




Artificial Intelligence: Machine Learning and Pattern Recognition



University of Venice, Italy

a.a. 2015/16

Prof. Marcello Pelillo



What is Artificial Intelligence (AI)?

There is no universally accepted definition of Artificial Intelligence. A tentative one is the following:

AI is the endeavor of building an intelligent artifact

But... what is “intelligence”?

Some definitions:

- ✓ It is the ability to learn (Buckingham, 1921)
- ✓ This faculty is judgment, otherwise called good sense, practical sense, initiative, the faculty of adapting one's self to circumstances (Binet and Simon, 1961)
- ✓ It is the ability to perform well in an intelligence test (Boring, 1961)




Alan Turing's Proposal





The Turing Test

In 1950, Alan M. Turing proposed an operational definition of intelligence (the “Turing test”).





An Imaginary Dialogue

Q: Please write me a sonnet on the subject of the Forth Bridge.

A : Count me out on this one. I never could write poetry.

Q: Add 34957 to 70764.

A: (Pause about 30 seconds and then give as answer) 105621.

Q: Do you play chess?

A: Yes.

Q: I have K at my K1, and no other pieces. You have only K at K6 and R at R1. It is your move. What do you play?

A: (After a pause of 15 seconds) R-R8 mate.



Passing the Turing Test

To pass the test a machine must possess the following skills:

Natural language processing

to interact with the interrogator

Knowledge representation

to memorize information before and during the dialogue

Automatic reasoning

to use the acquired knowledge to answer the question and draw conclusions

Learning

to adapt to new situations



The “Total” Turing Test

The machine can access an audio/video feed so that the interrogator can test its perception skills; further, the interrogator can pass objects to be manipulated.

This requires:

Perception


to analyze and comprehend images and sounds)

Robotics

to manipulate objects and navigate



An Interdisciplinary Endeavor





Two Approaches to AI

Symbolic (declarativism)



Deals with:
Theorem proving,
problem solving,
games, reasoning, etc.



Psychology



Serial systems

**Sub-symbolic
(non-declarativism)**



Deals with:
Pattern recognition,
perception, learning,




Neurobiology



Parallel systems



Some History





Early Attempts (1943-1956)

1943: McCulloch and Pitts propose a model for an artificial neuron and analyze its properties

1949: Donald Hebb proposes a learning mechanism in the brain, still of great interest

1950-53: Shannon and Turing work (independently) on chess-playing programs

1951: Minsky and Edmonds develop the first “neural” computer

1956: *Newell e Simon develop the “Logic Theorist”*



Hanover, 1956: The Birth of AI

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C. E. Shannon, Bell Telephone Laboratories

August 31, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. 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. 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. [...]



First successes...

1961: Newell and Simon develop General Problem Solver (GPS)

1952-: Samuel develops a checker playing game

1957: First attempts at automatic translation

1958: McCarthy invents LISP

1963 - : Minsky and students study problems on micro-worlds (es., ANALOGY, SHRDLU)

1962: Rosenblatt develops the Perceptron, a neural net that learns from examples



... and first failures

1966: Financing to “automatic translation” projects in the USA is canceled

1969: Minsky and Papert publish Perceptrons, where they show that the Rosenblatt model cannot solve some very simple problems

1971-72: Cook and Karp develop the computational complexity theory, showing that a lot of problems are “intractable” (NP-complete).



The Expert-System Boom

1969: Feigenbaum *et al.* (Stanford) develop DENDRAL, an ES for making predictions on molecular structures

1976: MYCIN, an ES with some 450 rules for the diagnosis of infectious diseases

1979: PROSPECTOR, an ES for mineral explorations

1982: R1, a commercial ES for configuring DEC VAX systems



The Resurgence of Neural Networks

1982: Hopfield (Caltech) develops a neural model based on the analogy with physical (ferromagnetic) systems

1985: Hopfield and Tank applied their model to “solve” intractable (NP-complete) problems

1986: The PDP group (re)introduces *back-propagation*, a learning algorithm for layered (feed-forward) neural networks, thereby overcoming the limitation of *Perceptrons*

1987: Sejnowski and Rosenberg develop *NETtalk*, a neural network that “learns” to talk

...

Today: “Deep learning” is the hottest topic in machine learning



NETtalk:

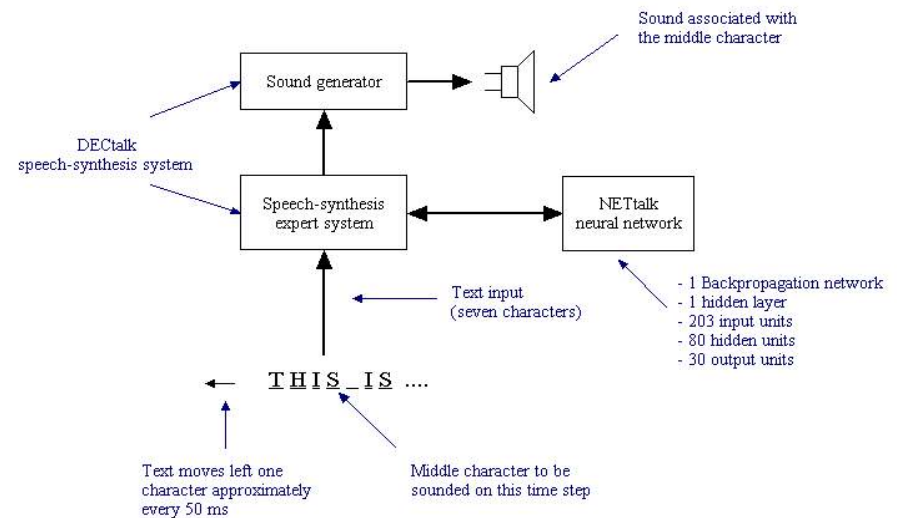

A Neural Net that Learns to Talk



T. J. Sejnowski and C. R. Rosenberg




“Parallel networks that learn to pronounce English text”

Complex Systems 1, 145-168 (1987)





Far away from HAL 9000 & Co., but...





IBM: Deep Blue vs. Kasparov (1997)





SONY and the “humanoids”




Robot SDR-4X II





EMI, *Experiments in Musical Intelligence*





Biometry

Biometry develops techniques for the automatic recognition of a person's identity.

Typical biometric information

Physiological:

- Face
- Fingerprints
- Voice
- Retina
- Iris
- Hands
- DNA
- ...

Behavioral:

- Signature
- Keystroke
- Gait
- ...



Biometry e anti-terrorism

« A comprehensive immigration reform must include a better system for verifying documents and work eligibility. A key part of that system should be a new identification card for every legal foreign worker.

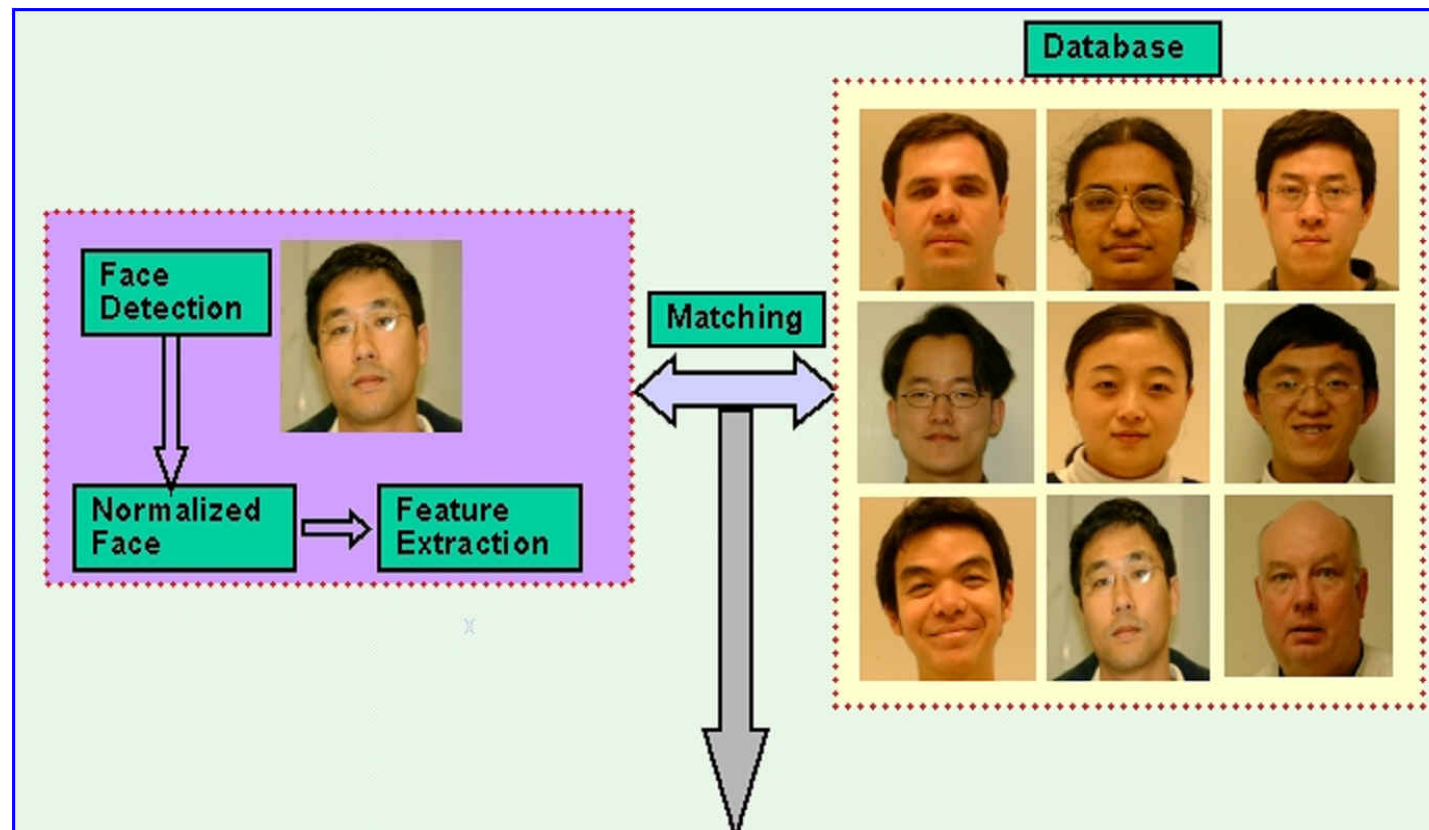
This card should use **biometric technology**. »



George W. Bush
May 15, 2006




Face Recognition






Face Recognition







Fingerprint Recognition







Winning at *Jeopardy!*





Assisting Car Drivers





Assisting Visually Impaired People





e-commerce

amazon.com
and you're done.™

ebay®

BARNES & NOBLE
BOOKSELLERS

MOOIE
GALLERY®
play on™

iLike Beta / Artists / eighte8
Sign in or Create your own profile

HOLLYWOOD
VIDEO®

NETFLIX®


bol.it
go

ibs.it internet bookshop



Recommended for You


Marcello Pelillo, [Amazon.com](https://www.amazon.com) has new recommendations for you based on [items](#) you purchased or told us you own.




[Elements of
Information Theory
\(Wiley Series in
Telecommunications
and Signal
Processing\)](#)




[Action in
Perception
\(Representation
and Mind\)](#)



[Topological
Vector Spaces,
Distributions
and Kernels](#)



[Applied
Exterior
Calculus
\(Dover Books
on
Mathematics\)](#)



[The Malliavin
Calculus \(Dover
Books on
Mathematics\)](#)



[Copernican
Questions: A
Concise
Invitation to
the Philosophy
of Science](#)



[Optimal
Filtering
\(Dover Books
on
Engineering\)](#)



[Riemannian
Geometry](#)

► [See More Recommendations](#)



Other Commercial Applications of AI

- Video-surveillance
- Traffic monitoring
- Plate recognition
- Road sign recognition
- Speech synthesis and recognition
- Web profiling
- Medical image analysis
- Virtual reality
- Man-machine interaction (e.g., gesture recognition)
- Expert systems
-



This Course





Topics Covered

Information theory and inference: source coding, channel coding.

Learning and inference in neural networks: feed-forward networks, deep learning architectures, Hopfield networks and related models.

Unsupervised and semi-supervised learning: K-means, spectral clustering, dominant sets, game-theoretic models.



Recommended Textbooks

D. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

S. Russell and P. Norving. *Artificial Intelligence: A Modern Approach* (2nd edition)

(trad it., *Intelligenza Artificiale: Un approccio moderno*)

C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, Springer 2007.

