Breve storia dell'intelligenza artificiale

dalla nascita a oggi

Maurizio Parton, Università di Chieti-Pescara

23 maggio 2019

John McCarthy, Dartmouth Conference on AI, 1956

We propose [. . .] a 2-month, 10-men study of artificial intelligence [. . .] [We 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.

An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

- The term *Artificial Intelligence* (AI) is born, together with the idea of *learning*. Big expectations!

First big AI successes

- Logic Theorist, 1956: proves theorems in Logic.
- Kolmogorov's Representation Theorem, 1957: the only continuous functions of several variable is the addition.
- First *Neural Network* (NN): *Perceptron*, 1957.

New York Times, 8 July 1958

The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

First big AI successes

New York Times, 8 July 1958

The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

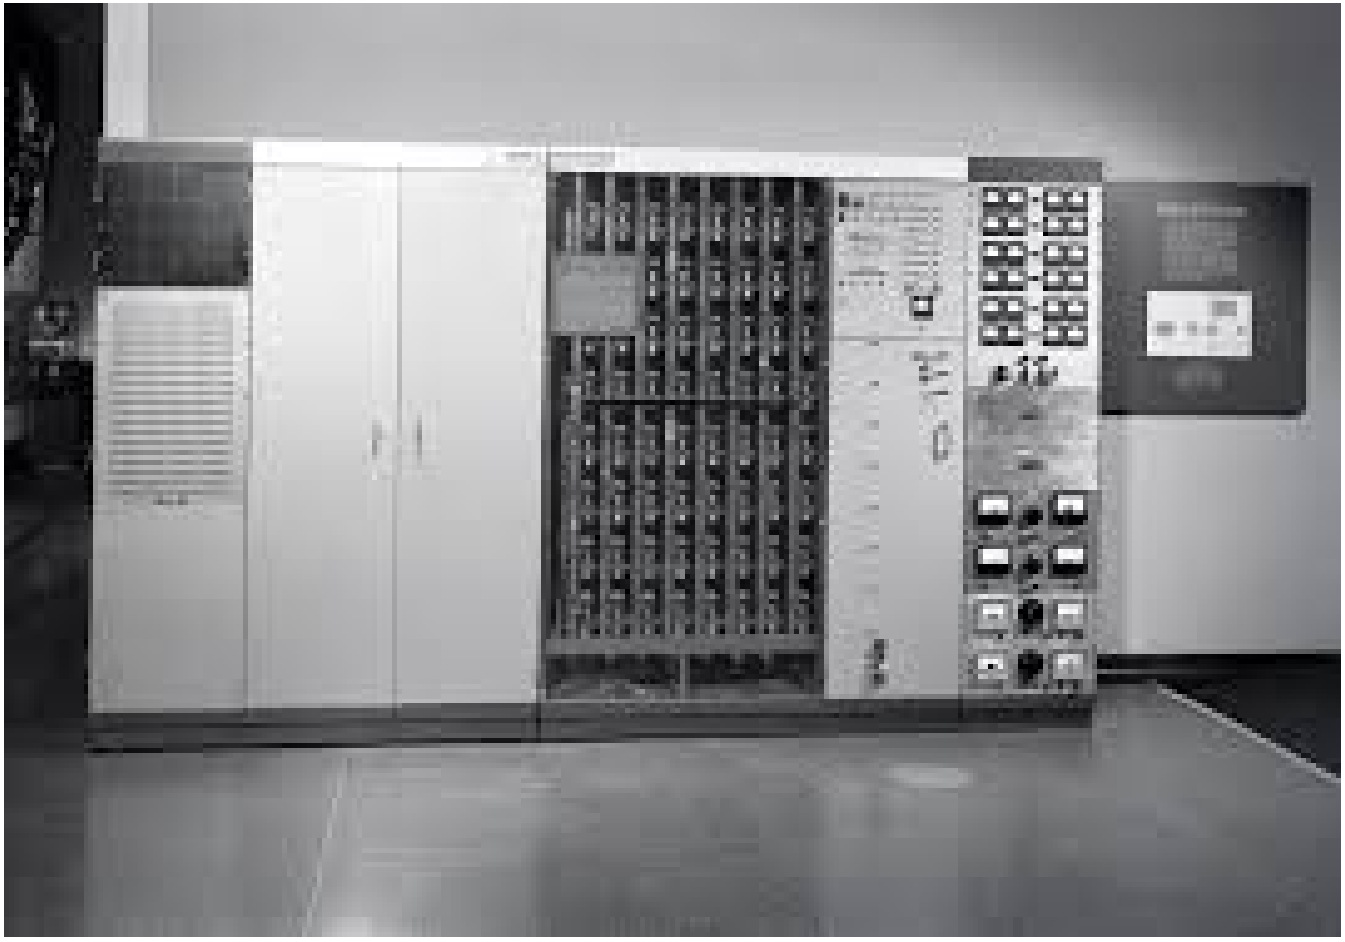

First big AI successes

- General Problem Solver, 1959: proves general theorems.

Jerome Wiesner, `https://youtu.be/aygSMgK3BEM`, 1961

I suspect if you come back in 4-5 years I'll say surely they really do think.

Herbert Simon, 1965

Machines will be capable, within twenty years, of doing any work a man can do.

Self-improving machines

John Good, Advances in Computers, 1965

Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion', and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control...

- *the intelligence of man would be left far behind* $\rightarrow$ Singularity.
- *provided that the machine is docile enough* $\rightarrow$ Need for control: the *Red Button*.

Self-improving machines

Marvin Minsky, 1967

Within a generation [. . .] the problem of creating 'artificial intelligence' will substantially be solved.

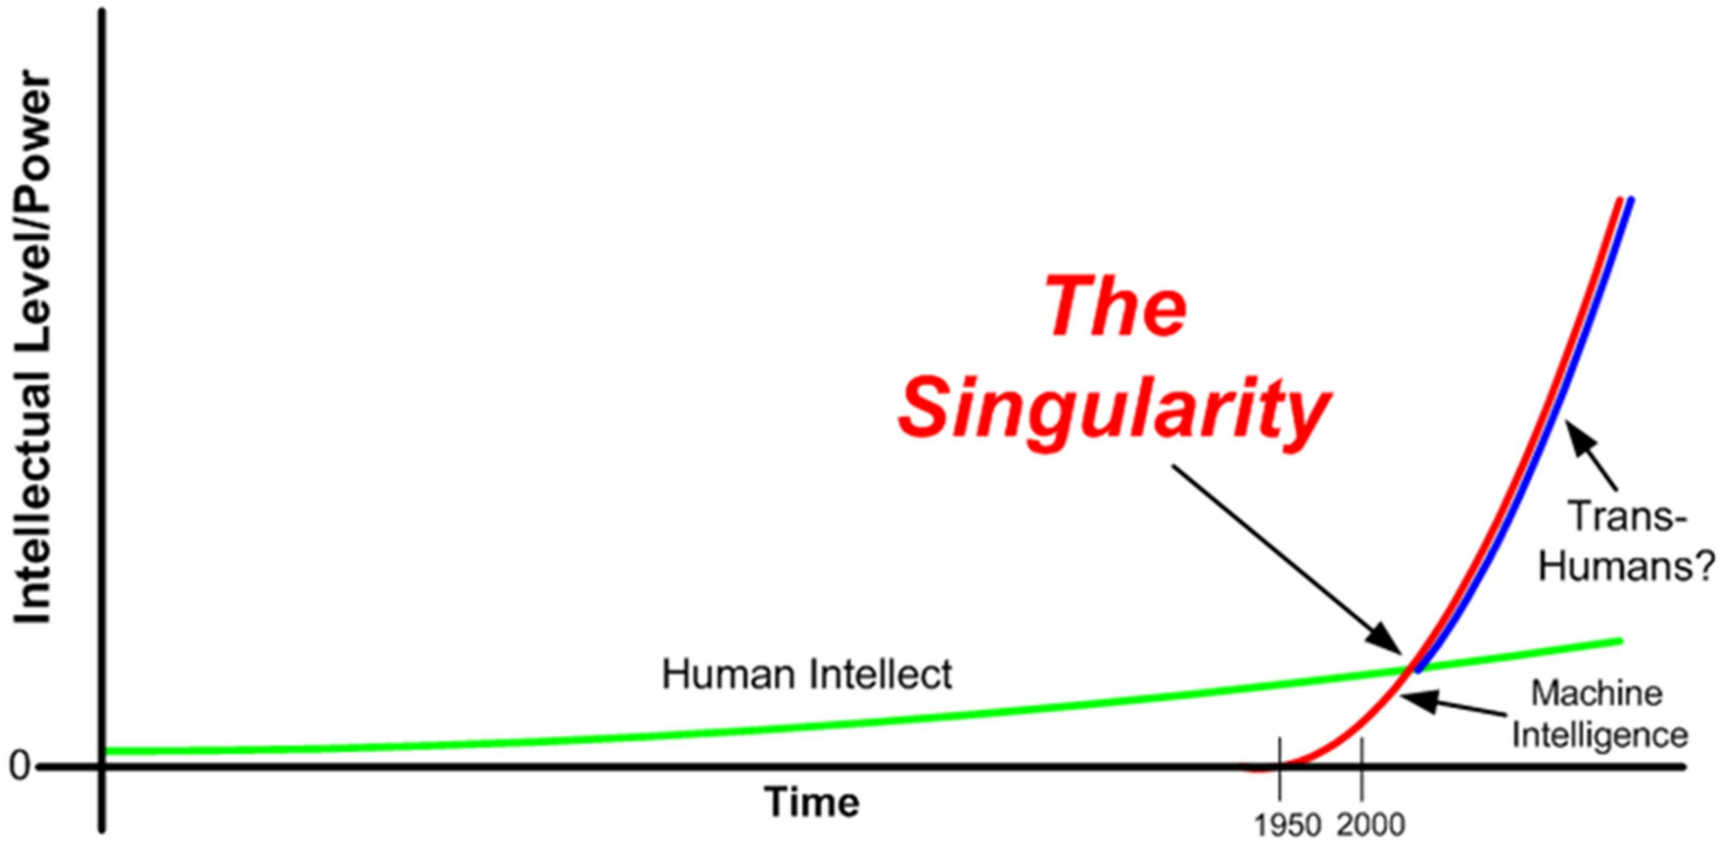

The “Perceptrons” book, a milestone, 1969

- “Perceptrons” Marvin Minsky, Seymour Papert, 1969. First mathematical analysis of a neural network. Beautiful reading, strongly suggested!
- It describes *failures* (XOR), *merits* (AND, OR, NOT) and *future* (multilayer networks) of perceptrons.

The “Perceptrons” book, a milestone, 1969

First AI winter: 1974-1980

- Limited computing power: exhaustive search doesn't work.
- Too much faith in advances in computing power.
- Cold war, failure in automatic translations Russian-English: "the spirit is willing but the flesh is weak" $\leftrightarrow$ "the vodka is good but the meat is rotten" (this citation is probably a myth, but gives the idea).

Failures $\Rightarrow$ Pessimism in AI community $\Rightarrow$ Pessimism in the press $\Rightarrow$ No more funding $\Rightarrow$ End of scientific research.

After the winter: 1980-1987

- *Expert systems*: AI specialized in one task.
- Mycin: identify bacterial infections and recommend antibiotics with the correct dosage.
- Dendral: identify unknown organic molecules.
- SID: CPU design.

After the winter: 1980-1987

- Paul Werbos, *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*, PhD thesis, 1974: *backpropagation* algorithm.
- With the backpropagation, the perceptron is now able to learn!
- *Deep NN* (DNN): NN with many layers can be trained.
- New successes $\Rightarrow$ New trust $\Rightarrow$ New money $\Rightarrow$ New research.

Q: "Is 1974 a typo?".
A: "No. It is just a very bad choice of time!".

After the winter: 1980-1987

Paul Werbos, 2006

In the early 1970s, I did in fact visit Minsky at MIT. I proposed [. . .] a way that uses the reverse method, which we now call backpropagation in the ANN field. But Minsky was not interested. In fact, no one at MIT or Harvard or any place else I could find was interested at the time.

- Werbos published his PhD thesis only in 1982, because of the AI winter!
- The backpropagation algorithm was rediscovered several times from 1980 to 1987.
- David Rumelhart, Geoffrey Hinton, and Ronald Williams rediscovered the technique in 1985-1986. They succeeded in making it widely known.

New approaches to learning

- Successes in speech recognition, with *Recurrent NN* (RNN).
- *Reinforcement Learning* (RL): a new learning technique that does not require a training set, 1989.
- *TD-gammon*, 1995: a reinforcement learning NN playing backgammon at world-champion level.

Second AI winter: 1987-(1994)-2006

- Sepp Hochreiter, *Untersuchungen zu dynamischen neuronalen Netzen*, Diploma thesis, 1991: backpropagation applied to DNN suffers of the *vanishing or exploding gradients* problem. *Perceptron-like effect* on fundings!
- *Shallow* NN can be used only for very simple tasks.
- The techniques used for Backgammon doesn't work for Chess and Go.
- Deep Blue, 1997: great success in Chess *without* using any AI technique!
- Other more promising techniques appear: random forests and support vector machines.

Second AI winter: 1987-(1994)-2006

The Star, 17 April 2015

in 2004, Hinton asked to lead a new program on neural computation. The mainstream machine learning community could not have been less interested in neural nets.

The Star, 17 April 2015

"It was the worst possible time," says Bengio [. . .] "Everyone else was doing something different."

The Deep Learning conspiracy

Table: Geoffrey Hinton, Yann LeCun, Yoshua Bengio

- They perseverate, thanks to *Canadian Institute for Advanced Research* (CIFAR).
- They rebrand the topic under the term *Deep Learning*.

Geoffrey Hinton, Simon Osindero, Yee-Whye Teh, 2006

A breakthrough paper that *solves the vanishing or exploding gradients problem*: a new era for DNN begins!

After 2006

- MNIST digits recognition with 1% error rate (2006).
- DNN are more efficient for difficult problems than shallow methods (2007).
- Speech recognition with 23% error rate (2011).
- Image recognition with *Convolutional NN* (CNN): win the ILSVRC-2012 (Olympics of computer vision) with 15.3% error rate - second place got 26.2%!
- Speech Recognition and instant translation: `https://www.youtube.com/watch?v=Nu-nlQqFCKg`, 2012.
- Skype real time translation: 2014.

After 2006

- Atari Breakout: https://www.youtube.com/watch?v=V1eYniJ0Rnk, 2015.
- AlphaGo, supervised learning: 2016. Wins against best human players.
- AlphaGo Zero, reinforcement learning: 2017. Wins against AlphaGo.
- Alpha Zero, plays Go, Chess and Shogi with reinforcement learning: 2017. Wins against best Chess software.
- AlphaStar, plays StarCraft II at professional level, 2019.

Why this is happening now?

- Backpropagation and ReLU.
- Moore's law: $x2$ every 18 months.
- Graphics cards (GPU): cheap computing power.
- Facebook, Google and Microsoft entered the game.
- 1998: MNIST, $60 * 10^3$ images. 2012: ImageNet, $1.2 * 10^6$ images.
- Big Data now available.
- A feedback of exciting new successes.


That's all Folks!

Licenza