- John McCarthy invites all scholars interested in computing towards intelligence
- among those
 - Marvin Minsky, co-founder of AI Lab at MIT
 - Alan Newell, Herb Simon, authors of Logic Theorist (an automatic theorem prover), likely the first AI program
 - John McCarthy, inventor of LISP, the first programming language for AI
- the term “Artificial Intelligence” was actually coined there, to describe the overall new field of research
- the dualism between AI as *intelligent behaviour* and AI as *intelligent process* already there, from the beginning

General AI I

General purpose AI

- building general-purpose intelligence machines is the goal of **General AI**
- we do have a poor understanding of human intelligence, and of intelligence in general
- early AI focussed then on *intelligent components*



General AI II

Components of intelligence

- perception
- machine learning
- problem solving & planning
- reasoning
- natural language understanding

Perception

- understanding the environment
- through sensors of any sort
- interpreting the overall situation
- ! one of the most difficult task of AI

General AI III

Machine learning

- learning from data
- building models (e.g., classification)
- making prediction
- e.g., face recognition through training

Problem solving & planning

- devising a course of actions towards a *goal*
- based on a *repertoire of actions*
- e.g., playing games



General AI IV

Reasoning

- inferring new knowledge from available one
- in a consistent and robust way

Natural language understanding

- ability to understand human languages
- either spoken or written
- possibly engage in conversations with humans



Next in Line...

- 1 The Idea of AI
- 2 The Raise of AI
- 3 Main Themes in Early AI**
- 4 The AI Winter
- 5 The Modern Era of AI
- 6 AI Today & Tomorrow



Search

Searching for solution

- from the current state of the world to a desired one, called *goal*
- based on a repertoire of available *actions*
- modelled in terms of their
 - *pre-conditions*, or, the state of the world in which they can be applied
 - *effects*, or, the state of the world they bring about
- a *search tree* can be built
 - from the *initial state* as its root—*forward search*
 - from the *final state* as its root—*backward search*



The Problem of Search

Explosion of the space of solutions

- search trees become too large
- *combinatory explosion*
- e.g., game of chess
- possible approach: *heuristics* as rule of thumb to direct search and reduce computational complexity
- e.g., IBM's Deep Blue against Gary Kasparov in 1996, combining computational power with heuristics



Planning in the Block World

SHRDLU

- the Blocks World is a virtual environment where a number of coloured objects (blocks, boxes, pyramids) are arranged on a surface (a table)
- SHRDLU was aimed at arranging the objects following user's instructions
- using a simulated robot arm with heuristics
- it worked: both the user interaction via a sort of natural language, and the practical reasoning (reasoning towards action)
- limits: language and environment actually very simplistic, not representative of real world interaction



Robots

SHAKY with STRIPS

- a remotely-controlled mobile robot
- moving objects in an office-like environment by pushing them
- planning with STRIPS
- it somehow worked
- limits: limited ability to perceive and understand real-world environment, so the environment had to be built around it to make SHAKY work
- basically, it looked like a real-world robot, but never worked in the real world



Next in Line...

- 1 The Idea of AI
- 2 The Raise of AI
- 3 Main Themes in Early AI
- 4 The AI Winter**
- 5 The Modern Era of AI
- 6 AI Today & Tomorrow



Complexity I

Microworld does not scale up to real-world

- solutions in microworld apparently do not work in real world
- qualitative or quantitative issues?

NP-complete problems

- *Non-deterministic Polynomial time complete* (NP-complete) problems
 - easy to check correctness of solutions
 - almost impossible to *find* a solution in a reasonable amount of time
 - e.g., Travelling Salesman Problem (TSP)
- ! almost all relevant problems in AI are (at least) NP-complete



Complexity II

AI Winter

- AI started with way too much hype
- e.g., General Problem Solver (GPS) never worked for real-world problems
- in a couple of decades, nothing was really accomplished
- ! Funding for AI research were cut everywhere, starting from the seventies



Tackling Complexity with Knowledge I

A way out

- humans do not explore enormous search spaces to find solutions
- they use (possibly extensive) *domain knowledge* to cut down the search space
- e.g., chess players
- ! the idea: exploiting explicit (domain) knowledge representation to make solutions computable



Tackling Complexity with Knowledge II

Expert systems

- a new class of AI systems emerged
 - no longer General AI, since they were tailored to specific domains
 - working only there
 - yet finally facing real-world problems successfully
- *expert systems* were built by
 - eliciting knowledge from domain experts
 - representing them in terms of *rules*
 - able to provide answers to human questions over a domain
 - along with an *explanation*
 - often coping with uncertainty



Tackling Complexity with Knowledge III

MYCIN

- nearly 200 rules
- about blood disease
- working as a doctor's assistant
- issues
 - difficult to extract knowledge from humans, and to precisely formalise it in terms of rules
 - operating only in a very narrow area of expertise



Tackling Complexity with Knowledge IV

CYC

- representing everything in the world in terms of rules
- an expert systems containing all the common sense human knowledge in terms of rules
- Doug Lenat's CYC started in 1984...
- ... and is still alive
- way too ambitious: some call it a failure
- yet it now lives out of its commercial applications



Computational Logic for AI

Logic Programming

- capturing knowledge, inferring new knowledge—with the precision and rigour of logic
- computing with logic: the PROLOG programming language [Colmerauer and Roussel, 1996]
 - knowledge and programs expressed in terms of *logic axioms* (Horn clauses)
 - altogether forming a *logic theory*
 - against which *goals* are demonstrated
 - *computation* as *goal demonstration* [Lloyd, 1984]
- issues
 - heavy computational load
 - not fitting (alone) real-world representation and computation



Bayesian reasoning

Coping with uncertainty

- humans are not good decision makers
- *rational decisions* better than human decisions in most fields, in particular where uncertain information is available
- e.g., when probabilistic reasoning is involved
- Bayes developed the math in the eighteen century, and a couple of centuries were required to make *Bayesian reasoning* usable for AI
- e.g., many NLP systems use Bayesian networks



Nouvelle AI

Against reasoning, against representation [Brooks, 1991a, Brooks, 1991b]

- frustration from absence of real-world results led to focus on real world itself
- against knowledge-based and logic-based AI
- perceive and react as the main cycle of robots in the real world, rather than (symbolic) reasoning



Behavioural AI

Robotics revolution [Brooks, 1986]

- focus on individual intelligent behaviours expected
- just build them, according to how they do relate with each other
- in particular, which ones take the precedence on the others
- e.g., Roomba vacuum cleaner



Next in Line...

- 1 The Idea of AI
- 2 The Raise of AI
- 3 Main Themes in Early AI
- 4 The AI Winter
- 5 The Modern Era of AI**
- 6 AI Today & Tomorrow



Autonomous Vehicles

Grand DARPA Challenges

- race for autonomous vehicles in the desert of Nevada (2005)
- won by STANLEY, a converted Volkswagen Touareg, equipped with seven onboard computers, interpreting sensor data from GPS, laser rangefinders, radar, and video feed



Machine Learning I

Computer that learn

- machine learning (ML) aims at computers that learn knowledge they have not be explicitly told
- they are typically *trained*

Supervised learning

- training by examples of the thing that the computer is trying to learn—*training data*
- e.g, face recognition



Machine Learning II

Reinforcement learning

- a system is able to experiment by making decisions
- receives feedback on those decisions
- and learns from that

Three factor for the success of ML

- scientific breakthroughs—deep learning dealing with complex problems
- training requires a lot of data data—nowadays data are hugely available
- training requires computational power—nowadays computational power is more and more available



Neural Nets

Nets of neurons

- *neural nets* implement ML by connecting many artificial neurons in complex networks
- each neuron gets input from neighbours and produces an output based on some *weights*
- learning occurs by adjusting weights
- two-phase research
 - first, simple neural networks show potential and are proven to have limits
 - then, more complex networks were proven able to overcome these problems
- e.g., character recognition
- ! issue: opaque, no explanation made available—expertise hidden in weights

Triumph of ML

Alpha Go

- in 2014 DeepMind demonstrated a system learning how to play arcade games just looking at the video and accessing the scores, using the same controls as humans
- acquired by Google, they built Alpha Go, which beat Go champion Lee Sedol 4 to 1 in 2016
- exploiting *deep neural networks* along with *self-training*
- Go search space is so huge that brute force just does not work: so, it was considered impossible for a machine to beat a human
- so, this also made everybody know that there were no known limits to the ability that machine intelligent could reach



Next in Line...

- 1 The Idea of AI
- 2 The Raise of AI
- 3 Main Themes in Early AI
- 4 The AI Winter
- 5 The Modern Era of AI
- 6 AI Today & Tomorrow**



AI Today

AI everywhere [Wooldridge, 2018]

- digital assistants like Siri, Alexa, and Cortana are classic AI
- face recognition software in social media is AI
- the software in car's satellite navigation system is AI
- product recommendations from online stores are made by AI
- cruise control systems in modern cars are AI
- AI-based automated translation systems are routinely used across the world
- ...

AI is just everywhere already



Challenges I

Challenges for society

- autonomous vehicles against millions of truck drivers
- conversational systems in contact centres
- intelligent machines will take many skilled roles
- what about unemployment?

AI is just everywhere already

ML vs. privacy

- ML can understand thing about us that are hidden in data, and invisible to humans
- a challenge for law-makers is to rule algorithms that have capabilities far beyond humans

Challenges II

Autonomous weapons

- a blessing or an unacceptable risk?
- more generally, *machine ethics* is nowadays a subject of scientific scrutiny

Fears of future AI

- singularity?
- conscious machines?



References I



Brooks, R. A. (1986).
Achieving artificial intelligence through building robots.
Technical report, Massachusetts Institute of Technology (MIT).



Brooks, R. A. (1991a).
Intelligence without reason.
In Mylopoulos, J. and Reiter, R., editors, *12th International Joint Conference on Artificial Intelligence (IJCAI 1991)*, volume 1, pages 569–595, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.




Brooks, R. A. (1991b).
Intelligence without representation.
Artificial Intelligence, 47:139–159.



Colmerauer, A. and Roussel, P. (1996).
The birth of Prolog.
In Bergin Jr., T. J. and Gibson Jr., R. G., editors, *History of programming languages—II*, volume II, pages 331–367. ACM, New York, NY, USA.



Lloyd, J. W. (1984).
Foundations of Logic Programming.
Springer, 1st edition.



References II



Searle, J. R. (1980).
Minds, brains, and programs.
The Behavioural and Brain Sciences, 3:417–457.



Turing, A. M. (1937).
On computable numbers, with an application to the entscheidungsproblem.
Proceedings of the London Mathematical Society, s2-42(1):230–265.



Turing, A. M. (1950).
I.—computing machinery and intelligence.
Mind, LIX(236):433–460.



Wooldridge, M. (2018).
Artificial Intelligence, volume 117 of *Ladybird Expert*.
Penguin Random House UK.



Artificial Intelligence: Evolution & Perspectives

Digital Culture

Andrea Omicini
andrea.omicini@unibo.it

Master in Finanza, Controllo e Auditing + Marketing and Sales / Bologna Business School

Academic Year 2018/2019

