



Artificial Intelligence

Maurizio Gabbrielli

DISI (Dipartimento di Informatica – Scienza e Ingegneria)
Università di Bologna
FOCUS Team INRIA



Informazioni sul corso

- ◇ Docente: *Maurizio Gabbrielli*
- ◇ Libro di testo: *Russel, Norvig. Intelligenza Artificiale: un approccio moderno, vol. I.* (AIMA Berkeley)
- ◇ Slides by Russel and Norvig

Cap 1 2

Programma

- Introduzione alla IA
- La nozione di agente (intelligente ?)
- Strategie di ricerca (non informata, informata, con avversari). Giochi
- Richiami di logica proposizionale e di logica del prim'ordine
- Unificazione. Risoluzione
- Programmazione logica
- Programmazione con vincoli
- Constraint Satisfaction Problems
- Constraint Handling Rules
- Modellazione e soluzione di problemi mediante agenti logici
- Planning

Cap 1 3

Credit for part of the slides



Paola Mello



Francesca Rossi



UNIVERSITÀ
DEGLI STUDI
DI PADOVA

2 BBS Tel Aviv 2016



Plan

- Some recent facts from the “real world”
- A short history
- Main applications
- Present and future challenges

3 BBS Tel Aviv 2016



But first of all .. what is AI?

- Officially: **The ability of a computer to perform tasks commonly associated with intelligent beings**
- Peter Norvig (Director of Research, Google): “Knowing what to do when you don’t know what to do”
 - We may not be able to observe everything
 - We may not know the exact result of an action
- As opposed to computer science
 - Solving problems we know how to solve

4 BBS Tel Aviv 2016



AI vs. Computer Science



5 BBS Tel Aviv 2016



However ...

- The definition of (natural) intelligence is not an easy one
- Several definitions of intelligence
- Even more definitions of AI
- More later (perhaps) ...

6 BBS Tel Aviv 2016



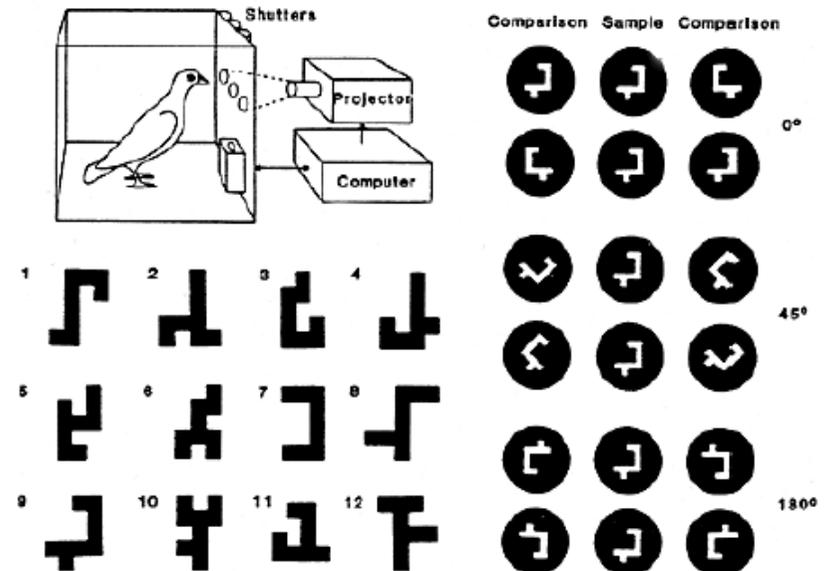
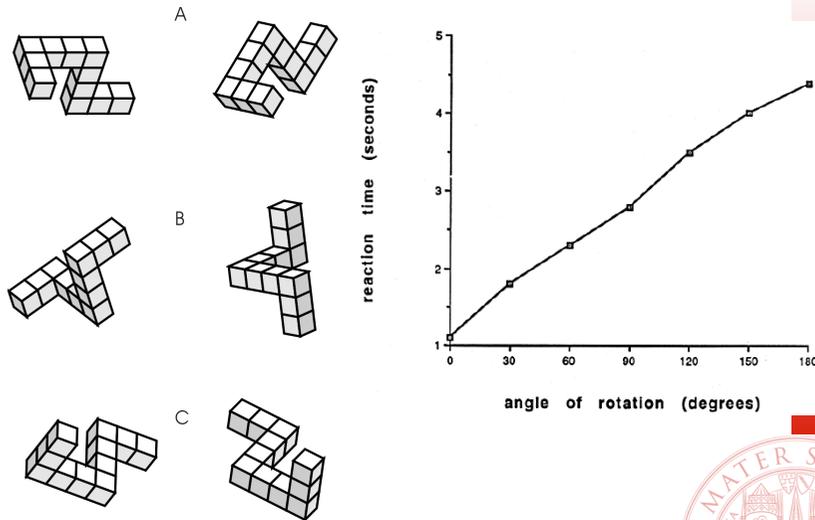
However ...

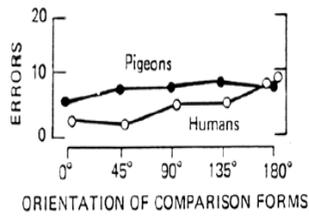
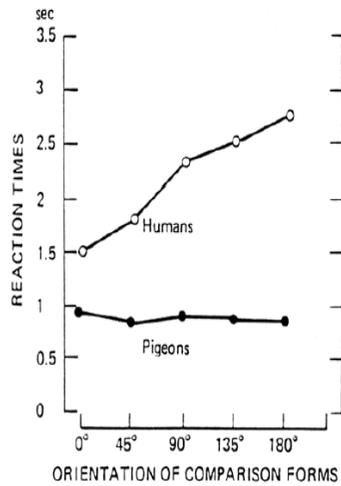
- The definition of (natural) intelligence is not an easy one
- Several definitions of intelligence
- Even more definitions of AI
- More later (perhaps) ...



Short digression on (natural) intelligence

- Is more intelligent the pigeon or the human being ?
- Probably the human being however ...
- For some tasks the pigeon is better than us

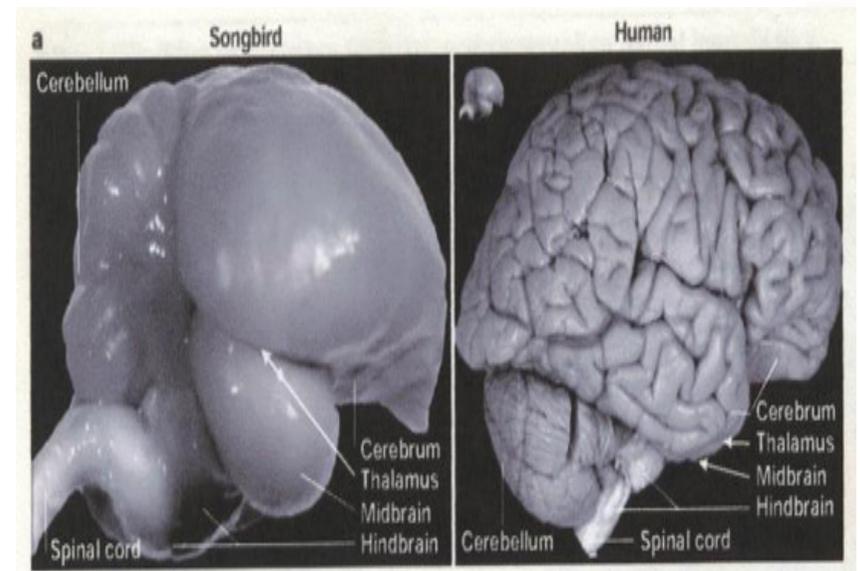
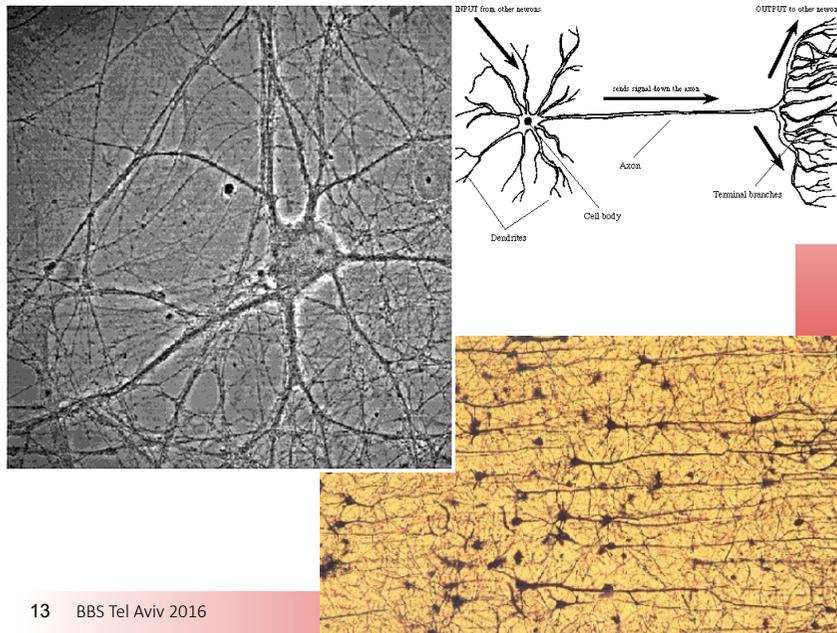




Learning abilities:

Number of reinforcements needed to learn a simple motor response

bee	2
carp	4
quail	8
pigeon	10
hen	18
rat	22
raccoon	24
rabbit	24
So who is more intelligent ?	
child	28 !



Brain dimensions matters ?

Whale: 8 Gg. - 200 billions of neurons

Human being : 1.2 – 1.4 Kg – 85 billions

....

....

....

Pigeon: 0.003 Kg. - 400.000.000 neurons

Bee: 1 mm³ - < 1 million of neurons



Some definitions of Intelligence

The Oxford English dictionary: 'quickness of understanding'

The Cambridge Advanced Learner's dictionary: 'the ability to learn, understand and make judgements or have opinions that are based on reason'

Enciclopedia Rizzoli-LaRousse: 'facoltà di conoscere, di comprendere, di intendere'

Wikipedia: 'L'intelligenza è l'insieme di funzioni conoscitive, adattative e immaginative generate dall'attività cerebrale dell'uomo e degli animali'



A possible definition

Intelligence is the ability to learn from experience, to apply knowledge to solve problems, and to adapt and survive in different environments (social and geographical)

In more simple terms, intelligence equals survival. If we show intelligence then we will survive and prosper within the social and geographical contexts that we reside in and encounter



G-factor

Spearman (1923) identified the g-factor (g for general) as the dewcription of the general intelligence of an individual. His idea was that the g-factor is the basis for the global cognitive ability of a person, and underlies all the various froms of specific abilities (e.g. mathematics, language etc.)

Other researchers (e.g. Guilford, 1967) claim that there exist no general intelligence.



Primary abilities

Thurstone (around 1930) identified several different factors describing the intelligence.

The most common classification uses 7 factors:

- Verbal fluidity
- Numerical ability
- Inference
- Spacial ability
- Velocity in perception
- Memory



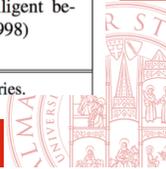
What is AI?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

Some definitions of AI (Russel and Norvig)

<p>Thinking Humanly</p> <p>“The exciting new effort to make computers think ... <i>machines with minds</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>Thinking Rationally</p> <p>“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)</p>
<p>Acting Humanly</p> <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>Acting Rationally</p> <p>“Computational Intelligence is the study of the design of intelligent agents.” (Poole <i>et al.</i>, 1998)</p> <p>“AI ... is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)</p>

Figure 1.1 Some definitions of artificial intelligence, organized into four categories.



Thinking humanly: Cognitive Science

1960s “cognitive revolution”: information-processing psychology replaced prevailing orthodoxy of **behaviorism** (this was concerned with “measures” of stimuli/answers only)

Requires scientific theories of internal activities of the brain

- What level of abstraction? “**Knowledge**” or “**circuits**”?
- How to validate? Requires

- 1) Predicting and testing behavior of human subjects (top-down)
- or 2) Direct identification from neurological data (bottom-up)

Both approaches (roughly, **Cognitive Science** and **Cognitive Neuroscience**) are now distinct from AI

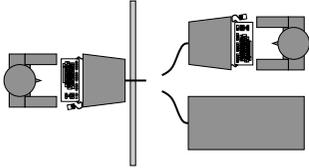
Both share with AI the following characteristic:

the available theories do not explain (or engender) anything resembling human-level general intelligence

Hence, all three fields share one principal direction!

Acting humanly: The Turing test

- Turing (1950) "Computing machinery and intelligence":
- ◇ "Can machines think?" → "Can machines behave intelligently?"
 - ◇ Operational test for intelligent behavior: the **Imitation Game**



- ◇ Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- ◇ Anticipated all major arguments against AI in following 50 years
- ◇ Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible, constructive**, or amenable to **mathematical analysis**

A comparison

	Computer	Brain
Computational units	10^8 logical ports	10^{11} neurons
Memory	$10^{10} + 10^{11}$ bit (RAM+ Hard)	$10^{11} + 10^{14}$ (neurons + synapses)
Clock	10^{-9} sec	10^{-3} sec (parallel !)
Bandwidth	10^{10} bit/sec	10^{14} bit/sec

Brain 100.000 times faster than computer in performing activities (all neurons work in parallel).

Thinking humanly: imitate the brain

Why not imitate the human brain ? We do not know enough:

- Broca (1861) relation between language and specific areas in the brain
 - Golgi (1873) first visualization of neurons
 - Berger (1929) EEG
 - Ogawa et al (1990) fMRI: measurement of cognitive processes.

Today

- We know some relations between areas in the brain and specific functions
 - do not know how a cognitive process works
 - do not know how memory works

Thinking rationally: Laws of Thought

Normative (or **prescriptive**) rather than **descriptive**

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of **logic**:
notation and **rules of derivation** for thoughts;
 may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:

- 1) Not all intelligent behavior is mediated by logical deliberation
- 2) **What is the purpose of thinking?** What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good

Cap 1 11

Rational agents

An **agent** is an entity that perceives and acts

This course is about designing **rational agents** (più o meno ...)

Abstractly, an agent is a function from percept histories to actions:

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

Caveat: **computational limitations make perfect rationality unachievable**

→ design best **program** for given machine resources

Cap 1 12

Rational agents

Cosa significa razionale ?

- provvisito di ragione come facoltà peculiare dell'uomo (Devoto - Oli)

Ragione: facoltà di pensare, peculiare dell'uomo, soprattutto in quanto capacità di discernere, di determinare rapporti logici e di formulare giudizi (Devoto - Oli).

.... non proprio chiaro cosa significhi in un contesto artificiale (senza assumere imitazione del comportamento umano ...)

Cap 1 13

AI is already here in the real world

21 BBS Tel Aviv 2016

AI in the media

La scalata di Facebook all'intelligenza artificiale

Zuckerberg pronto a scommettere nel settore investendo in Vicarious. I suoi alleati sono Elon Musk, il fondatore di Tesla e l'attore Ashton Kutcher



Il CEO di Facebook Mark Zuckerberg



Quando un romanzo scritto da un computer 'rischia' di vincere un premio letterario; e se il nuovo Balzac fosse di plastica e metallo?
di Maurizio Di Fazio

Nell'era dello storytelling di massa, l'intelligenza artificiale sceglie la scrittura di fantasia per tornare a mostrare i propri muscoli di microchip. La notizia viene dal Giappone: un importante premio letterario fantascientifico nazionale. Il Nikkei Hoichi Shintachi Literary Award, per un pelo non è stato vinto da un romanzo scritto da un robot. Il libro è intitolato programmaticamente The Day A Computer Writes A novel
di Maurizio Di Fazio | 30 marzo 2016

AI in the media: personal assistant and expert systems

Jill Watson

The Teaching Assistant of the 2016 Artificial Intelligence course at Georgia Tech was a program (based on IBM Watson). It was answering students questions online with a success rate of 97%. Students were not aware to talk to a computer.

Watson save a life.

In June 2016 IBM Watson saved the life of a Japanese woman by correctly identifying her disease. The system looked at the woman's genetic information and compared it to 20 million clinical oncology studies. After doing so, it determined that the patient had an exceedingly rare form of leukemia.

Virtual Judge

UCL AI algorithm examined English language data sets for 584 cases relating to torture and degrading treatment, fair trials and privacy. In each case, the software analysed the information and made its own judicial decision. In 79% of those assessed, the AI verdict was the same as the one delivered by the court.

AI in the media: art and ethics



Magenta (Google 2016)

Automatic music and art generation. Uses Tensor Flow machine learning system. Released Open source on GitHub.

Sunspring: first screenplay written by a computer (a neural network trained with tens of SF movies). Director Oscar Sharp and AI Expert Ross Goodwin.

Moral algorithms for autonomous vehicles.

Jean-Francois Bonnefon at the Toulouse School of Economics made a study asking ordinary people. People are in favour of cars that sacrifice the occupant to save other lives—as long they don't have to drive one themselves.

AlphaGo beat the world champion of Go

Go: game 2500 years old, very common in Asia (40M players)

Very simple rules but very complex game: possible positions $10^{365} > \#$ atoms in the universe 10^{85}



March 2016: AlphaGo win against Lee Sedol



New research centers

Leverhulme Centre for the Future of Intelligence (LCFI) at Cambridge University opened in October 2016

A multi-disciplinary institute that will attempt to tackle some of the open-ended questions raised by the rapid pace of development in AI research.

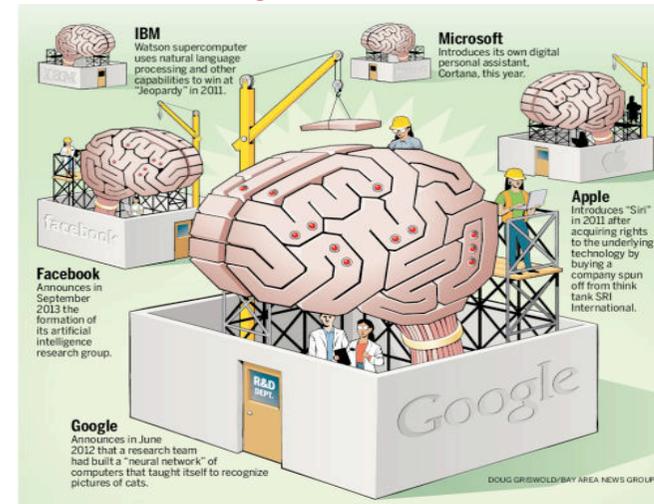
Stephen Hawking

“We spend a great deal of time studying history, which, let’s face it, is mostly the history of stupidity. So it’s a welcome change that people are studying instead the future of intelligence.”

The creation of powerful artificial intelligence will be “either the best, or the worst thing, ever to happen to humanity”



Increasing interest for AI in the companies



Recent facts



- 2014. Acquired AI startup DeepMind for \$ 400 M (Deep learning, Differentiable Neural Computer)
- Bought 14 AI and robotics companies in the past few years
- Machine learning system TensorFlow freely available



- Launched a fund and an incubator for AI startup (2016)
- Oxford project for emotion recognition now open to developers: free APIs for speech recognition, vision etc. including the ability to recognize sad and happy faces



Recent facts



- Famous for Watson (2011)
- Developing a teaching assistant



- 2015. Acquired Vocal IQ: speech recognition for development of SIRI
- 2016. Acquired AI startup Emotient: facial recognition and reaction to ads



Recent facts

facebook

- AI Research lab. (FAIR) with 30 top scientist from various countries, headed by the inventor of CNN. 2013
- Face recognition with accuracy of 97% used for automatic tagging of photos videos (can recognize the sport being played). 2014.
- Natural language processing and translation (493 directions)
- Personal assistant M, based on the work of the startup WIT.ai acquired by FB in 2015. Can book a flight and make travel plans.
- Embed the world project: “replacing reasoning with algebra” (vectors)



- 2016. Acquired Meta Mind, and AI setarup specialized in deep learning. Tecniche AI usate per CRM (Customer Relationship Management)
- Given a large text should answer questions like: What is the overall sentiment ?



Recent facts

U B E R

- December 2016. Acquired Geometric Intelligence, an AI startup specialized in machine learning. Set up of an AI Lab. with 15 scientists
- Investing heavily in autonomous driving technology



- Launching virtual assistant to answer on-line customer questions



- Real-time language translation in 50 languages.



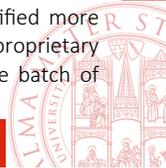
Success stories for data analytics

Food Genius and open data. A foodservice data provider that scrapes data from restaurant menus posted online and searches for local trends to help companies like **Kraft Foods** to develop and market products.

Westpac Banking Corp. and the 360-degree view of the customer. The Australian bank used SAS technology capturing and analyzing the behaviour (ATM, call center interaction) of its 12 M customer to build new offerings. In 9 months, customer engagement from 1% to 25%.

Oreal Group and customer engagement. Using CRM technology from Clarabridge, L'Oreal is analyzing tweets, Facebook posts, product reviews and news stories in a “command center”, claiming this has “transformed how brand awareness and loyalty are leveraged”.

Coca-Cola Co. and product consistency. Coke engineers identified more than 600 (!) possible flavors of oranges and then built a proprietary algorithm to ensure consistent taste and texture from one batch of orange juice to the next.

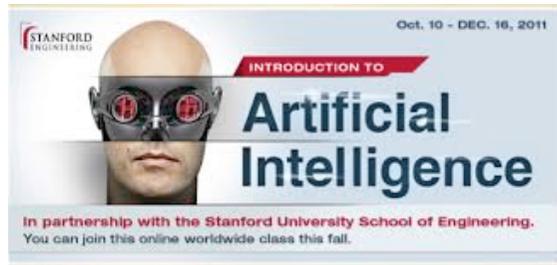


Investments in AI

- More than **\$17 billion** invested in the period 2009-2015
- 2013: more than \$2 billion invested by 322 companies with AI-like technology
- Private investment in the AI sector expanding 62% a year on average in the past 4 years
- Apple, Google, Facebook, Uber, Yahoo, Intel, Dropbox, LinkedIn, Pinterest, Twitter, Salesforce all have purchased AI companies



More than 160,000 students!



- Some recent facts from the “real world”:
- A short history
- Main applications
- Present and future challenges



René Descartes 1596-1610

Influenced by the automatons on display throughout the city of Paris, began to investigate the connection between the mind and body. Dualism:

- body works like a machine, that it has material properties.
- mind is not material and does not follow the laws of nature.
- mind interacts with the body at the pineal gland

Cartesio, il filosofo della machine de terre può essere considerato l'antesignano della scienza cognitiva e dell'Intelligenza Artificiale, in virtù della portata conoscitiva che ha attribuito alla costruzione degli automi. ¹

1) Bruno G. Bara, Scienza cognitiva. Un approccio evolutivo alla simulazione della mente, Bollati Boringhieri, Torino 1982

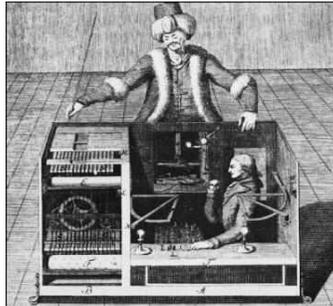


Cartesio (Discorso sul metodo - 1637)

“Qui in particolare mi ero fermato per far vedere che se ci fossero macchine con organi e forma di scimmia o di qualche altro animale privo di ragione, non avremmo nessun mezzo per accorgerci che non sono in tutto uguali a questi animali; mentre se ce ne fossero di somiglianti ai nostri corpi e capaci di imitare le nostre azioni per quanto di fatto possibile, ci resterebbero sempre due mezzi sicuri per riconoscere che, non per questo, sono uomini veri. **In primo luogo, non potrebbero mai usare parole o altri segni combinandoli come facciamo noi per comunicare agli altri i nostri pensieri. Perché pur nel concepire che una macchina sia fatta in modo tale da proferire parole, e ne proferisca anzi in relazione a movimenti corporei che provochino qualche cambiamento nei suoi organi; che chieda, ad esempio, che cosa si vuole da lei se la si tocca in qualche punto, o se si tocca in un altro gridi che le si fa male e così' via; ma non si può immaginare che possa combinarle in modi diversi per rispondere al senso di tutto quel che si dice in sua presenza, come possono fare gli uomini, anche i più ottusi.** L'altro criterio è ' che quando pure facessero molte cose altrettanto bene o forse meglio di qualcuno di noi, fallirebbero inevitabilmente in altre, e si scoprirebbe così' che **agiscono non in quanto conoscono**, ma soltanto per la disposizione degli organi.”



AI ante litteram ... fake !



An image from
18th century



AI (and CS) ante litteram in the XVII century

- G.W. Leibniz (1646-1716) Calculus raziocinator.
- Jacques de Vaucanson(1709-1782) built an artificial duck that could swim and “eat” grain (400 parts in each wing)
- Pierre Jacquet-Droz (1721–1790) swiss-born watchmaker who designed and built animated dolls and automata including a writer and a musician
 - Movements controlled by an input coded on a metallic disc !

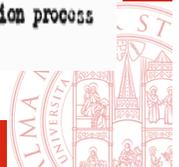
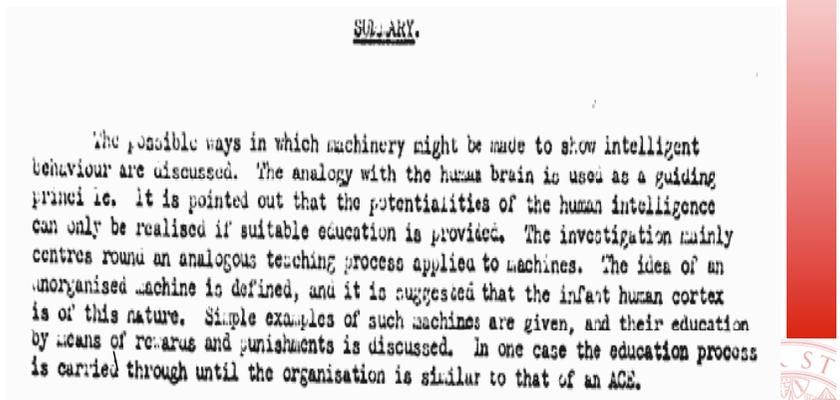


AI before AI term was coined, i.e. Turing

- 1935 Universal Turing Machine
- 1947 Talk by A. Turing in London: "What we want is a machine that can learn from experience ... the possibility of letting the machine alter its own instructions provides the mechanism for this".
- 1948. *Intelligent Machinery*. Never published during his life, a **real AI manifesto** contain several ideas later re-invented (including symbolic systems and neural networks).
- 1950. *Computing Machinery and Intelligence*. Including the imitation game, i.e. Turing test.



Intelligent Machinery (1948) http://www.alanturing.net/turing_archive/



Turing, A.M. (1950). Computing machinery and intelligence. Mind, 59, 433-460.

1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

The new form of the problem can be described in terms of a game which we call the 'imitation game.' It is played with three people, a man (A), a

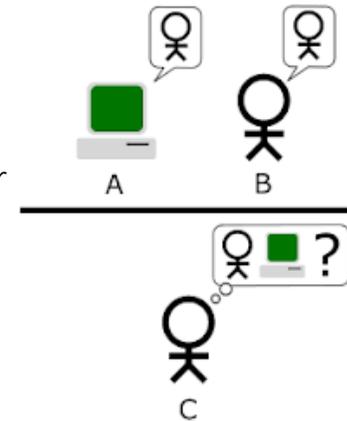


Turing test, 1950

Can machines think?

Behavioral test

If it is impossible to distinguish the behavior of a machine from that of a human, then the machine is intelligent



Neurons and learning

- W. S. McCulloch and W. H. Pitts (1943). First computational model of the neuron (threshold logic)
At the basis of the research on neural networks

- D.O. Hebb (1949). Hebbian learning.

A theory in neuroscience that proposes an explanation for the adaptation of neurons in the brain during the learning process, describing a basic mechanism for synaptic plasticity neuronal basis of unsupervised learning

- F Rosenblatt (1958). Perceptron.

An algorithm for pattern recognition based on a two-layer computer learning network using simple addition and subtraction



The official start of AI: Dartmouth conference 1956

John McCarthy coined the term AI: "The study is to proceed on the basis of the conjecture 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".



Trenchard More, John McCarthy, Marvin Minsky, Oliver Selfridge, Ray Solomonoff



They were very optimistic

“In 1968 a computer will be the world chess champion”, Simon and Newell 1958

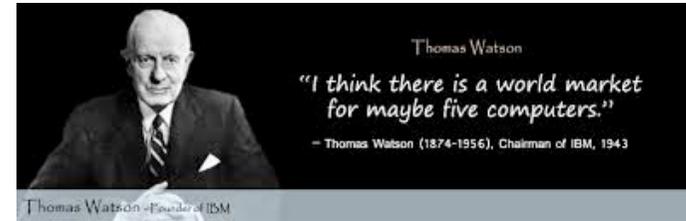
“Machine will be capable of doing any work a man can do in 20 years”, Simon 1965



IBM Deep Blue won against Kasparov, 1997



No idea of what computers could do!



In general the future is difficult to forecast....

“There is no reason anyone would want a computer in their home.”

Ken Olson, DEC co-founder, 1987

“The problem of viruses is temporary and will be solved in two years.”

John Mac Affee, 1988

“The Internet? We are not interested in it.”

Bill Gates, 1993

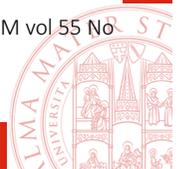


Don't feel bad if you cannot predict the future ..

- The prediction performance of experts is not better than random guessing.
- Correct predictions are often **way off in timing**

Based on a study of Philip Tetlock on 27,450 predictions by 284 experts in many fields.

Don't Feel Bad if You Can't Predict the Future, Peter J. Denning, CACM vol 55 No 9, September 2012



AI prehistory

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms, computation, (un)decidability, (in)tractability probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

Cap 1 14

Potted history of AI

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1952–69	Look, Ma, no hands!
1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist Proved most of theorems in Russell's Principia Mathematica Gelernter's Geometry Engine
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity No domain information Failure of automatic translation: Lo spirito è forte ma la carne e' debole - = La vodka è buona ma la bistecca è marcia Computational complexity Models not adequate (e.g. perceptrons) Neural network research almost disappears

Cap 1 15

Potted history of AI

1969–79	Early development of knowledge-based systems (DENDRAL, MYCIN) Domain Knowledge
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural networks return to popularity
1988–	Resurgence of probability; general increase in technical depth "Nouvelle AI": ALife, GAs, soft computing
1995–	Agents, agents, everywhere . . .
2003–	Human-level AI back on the agenda

Cap 1 16

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis

Cap 1 17

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge

Cap 1 22

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem

Cap 1 23

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology

Cap 1 24

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story

Cap 1 25

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law

Cap 1 26

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law
- ◇ Translate spoken English into spoken Swedish in real time

Cap 1 27

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law
- ◇ Translate spoken English into spoken Swedish in real time
- ◇ Converse successfully with another person for an hour

Cap 1 28

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law
- ◇ Translate spoken English into spoken Swedish in real time
- ◇ Converse successfully with another person for an hour
- ◇ Perform a complex surgical operation

Cap 1 29

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law
- ◇ Translate spoken English into spoken Swedish in real time
- ◇ Converse successfully with another person for an hour
- ◇ Perform a complex surgical operation
- ◇ Unload any dishwasher and put everything away

Cap 1 30

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Telegraph Avenue
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Berkeley Bowl
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law
- ◇ Translate spoken English into spoken Swedish in real time
- ◇ Converse successfully with another person for an hour
- ◇ Perform a complex surgical operation
- ◇ Unload any dishwasher and put everything away

Cap 1 31

AI News

Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997

OTTER proved a mathematical conjecture (Robbins conjecture) unsolved for decades

No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego, 2797 miles)

During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people

NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft

Proverb solves crossword puzzles better than most humans

Cap 1 32

Unintentionally funny stories

One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe threatened to hit Irving if he didn't tell him where some honey was. The End.

Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.

Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.

Cap 1 33

Unintentionally funny stories

Joe Bear was hungry. He asked Irving Bird where some honey was. Irving refused to tell him, so Joe offered to bring him a worm if he'd tell him where some honey was. Irving agreed. But Joe didn't know where any worms were, so he asked Irving, who refused to say. So Joe offered to bring him a worm if he'd tell him where a worm was. Irving agreed. But Joe didn't know where any worms were, so he asked Irving, who refused to say. So Joe offered to bring him a worm if he'd tell him where a worm was . . .