ARTIFICIAL INTELLIGENCE
HISTORY, ACHIEVEMENTS AND OPEN PROBLEMS

Eduardo Bezerra
ebezerra@cefet-rj.br

What is AI?

[hype.]

What is AI?

Field of CS that seeks to explain and simulate, through mechanical or computational processes, some or all aspects of human intelligence.

"AI is the science of how to get machines to do the things they do in the movies."

Astro Teller

History

Artificial Neuron

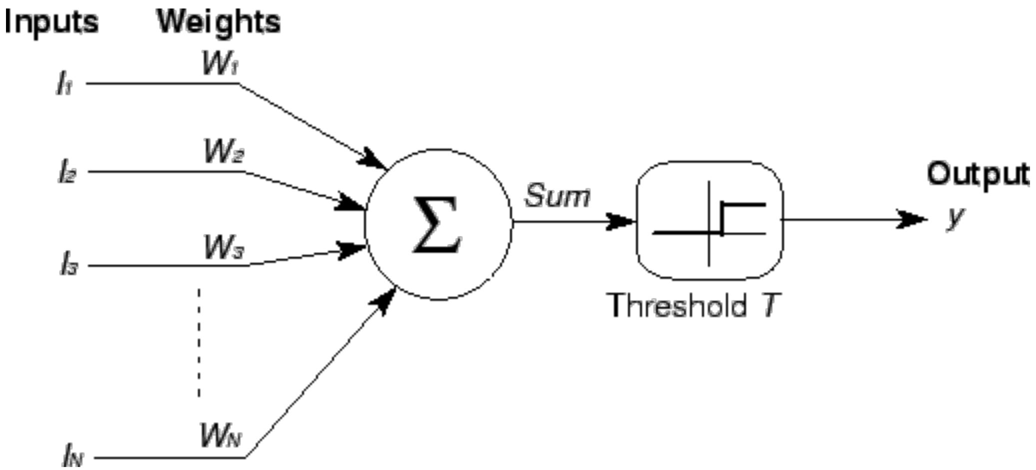

McCulloch-Pitts model (boolean functions, cycles)

1943

The Imitation Game

"I propose to consider the question, 'Can machines think?'"

VOL. LIX. No. 236.] [October, 1950

MIND
A QUARTERLY REVIEW
OF
PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND INTELLIGENCE

BY A. M. TURING

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

1950

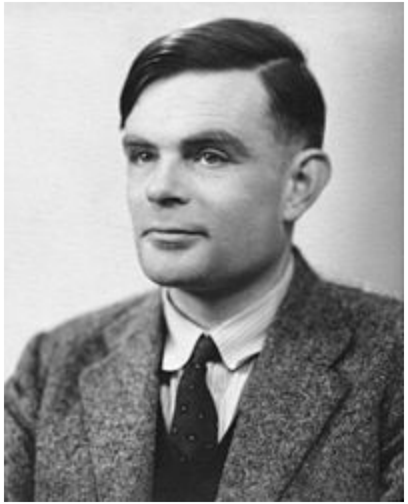

Alan Turing

Georgetown-IBM experiment

- Machine translation: automatic translation of Russian sentences into English.

"The experiment was considered a success and encouraged governments to invest in computational linguistics. The project managers claimed that machine translation would be a reality in three to five years."

1954

Logic Theorist

The first (well, not the first!) AI program

Proved 38 of the first 52 theorems in chapter 2 of the Principia Mathematica.

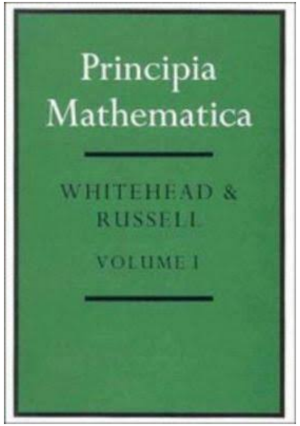

1956

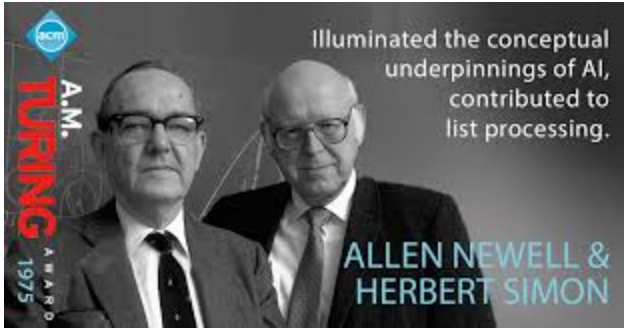

Dartmouth meeting

- John McCarthy, Claude Shannon, Marvin Minsky, …

"We propose that a […] study of artificial intelligence be carried out […]. 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. 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."

1956

Predictions…

"It is not my aim to surprise or shock you –but … there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until –in a visible future – the range of problems they can handle will be coextensive with the range to which human mind has been applied. More precisely: within 10 years a computer would be chess champion, and an important new mathematical theorem would be proved by a computer.

--Herbert Simon

1957

Lisp

list processing, lambda expressions, function definition, …

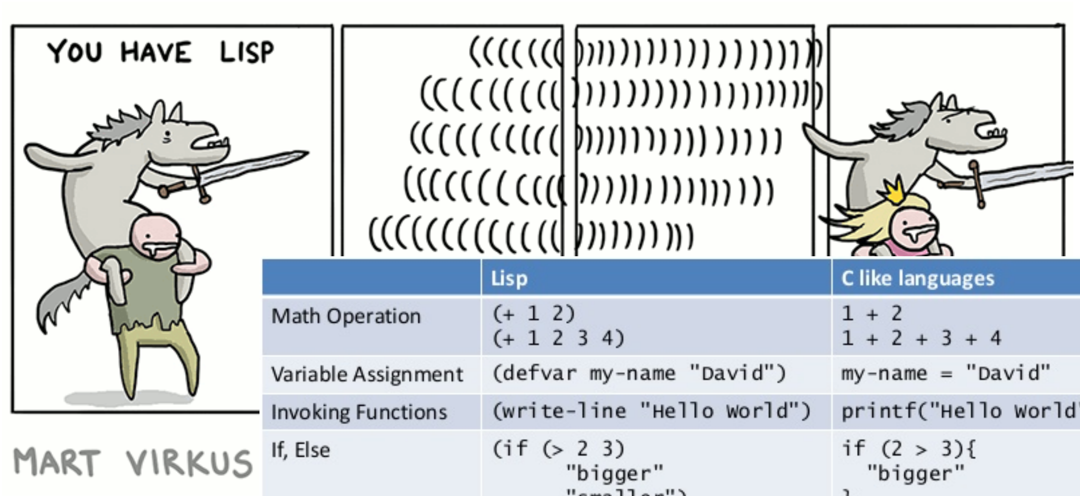

	Lisp	C like languages
Math Operation	`(+ 1 2)` `(+ 1 2 3 4)`	`1 + 2` `1 + 2 + 3 + 4`
Variable Assignment	`(defvar my-name "David")`	`my-name = "David"`
Invoking Functions	`(write-line "Hello World")`	`printf("Hello World")`

If, Else	`(if (> 2 3)`
	`"bigger"`
	`"smaller")`
		`if (2 > 3){`
		`"bigger"`
		`}`
		`else`
		`"smaller"`
		`}`

Function definition	`(defun add (a b)`
	`(+ a b))`
		`int add(int a,int b){`
		`return a + b;`
		`}`

John McCarthy

1958

Machine Learning (to play Checkers)

Coined the term Machine Learning ("Field of study that gives computers the ability to learn without being explicitly programmed.")

"it will learn to play a better game of checkers than can be played by the person who wrote the program."

search tree
alpha-beta pruning
scoring functions

Arthur Samuel

1959

Mark I Perceptron

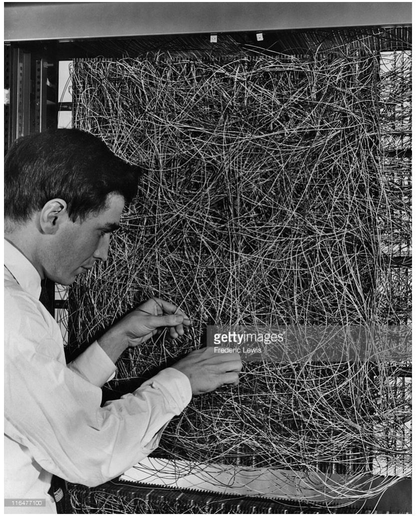

NY Times: "the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

Frank Rosenblatt

1960

ELIZA

"Natural language" conversation.

"...sister..." → "They me more about your family."

1966

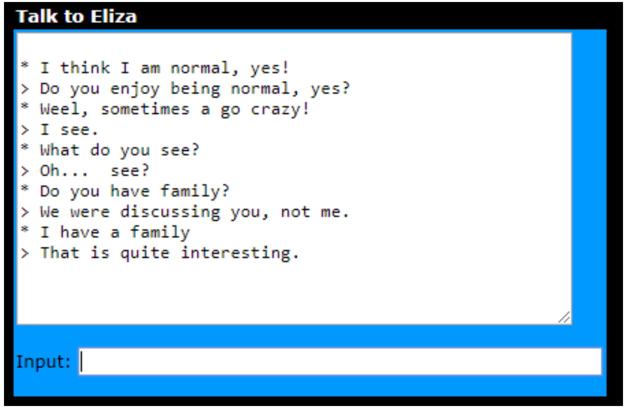

Perceptrons deprecated!

"[…] by the mid-1960s there had been a great many experiments with perceptrons, but no one had been able to explain why they were able to recognize certain kinds of patterns and not others."

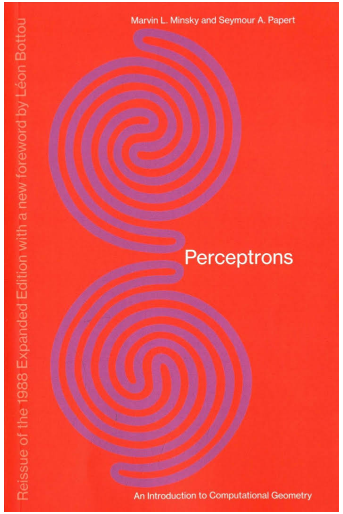

1969

siblings(X, Y) :– parent(Z, X) <u>and</u> parent(Z, Y),

What Computers Can't Do

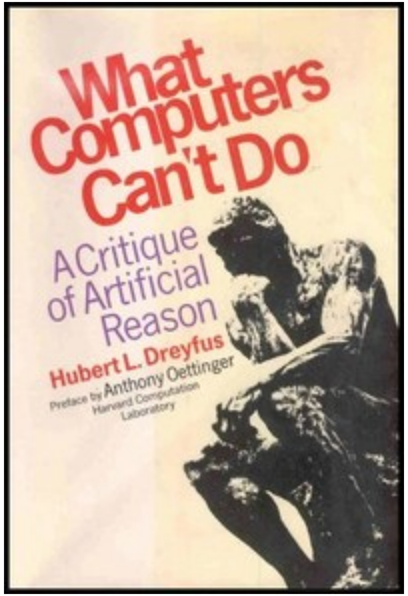

1972

DARPA cuts AI funding

1974

Expert Systems

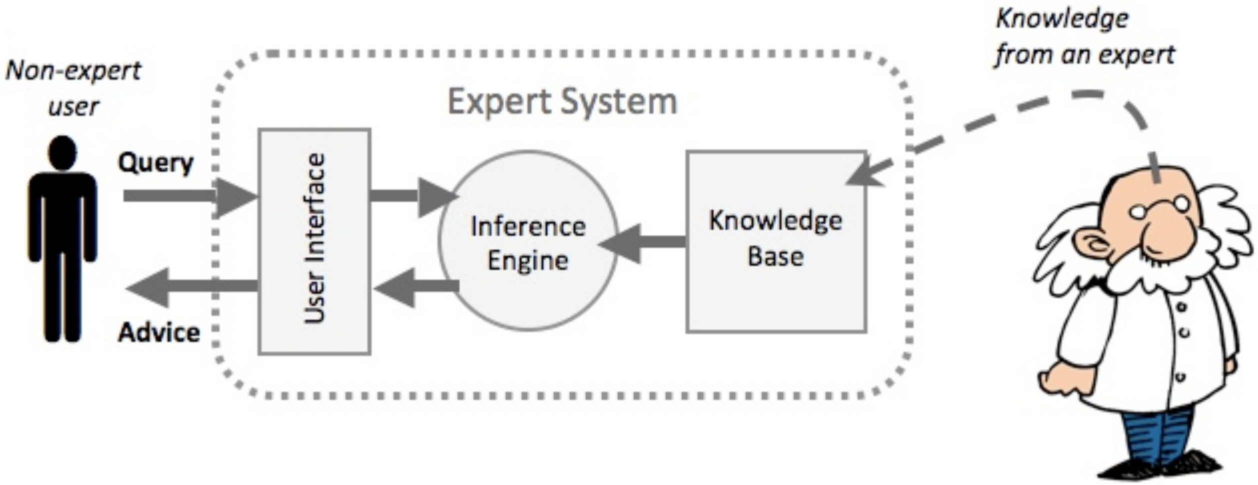

1980s

Cyc Project

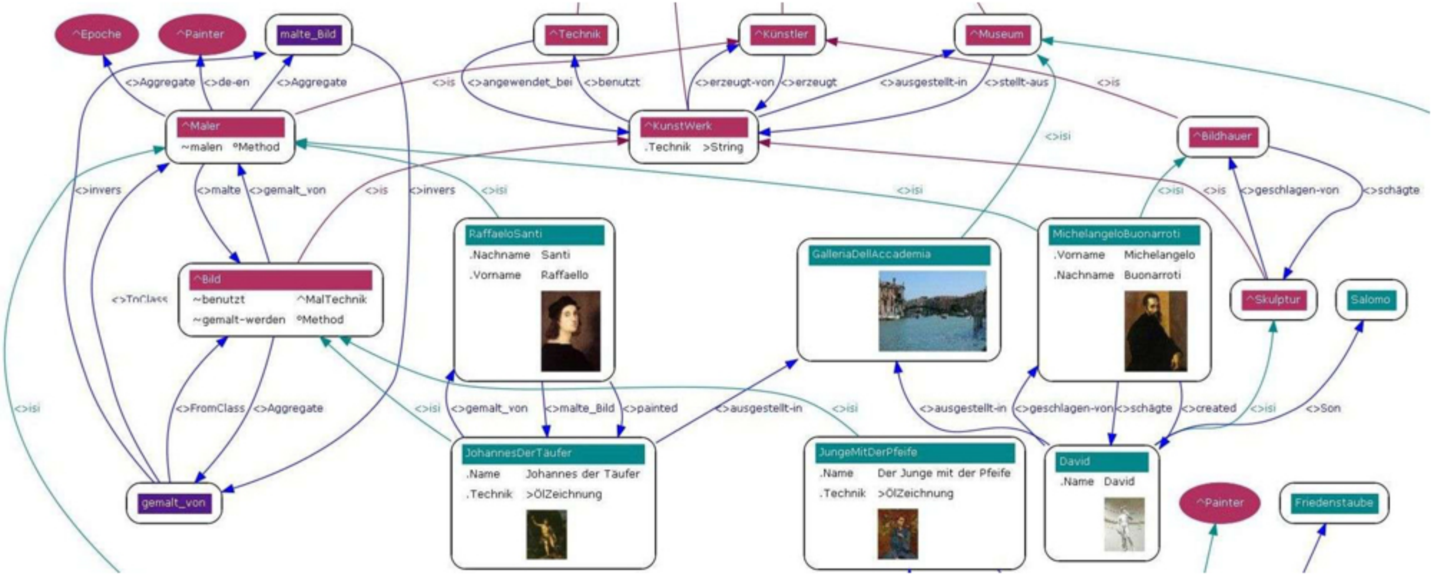

1984

Backpropagation

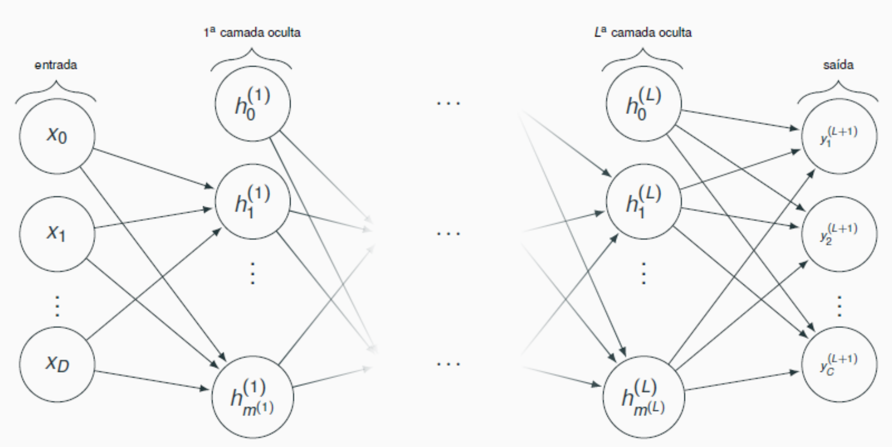

Learning representations by back-propagating errors

David E. Rumelhart*, Geoffrey E. Hinton† & Ronald J. Williams*

* Institute for Cognitive Science, C-015, University of California, San Diego, La Jolla, California 92093, USA
† Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Philadelphia 15213, USA

1986

Moravek’s paradox

- "it is comparatively easy to make computers exhibit […] intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility."

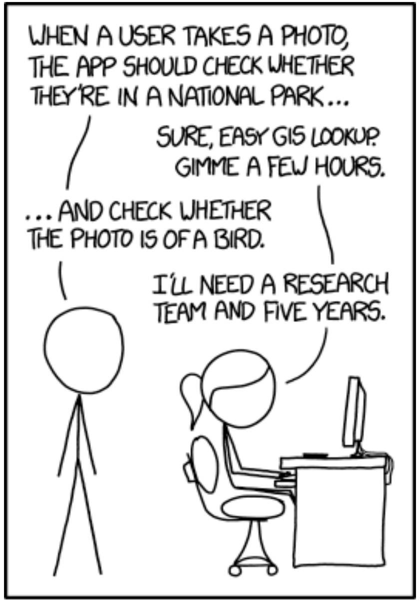

IN CS, IT CAN BE HARD TO EXPLAIN THE DIFFERENCE BETWEEN THE EASY AND THE VIRTUALLY IMPOSSIBLE.

1988

LeNet

Yan Lecun

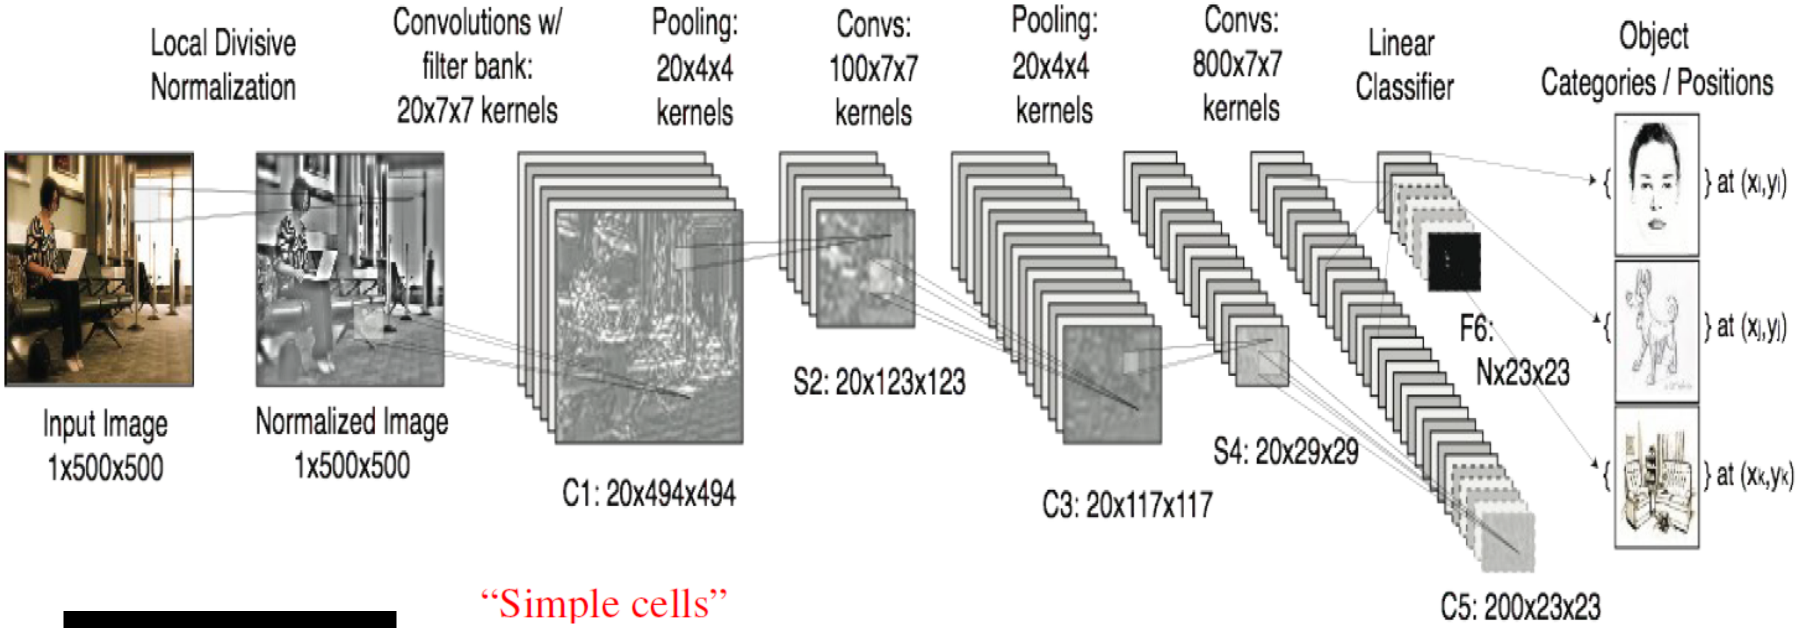

"Simple cells"

1989

Yan Le Cun et al, Convolutional Networks and Applications in Vision (2010)

Support Vector Machines

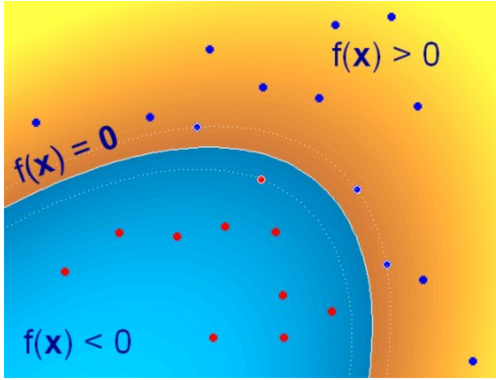

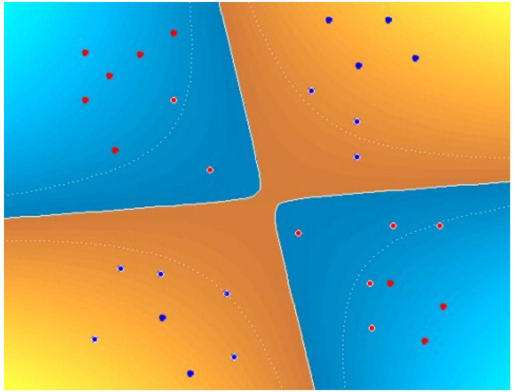

Isabelle Guyon

Vladimir Vapnik

1991

Larry Jackel vs Vladimir Vapnik

1. Jackel bets (one fancy dinner) that by March 14, 2000, people will understand
quantitatively why big neural nets working on large databases are not so bad.
(Understanding means that there will be clear conditions and bounds)

Vapnik bets (one fancy dinner) that Jackel is wrong.

But .. If Vapnik figures out the bounds and conditions, Vapnik still wins the bet.

****

2. Vapnik bets (one fancy dinner) that by March 14, 2005, no one in his right mind will use neural nets that are essentially like those used in 1995.

Jackel bets (one fancy dinner) that Vapnik is wrong

1995

Deep Blue

IBM 704 (chess): 50000 computations/second
Deep Blue: 50 bilions computations/second!

VS

1997

Dedicated hardware, minimax algorithm with alpha-beta cut, evaluation function: 8000 features!

Searched 200 millions configurations per second; secret methods to extend the search to depth 40.

LSTM Neural Nets

"if you can understant the paper, you are better than many people in machine learning. It took 10 years until people understand what they were talkink about". – Jeff Hinton

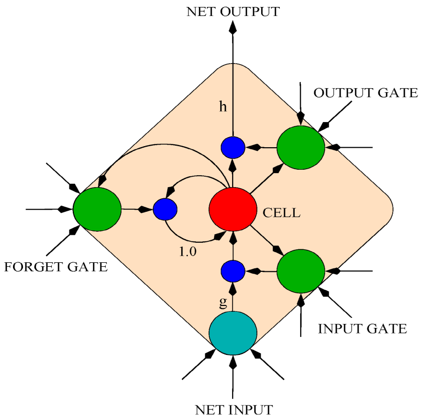

Sepp Hochreiter

https://scholar.google.de

Juergen Schmidhuber

1997

Citações: 2012: 75; out/2016: 1027; nov/2017: 2846

1st DARPA Grand Challenge

2004

Mojave Desert (USA), 150-mile (240 km) route

IBM's Watson wins *Jeopardy!*

2011

Deep Learning explosion

2012-now

Also known as: ***Revenge of the ~~Sith~~ Neural Nets!***

Revenge of the ~~Sith~~ Neural Nets

Computer Vision

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

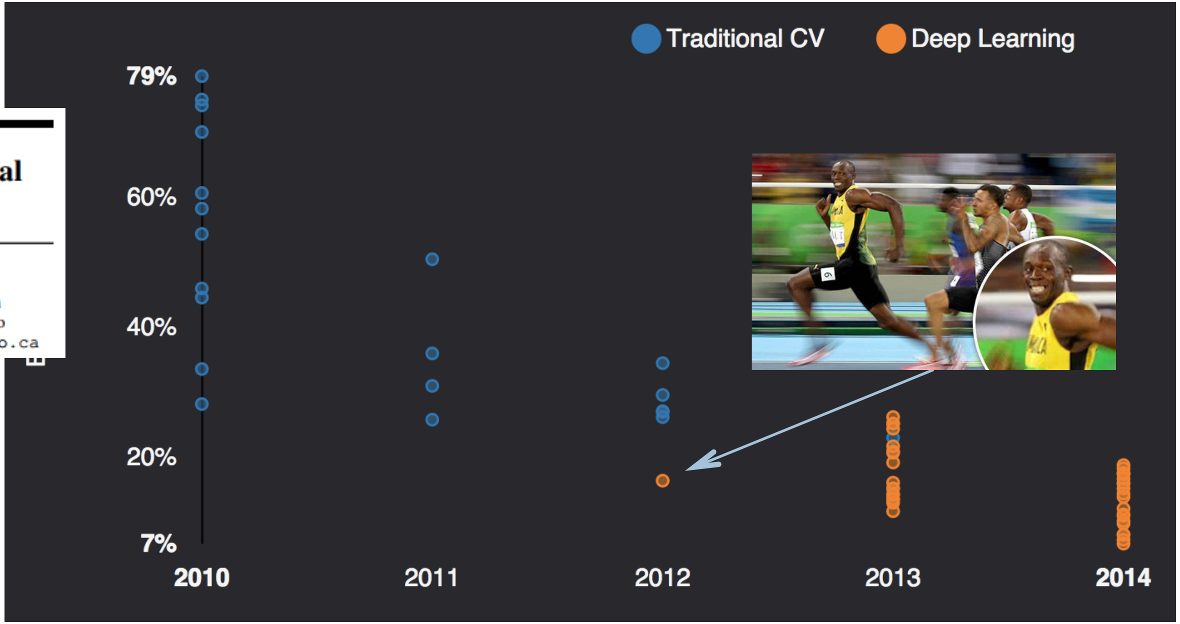

Credits: Mathew Zeiler (Clarifai)

Revenge of the ~~Sith~~ Neural Nets

Computer Vision
+
Text Processing

A close up of a hot dog on a bun.

A bath room with a toilet and a bath tub.

A vase filled with flower sitting on a table.

Long-term Recurrent Convolutional Networks for Visual Recognition and Description, 2016.

Revenge of the ~~Sith~~ Neural Nets

Deep Neural Networks for Acoustic Modeling in Speech Recognition

Geoffrey Hinton, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara Sainath, and Brian Kingsbury

Revenge of the ~~Sith~~ Neural Nets

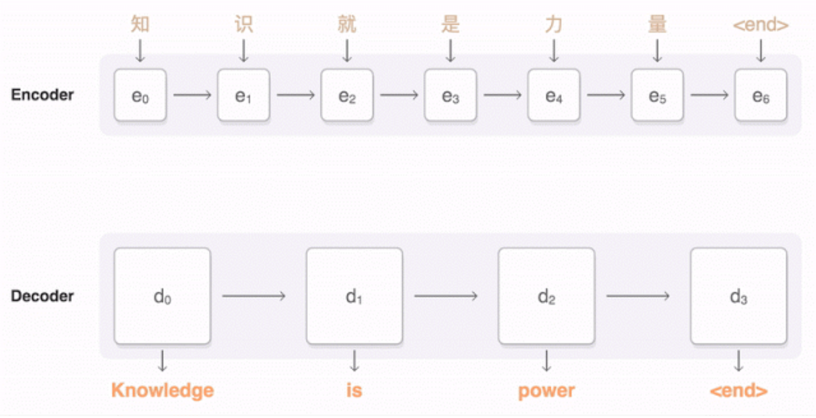

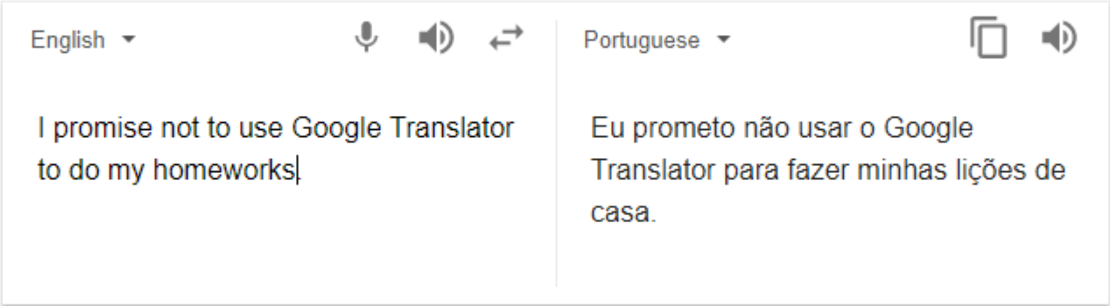

Revenge of the ~~Sith~~ Neural Nets

Computer Vision
+
Reinforcement Learning

Deep Q-Learning

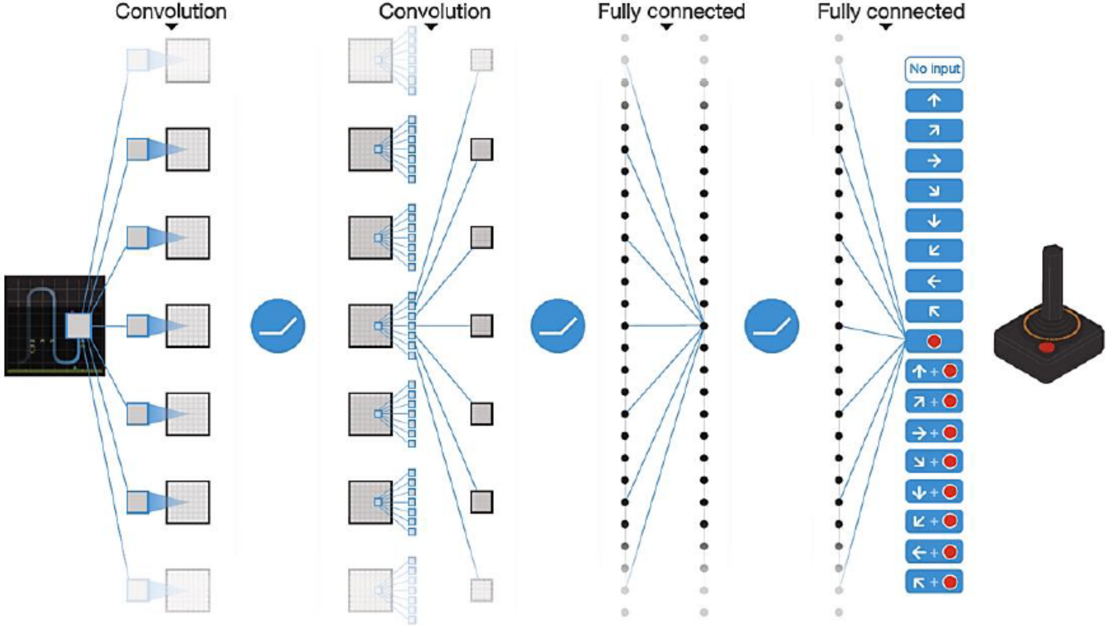

Deterministic Games: Go

Computer Vision
+
Reinforcement Learning

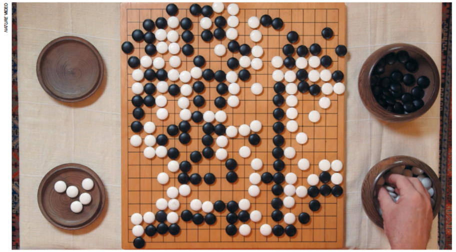

Go, a complex game popular in Asia, has frustrated the efforts of artificial-intelligence researchers for decades.

ARTIFICIAL INTELLIGENCE

Google masters Go

Deep-learning software excels at complex ancient board game.

Deep Learning: success factors

- Big Data (e.g, MNIST ~ 70k; ImageNet ~ 10^6)
- Hadware improvements
- Crowdsourcing

"What was wrong in the 80's is that we didn't have enough data and we didn't have enough computer power"

Geoffrey Hinton

Some trends

- Multimodal learning (networks trained with text + images, audio + video, etc...)
- Atention models
- Modularization – reuse and composition of models
- Deep Q-Learning

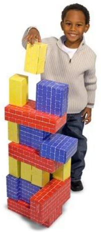

Open Problems

What is difficult is easy, and vice-versa

- […] hard problems are easy and the easy problems are hard. The mental abilities of a four-year-old – recognizing a face, lifting a pencil, walking across a room, answering a question – in fact solve some of the hardest engineering problems [...], it will be the stock analysts and petrochemical engineers and parole board members who are in danger of being replaced by machines. The gardeners, receptionists, and cooks are secure in their jobs for decades to come.

Steven Pinker

Natural Language Processing

- Headlines:
  - Enraged Cow Injures Farmer With Ax
  - Hospitals Are Sued by 7 Foot Doctors
  - Ban on Nude Dancing on Governor's Desk
  - Iraqi Head Seeks Arms
  - Local HS Dropouts Cut in Half
  - Juvenile Court to Try Shooting Defendant
  - Stolen Painting Found by Tree
  - Kids Make Nutritious Snacks

- Why are these funny?

Source: CS188

Reinforcement Learning

Common Sense Knowledge

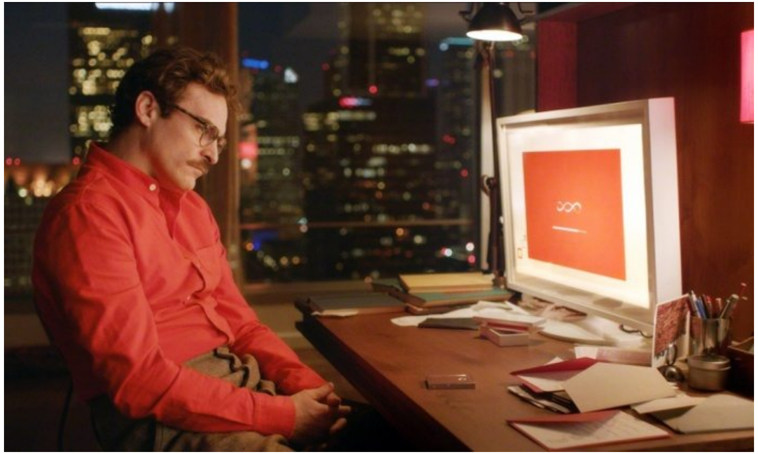

"If a mother has a son, then the son is younger than the mother and remains younger for his entire life."

"If President Trump is in Washington, then his left foot is also in Washington,"

“There’ll be a lot of people who argue against it, who say you can’t capture a thought like that. But there’s no reason why not. I think you can capture a thought by a vector.” – Geoff Hinton

Supervised Learning

Unsupervised Learning

Perhaps, the great frontier to be reached!

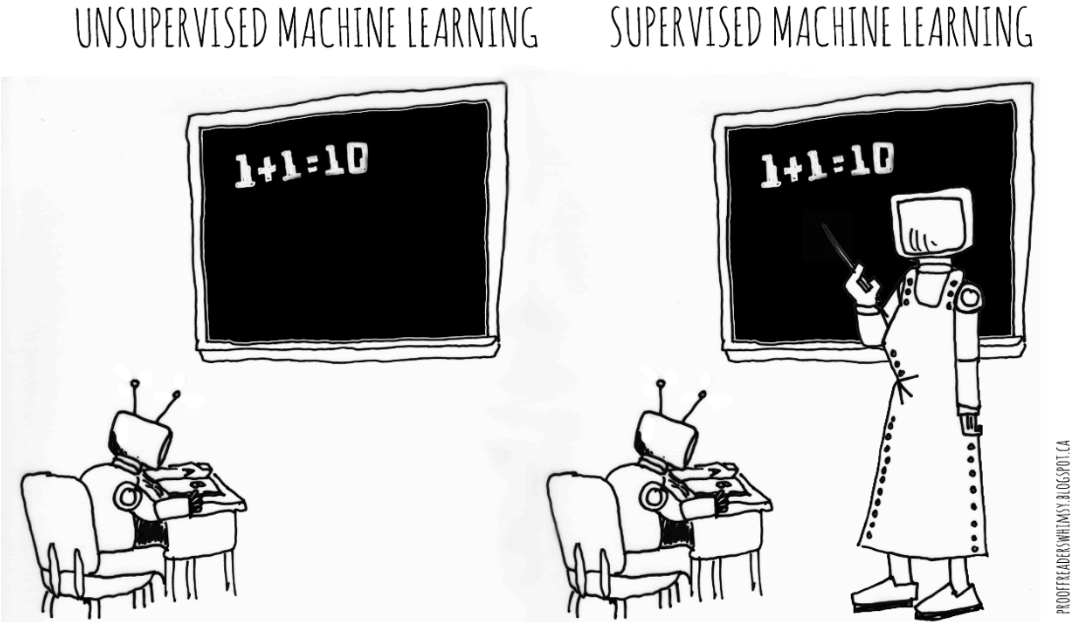

INTELIGÊNCIA ARTIFICIAL
HISTÓRICO, AVANÇOS E PROBLEMAS EM ABERTO

OBRIGADO!

Eduardo Bezerra (ebezerra@cefet-rj.br)