

Artificial Intelligence and the Singularity

piero scaruffi




www.scaruffi.com

October 2014 - Revised 2016

*"The person who says it cannot be
done should not interrupt the person
doing it" (Chinese proverb)*

Piero Scaruffi




- piero scaruffi
p@scaruffi.com
scaruffi@stanford.edu




Olivetti AI Center, 1987

Piero Scaruffi

- Cultural Historian
- Cognitive Scientist
- Blogger
- Poet
- www.scaruffi.com



www.scaruffi.com



This is Part 1

- See <http://www.scaruffi.com/singular> for the index of this Powerpoint presentation and links to the other parts
 1. Classic A.I. - The Age of Expert Systems
 2. The A.I. Winter and the Return of Connectionism
 3. Theory: Knowledge-based Systems and Neural Networks
 4. Robots
 5. Bionics
 6. Singularity
 7. Critique
 8. The Future
 9. Applications
 10. Machine Art
 11. The Age of Deep Learning
 12. Natural Language Processing

Classic A.I.

The Age of Expert Systems

A Brief History of Logic

George Boole's "The Laws Of Thought" (1854): the laws of logic “are” the laws of thought

Propositional logic and predicate logic: true/false!

$F \vee G = G \vee F;$ $F \wedge G = G \wedge F.$ **commutativity**

$(F \vee G) \vee H = F \vee (G \vee H);$ $(F \wedge G) \wedge H = F \wedge (G \wedge H).$ **associativity**

$F \vee (G \wedge H) = (F \vee G) \wedge (F \vee H).$ **distributivity**

$F \wedge (G \vee H) = (F \wedge G) \vee (F \wedge H).$ **distributivity**

$\neg (F \vee G) = \neg F \wedge \neg G;$ $\neg (F \wedge G) = \neg F \vee \neg G.$ **De Morgan**

connectives = $\{\supset, \neg, \vee, \wedge, \equiv\}$

quantifiers = $\{\forall, \exists\}$

$(\exists x)\text{Tet}(x)$

$(\forall x)(\text{Tet}(x) \rightarrow \text{Large}(x))$

$((\forall x)(\text{Tet}(x) \rightarrow \text{Large}(x)) \rightarrow (\text{Tet}(x) \rightarrow \text{Large}(x)))$

$\text{Tet}(x) \rightarrow \text{Large}(x)$

$\text{Large}(x) \rightarrow (\exists x)\text{Large}(x)$

$\text{Tet}(x) \rightarrow (\exists x)\text{Large}(x)$

$(\exists x)\text{Tet}(x) \rightarrow (\exists x)\text{Large}(x)$

$(\exists x)\text{Large}(x)$

A Brief History of Logic

Axiomatization of Thought:

Gottlob Frege's "Foundations of Arithmetic"
(1884)

Giuseppe Peano's "Arithmetices Principia
Nova Methodo Exposita" (1889)

Bertrand Russell's "Principia Mathematica"
(1903)



A Brief History of Logic

- David Hilbert (1928)
 - Entscheidungsproblem problem: the mechanical procedure for proving mathematical theorems
 - An algorithm, not a formula
 - Mathematics = blind manipulation of symbols
 - Formal system = a set of axioms and a set of inference rules



The Cultural Context


- 1910-1950 Everything changed:
 - Everyday life
 - The foundations of science
 - The concept of art
 - The geopolitical order

The 1910s


- Electricity
- Refrigerator
- Automobile
- Airplane
- Telegraph
- Telephone
- Phonograph
- Camera
- Cinema
- Radio
- Typewriter
- Calculator
- Skyscraper
- Plastic



Everything
is relative




You are a
formula



You are and
you are not



You are just
a reflex




You are a
probability



Everything is
moving away
from you

Everything is
uncertain






*The
emancipation of
the dissonance*



*History is a
nightmare from
which I am trying
to awake.*



There will always
be something you
cannot prove


Your mind
creates reality

Truth is an
opinion



Life and machines
obey the same
laws of nature

Mind is a
symbol
processor




Everything
comes from just
one point

Everything is
information




Cultural Context

- Bottom line:
 - Nonconformism
 - Anxiety
 - Noise
 - Freedom



Cultural Context

- World War II (1939-45)
- The Holocaust
- Hiroshima
- Disintegration of the British Empire
- Rise of the USA and Soviet Union



Alan Turing


- Hilbert's challenge (1928): an algorithm capable of solving all the mathematical problems
- Turing Machine (1936): a machine whose behavior is determined by a sequence of symbols and whose behavior determines the sequence of symbols
- A universal Turing machine (UTM) is a Turing machine that can simulate an arbitrary Turing machine




$\forall M[M \in \mathcal{M}] \Leftrightarrow$
 $M: (S_M, I_M, O_M: S_M \times I_M \rightarrow I_M, N_M: S_M \times I_M \rightarrow S_M, D_M: S_M \times I_M \rightarrow d)$
 $\mathcal{N} = \{0, \dots, \infty\}$ (the "natural" numbers)
 $\mathcal{J} = \{1, \dots, \infty\}$ (the positive "integers")
 $S_M = \{s_0, \dots, s_n\}, n \in \mathcal{J}$ ($\mathcal{M}$ states)
 $I_M = \{i_0, \dots, i_j\}, j \in \mathcal{J}$ ($\mathcal{M}$ tape symbols)
 where
 $d = \{-1, 0, +1\}$ ($\mathcal{M}$ head motions)
 $\delta: \mathcal{N} \rightarrow S_M$ ($\mathcal{M}$ state over time)
 $\square_M: \mathcal{N} \times \mathcal{N} \rightarrow I_M$ ($\mathcal{M}$ tape contents over time)
 $P_M: \mathcal{N} \rightarrow \mathcal{N}$ (current $\mathcal{M}$ cell at each time)
 Box 1
 $\forall M \forall V(M, V) \in \mathcal{V} \Leftrightarrow$
 $[V \subset I^*] \text{ and } [M \in \mathcal{M}] \text{ and } \forall v \in V \forall H \forall t, j \in \mathcal{N}$
 $[P_t = j] \text{ and } [S_t = S_0] \text{ and } [\square_{t,j}, \dots, \square_{t,j+|v|-1}] = v \Rightarrow$
 $\exists v' \in V, \exists t', t'', i', j' \in \mathcal{N} \text{ and } t' > t$
 (1) $[(j' + |v'| \leq j) \text{ or } [(j + |v|) \leq j']] \text{ and}$
 (2) $[\square_{t',j'}, \dots, \square_{t',j'+|v'|-1}] = v' \text{ and}$
 (3) $[\exists t'' [t < t'' < t'] \text{ and } [P_{t''} \in j', \dots, j' + |v'|-1]]$
 Box 2

Alan Turing

- Alan Turing (1936)
 - Universal Turing Machine: a Turing machine that is able to simulate any other Turing machine
 - The universal Turing machine reads the description of the specific Turing machine to be simulated



Turing Machine



Alan Turing

“It is of fundamental importance for the character of this problem that only mechanical calculations according to given instructions [...] are admitted as tools for the proof.”

(BTW, the halting problem is undecidable, i.e. Hilbert's Entscheidungsproblem is impossible)

Alan Turing

- Turing machines in nature: the ribosome, which translates RNA into proteins
 - Genetic alphabet: nucleotides ("bases"): A, C, G, U
 - The bases are combined in groups of 3 to form "codons"
 - RNA is composed of a string of nucleotides ("bases") according to certain rules
 - There are special carrier molecules ("tRNA") that are attached to specific aminoacids (proteins)
 - The start codon encodes the aminoacid Methionine
 - A codon is matched with a specific tRNA
 - The new aminoacid is attached to the protein
 - The tape then advances 3 bases to the next codon, and the process repeats
 - The protein keeps growing
 - When the "stop" codon is encountered, the ribosome dissociates from the mRNA

Alan Turing

- World War II:
 - Breaking the Enigma code (Bombe)
 - Turing worked at Bletchley Park where the Colossus was built but it was not a universal Turing machine (not general purpose)



Replica of the Bombe



The Turing Century

- Can you name any achievement of the last 50 years (from the Moon landing to animal cloning) that would have happened even without programmable computers?

The Turing Century



- Are thinking machines possible?



1934: 17 years before the first commercial computer


Electronic Brains

- 1941: Konrad Zuse's Z3 programmable electromechanical computer, the first Turing-complete machine
- 1943: Tommy Flowers and others build the Colossus, the world's first programmable digital electronic computer



Electronic Brains

- 1944: Howard Aiken of IBM unveils the first computer programmed by punched paper tape, the electromechanical Harvard Mark I




- 1945: John Von Neumann, John Mauchly and Presper Eckert design a computer that holds its own instructions, the "stored-program architecture"



Electronic Brains


1945: John Von Neumann's computer architecture

- Separation of instructions and data (although both are sequences of 0s and 1s)
- Sequential processing



Control unit:


- reads an instruction from memory
- interprets/executes the instruction
- signals the other components what to do




Don Fussell

Electronic Brains


1946: The first non-military computer, ENIAC, or "Electronic Numerical Integrator and Computer", is unveiled, built by John Mauchly and Presper Eckert at the University of Pennsylvania



Electronic Brains



The first book on electronic computers (1949)



A GIANT "thinking machine," able to apply electronic speeds for the first time to mathematical tasks too difficult for previous solution, is now in use by the U. S. Army. It can compute 1,000 times faster than the most advanced calculating device previously known, can solve in hours problems which would take years on a mechanical machine and shows promise of revolutionizing long-range

weather prediction and many other highly complicated sciences.

Known as the ENIAC (Electronic Numerical Integrator and Computer), the revolutionary robot weighs 30 tons and occupies a room 30 by 50 feet in size. Temporarily housed at the University of Pennsylvania, the machine's permanent home will be the Army Ordnance Proving Grounds at Aberdeen, Md.

The first all-electronic, general purpose computer ever developed, it was designed and constructed for Army Ordnance at the Moore School of Electrical Engineering, University of Pennsylvania. The electronic robot is capable of solving thousands of technical and scientific problems so complex and difficult that all previous methods of solution were considered impractical, if not impossible.

Overall view of the ENIAC being set up for a complicated math problem.



"thinking" for Canadian business ...
the **UNIVAC 120**

Does the UNIVAC 120 really think?

among other things, the most complicated payroll, production and billing problems.

Does Univac 120 eliminate intelligence? Certainly not. But it does free your skilled personnel for creative thinking. Wherever you may find clerical chores that

New N.P.L. Wonder

ELECTRIC BRAIN TO BE MADE AT TEDDINGTON

34 YEARS-OLD DESIGNER TALKS TO SURREY COMET

£100,000 A.C.E. WILL BE OBSOLETE BEFORE COMPLETED

SOME of the feats that will be able to be performed by Britain's new electronic brain, which is being developed at the N.P.L., Teddington, were described to the SURREY COMET yesterday by Dr. A. M. Turing, 34 year-old mathematics expert, who is pioneer of the scheme in this country.

The machine is to be an improvement on the American ENIAC, and it was in the brain of Dr. Turing that the more efficient model was developed.

Electronic Brains

- Apr 1949: The Manchester Mark 1, the first stored-program electronic computer
- May 1949: Cambridge's EDSAC, the second stored-program electronic computer
- Aug 1949: Philadelphia's EDVAC, the third stored-program electronic computer
- 1950: The Pilot ACE computer

New N.P.L. Wonder

ELECTRIC BRAIN TO BE MADE AT TEDDINGTON

**34 YEARS-OLD DESIGNER TALKS
TO SURREY COMET**



**£100,000 A.C.E. WILL BE OBSOLETE
BEFORE COMPLETED**

SOME of the feats that will be able to be performed by Britain's new electronic brain, which is being developed at the N.P.L., Teddington, were described to the SURREY COMET yesterday by Dr. A. M. Turing, 34 years-old mathematics expert, who is pioneer of the scheme in this country.

The machine is to be an improvement on the American ENIAC, and it was in the brain of Dr. Turing that the more efficient model was developed.


Electronic Brains

- May 1950: The first stored-program electronic computer to be deployed in the USA, the SEAC, and the first to use semiconductors instead of vacuum tubes
- Feb 1951: The Ferranti Mark 1, the first commercial computer, an evolution of the EDSAC
- 1952: A Univac 1 correctly predicts that Eisenhower would win the elections



Electronic Brains


- Goldstine and Eckert with the electronics needed to store a single decimal digit



(Computer History Museum, Mountain View)

Electronic Brains

Computer programmers of 1951: Patsy Simmers (holding an ENIAC board) Gail Taylor (holding an EDVAC board), Milly Beck (holding an ORDVAC board), Norma Stec (holding a BRLESC-I board)




Electronic Brains


1954: IBM's first “mass-produced” computer, the 650
(1,800 units sold - \$200-400,000 each)



Electronic Brains



Documentary of 1958



Science-fiction tales of 1954


Electronic Brains

- USA/ Semiconductors
 - 1947: AT&T's Bell Labs invent the transistor (William Shockley, John Bardeen, Walter Brattain)
 - 1949: The USA files an antitrust lawsuit against AT&T
 - 1952: AT&T's symposium on the transistor, open to everybody
 - 1954: Texas Instruments introduces the first commercial transistor
 - 1954: The first transistor radio (“Regency”)




Electronic Brains

- USA/ Semiconductors
 - 1961: Texas Instruments introduces the first commercial integrated circuit
 - Military and space applications use the integrated circuit
 - 1965: Gordon Moore predicts that the processing power of computers will double every 18 months
 - 1971: Intel invents the microprocessor
 - Universities are irrelevant in semiconductor progress because the manufacturing process is too costly
 - Universities are crucial for progress in computers



Jack Kilby's I.C.



Intel 4004

r

r



Electronic Brains

Software

- 1958: Jim Backus (at IBM) invents the FORTRAN programming language, the first machine-independent language
- 1964: IBM introduces the first "operating system" for computers (the OS/360)
- 1968: The Arpanet is established based on Baran's idea (four nodes: UCLA, Stanford Research Institute, UCSB, University of Utah)
- 1969: the Unix operating system is born

Electronic Brains

Democratizing technology


- Antitrust policies contribute to the rapid diffusion of intellectual property throughout the computer and semiconductor industries
- 1956: IBM and AT&T settle antitrust suits by licensing their technologies to competitors
- 1969: The “unbundling” of software by IBM creates the software industry

Cybernetics

The Steam Engine

- Biggest impact on daily life since the printing press
- Inventors are ordinary people, not academics
- The automation of manufacturing begins in Lancashire, not at a university

James Watt (1776)



Cybernetics

- Mathematician Norbert Wiener, physiologist Arturo Rosenblueth and engineer Julian Bigelow: "*Behavior, Purpose and Teleology*" (1943)

Behavior, Purpose and Teleology

Arturo Rosenblueth, Norbert Wiener and Julian Bigelow

in: *Philosophy of Science*, 10(1943), S. 18–24.

This essay has two goals. The first is to define the behavioristic study of natural events and to classify behavior. The second is to stress the importance of the concept of purpose.

Given any object, relatively abstracted from its surroundings for study, the behavioristic approach consists in the examination of the output of the object and of the relations of this output to the input. By output is meant any change produced in the surroundings by the object. By input, conversely, is meant any

object is the source of the output energy involved in a given specific reaction. The object may store energy supplied by a remote or relatively immediate input, but the input does not energize the output directly. In passive behavior, on the contrary, the object is not a source of energy; all the energy in the output can be traced to the immediate input (e.g., the throwing of an object), or else the object may control energy which remains external to it throughout the reaction (e.g., the soaring flight of a bird).




Cybernetics



- Warren McCulloch's and Walter Pitts' binary neuron (1943)

McCulloch-Pitts model of an artificial neuron



$$y = f(W \cdot p + b)$$


$p = (p_1, \dots, p_R)^T$ is the input column-vector

$W = (w_1, \dots, w_R)$ is the weight row-vector


Emil M. Petriu

Some transfer functions "f"


Hard Limit: $y = 0$ if $z < 0$
 $y = 1$ if $z \geq 0$




Symmetrical Hard Limit: $y = -1$ if $z < 0$
 $y = +1$ if $z \geq 0$



Log-Sigmoid:
 $y = 1/(1 + e^{-z})$



Linear:
 $y = z$



Cybernetics

- Macy Conference on Cybernetics (March 1946, New York)
 - John von Neumann (computer science)
 - Rafael Lorente de No (neurophysiology)
 - Norbert Wiener (mathematics)
 - Arturo Rosenblueth (physiology)
 - Warren McCulloch (neuropsychiatry)
 - Gregory Bateson (anthropology)
 - Margaret Mead (anthropology)
 - Walter Pitts (mathematics)
 - Ralph Gerard (neurophysiology)
 - Heinrich Kluever (psychology)
 - Lawrence Frank (sociology)
 - Molly Harrower (psychology)
 - Lawrence Kubie (psychoanalysis)
 - Filmer Northrop (philosophy)
 - Paul Lazarsfeld (sociology)



Cybernetics

Norbert Wiener (1947)

- Bridge between machines and nature, between "artificial" systems and natural systems
- Feedback, by sending back the output as input, helps control the proper functioning of the machine
- A control system is realized by a loop of action and feedback
- A control system is capable of achieving a "goal", is capable of "purposeful" behavior
- Living organisms are control systems





Computational Neuroscience

1947: Kacy Cole's voltage clamp technique

1952: Alan Hodgkin and Andrew Huxley: computational model of a spiking neuron

1962: Wilfrid Rall simulates a dendritic arbor

1963: Donald Perkel simulates the working of the neuron



The Turing Test

1950: Alan Turing's "Computing Machinery and Intelligence" (the "Turing Test")

A. M. Turing (1950) Computing Machinery and Intelligence. *Mind* 49: 433-460.

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 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 Turing Test




The Turing Test (1950)

- Hide a human in a room and a machine in another room and type them questions: if you cannot find out which one is which based on their answers, then the machine is intelligent



The Turing Test

The “Turing point”: a computer can be said to be intelligent if its answers are indistinguishable from the answers of a human being



Disembodied Intelligence

- What Turing and Wiener did
 - Removed the body from intelligence
 - Intelligence has to do with manipulating, transmitting, information
 - Intelligence is independent of the material substrate
 - They did not interpret machines as humans, but humans as (information-processing) machines
 - They moved humans closer to machines, not machines closer to humans

Fear of the Technocracy

Wright Mills: "The Power Elite" (1956)

Jacques Ellul: "The Technological Society" (1964)

Herbert Marcuse: "One-dimensional Man" (1964)

John Kenneth Galbraith: "The New Industrial State" (1967)

Lewis Mumford: "The Myth of the Machine" (1967)

Theodore Roszak: "The Making of a Counterculture" (1969)

Charles Reich: "The Greening of America" (1970)

Before Artificial Intelligence

- 1949: Warren Weaver's "Translation" memorandum
- 1950: Claude Shannon's "Programming a Computer for Playing Chess"
- 1951: AI programs at Manchester on the Ferranti Mark:
 - A draughts-playing program by Christopher Strachey
 - A chess-playing program by Dietrich Prinz

Before Artificial Intelligence

- 1951: Calculating Machines and Human Thought (Paris)
- 1952: First International Conference on Machine Translation organized by Yehoshua Bar-Hillel
- 1954: Demonstration of a machine-translation system by Leon Dostert's team at Georgetown University and Cuthbert Hurd's team at IBM, possibly the first non-numerical application of a digital computer



Wiener playing with Torres y Quevedo's automaton

Before Artificial Intelligence



- 1954: Wesley Clark and Belmont Farley build the first computer simulation of a neural network
- 1955: The Western Joint Computer Conference with papers by Newell, Selfridge, Clark, etc



1955 WESTERN JOINT COMPUTER CONFERENCE

Generalization of Pattern Recognition in a Self-Organizing System*

W. A. CLARK† AND B. G. FARLEY†



Before Artificial Intelligence


1954: Demonstration of a machine-translation system by Leon Dostert's team at Georgetown University and Cuthbert Hurd's team at IBM



Before Artificial Intelligence

Mathematical models of the brain (Britain)

- William Grey-Walter's Elmer and Elsie (1948)




- Ross Ashby's homeostat (1948)



- Alan Turing's "Intelligent Machinery" (1948)
- Ross Ashby's "Design for a Brain" (1952)

Before Artificial Intelligence

The Ratio Club (1949-55)



Before Artificial Intelligence

Mathematical models of the brain (Britain)

- Jack Allanson (1956)

Jack Allanson

The first Communist President of the National Union of Students.

- Raymond Beurle (1956)

PHILOSOPHICAL TRANSACTIONS
OF THE ROYAL SOCIETY B
BIOLOGICAL SCIENCES

**Properties of a Mass of Cells Capable of
Regenerating Pulses**

R. L. Beurle

Published 23 August 1956.

- Pete Uttley (1956)



Before Artificial Intelligence

Algorithms

- 1933: principal component analysis (Harold Hotelling)
- 1944: logistic regression (Joseph Berkson)
- 1950: probability theory (Andrei Kolmogorov)
- 1951: the "k-nearest-neighbors" classifier (Evelyn Fix and Joseph Hodges)
- 1951: stochastic gradient descent (Herbert Robbins)
- 1953: the first Markov Chain Monte Carlo method, the "Metropolis algorithm" (Marshall & Arianna Rosenbluth)
- 1957: Dynamic Programming" (Richard Bellman)
- 1957: k-means clustering (Stuart Lloyd)



Kolmogorov



Robbins



Rosenbluth



Teller



Metropolis



Arianna



Bellman

Before Artificial Intelligence

Algorithms

- Gradient methods

Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be twice continuously differentiable ($\mathcal{C}^2$), then

- The **gradient** of f is

$$\nabla f = \left[\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_N} \right]^T$$

- The **Hessian** of f is

$$\nabla^2 f = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_N} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_N \partial x_1} & \cdots & \frac{\partial^2 f}{\partial x_N^2} \end{bmatrix}$$

Batch gradient descent

$$\theta = \theta - \eta \cdot \nabla_{\theta} J(\theta).$$

Stochastic gradient descent

$$\theta = \theta - \eta \cdot \nabla_{\theta} J(\theta; x^{(i)}; y^{(i)}).$$



Cauchy 1847



Robbins 1951

Artificial Intelligence



AI Magazine Volume 27 Number 4 (2006) (© AAAI)


A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence

August 31, 1955

*John McCarthy, Marvin L. Minsky,
Nathaniel Rochester,
and Claude E. Shannon*

Artificial Intelligence?

- Definitions!
 - When does Computer Science become Artificial Intelligence?
 - When does Automation become AI?
 - When does technology become AI?
 - What is the difference between an algorithm and an AI algorithm?
 - What is “intelligence”?
 - What is “artificial”?



Artificial Intelligence

- A.I. = All the technologies needed for a machine to pass the Turing Test that could not de facto be implemented on a Von Neumann architecture: speech recognition, computer vision, natural language processing, reasoning, learning, common sense...

Artificial Intelligence

- What is Human Intelligence?
- What is Animal Intelligence?
- What is Machine Intelligence?
- How do we compare Human, Animal and Machine intelligences?

Universal Intelligence: A Definition of Machine Intelligence

Shane Legg

IDSIA, Galleria 2, Manno-Lugano CH-6928, Switzerland

Marcus Hutter

RSISE @ ANU and SML @ NICTA, Canberra, ACT, 0200, Australia

December 2007



we can now define our formal measure of intelligence for arbitrary systems. Let E be the space of all computable reward summable environmental measures with respect to the reference machine $\mathcal{U}$, and let K be the Kolmogorov complexity function. The expected performance of agent π with respect to the universal distribution $2^{-K(\mu)}$ over the space of all environments E is given by,

$$\Upsilon(\pi) := \sum_{\mu \in E} 2^{-K(\mu)} V_{\mu}^{\pi}.$$

We call this the *universal intelligence* of agent π .

A.I. (well, not really)

- Science fiction
 - Isaac Asimov's I Robot (1940-50)
 - Robert Wise's "The Day the Earth Stood Still" (1951)
 - Groff Conklin's "Science Fiction - Thinking Machines" (1954)
 - Forbidden Planet (1956)






The two schools of A.I.






Artificial Intelligence (1956)

- Knowledge-based approach uses mathematical logic to simulate the human mind




- Neural-net approach simulates the structure of the brain



Artificial Intelligence

1956: Allen Newell and Herbert Simon demonstrate the "Logic Theorist", the first A.I. program, that uses "heuristics" (rules of thumb) and proves 38 of the 52 theorems in Whitehead's and Russell's "Principia Mathematica"

1957: "General Problem Solver" (1957): a generalization of the Logic Theorist but now a model of human cognition



RM-3731-CC
JUNE 1963

THE LOGIC THEORY MACHINE:
A MODEL HEURISTIC PROGRAM

Einar Stefferud




Artificial Intelligence

1956: Ray Solomonoff's inductive inference engine



Artificial Intelligence

1957: Frank Rosenblatt's Perceptron, the first artificial neural network





Artificial Intelligence



1957: Frank Rosenblatt's Perceptron, the first artificial neural network



Activation Functions:



Perceptron = Linear Regression



Perceptron = Logistic Regression

Artificial Intelligence

1957: Frank Rosenblatt's Perceptron
Running on a \$2,000,000 IBM 704
and only capable of learning the
difference between right and left

The New York Times

July 13, 1958

Electronic 'Brain' Teaches Itself

The Navy last week demonstrated the embryo of an electronic computer named the Perceptron which, when completed in about a year, is expected to be the first non-living mechanism able to "perceive, recognize and identify its surroundings without human training or control." Navy officers demonstrating a preliminary form of the device in Washington said they hesitated to call it a machine because it is so much like a "human being without life."

Dr. Frank Rosenblatt, research psychologist at the Cornell Aeronautical Laboratory, Inc., Buffalo, N. Y., designer of the Perceptron, conducted the demonstration. The machine, he said, would be the first electronic device to think as the human brain. Like humans, Perceptron will make mistakes at first, "but it will grow wiser as it gains experience," he said.

The first Perceptron, to cost about \$100,000, will have about 1,000 electronic "association cells" receiving electrical impulses from an eyelike scanning device with 400 photocells. The human brain has ten billion responsive cells, including 100,000,000 connections with the eye.

Difference Recognized

The concept of the Perceptron was demonstrated on the Weather Bureau's \$2,000,000 IBM 704 computer. In one experiment, the 704 computer was shown 100 squares situated at random either on the left or the right side of a field. In 100 trials, it was able to "say" correctly ninety-seven times whether a square was situated on the right or left. Dr. Rosenblatt said that after having seen only thirty to forty squares the device had learned to

recognize the difference between right and left, almost the way a child learns.

When fully developed, the Perceptron will be designed to remember images and information it has perceived itself, whereas ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons, Dr. Rosenblatt said, will be able to recognize people and call out their names. Printed pages, longhand letters and even speech commands are within its reach. Only one more step of development, a difficult step, he said, is needed for the device to hear speech in one language and instantly translate it to speech or writing in another language.

Self-Reproduction


In principle, Dr. Rosenblatt said, it would be possible to build Perceptrons that could reproduce themselves on an assembly line and which would be "conscious" of their existence.

Perceptron, it was pointed out, needs no "priming." It is not necessary to introduce it to surroundings and circumstances, record the data involved and then store them for future comparison as is the case with present "mechanical brains." It literally teaches itself to recognize objects the first time it encounters them. It uses a camera-eye lens to scan objects or survey situations, and an electrical impulse system, patterned point-by-point after the human brain does the interpreting.

The Navy said it would use the principle to build the first Perceptron "thinking machines" that will be able to read or write.


Artificial Intelligence

1957: Noam Chomsky's "Syntactic Structures"



- (1) S → NP + VP
- (2) VP → Verb + NP
- (3) NP → Det + N
- (4) Verb → Aux + V
- (5) Det → *the, a, ...*
- (6) N → *man, ball, ...*
- (7) Aux → *will, can, ...*
- (8) V → *hit, see, ...*

- (1) S → A + B
- (2) B → C + D
- (3) A → a + b
- (4) C → c + e + f
- (5) D → d + g + h



S stands for Sentence, NP for Noun Phrase, VP for Verb Phrase, Det for Determiner, Aux for Auxiliary (verb), N for Noun, and V for Verb stem

Artificial Intelligence



Zellig Harris

- 1952 Discourse Analysis
- 1959 The first parser (for the Transformations and Discourse Analysis Project or TDAP)

Artificial Intelligence

1957: Richard Bellman 's Dynamic Programming



$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | \pi, s_0 = s \right] = \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in \mathcal{S}} \mathcal{P}_{s,s'}^a (\mathcal{R}_{ss'}^a + \gamma V^\pi(s'))$$

Bellman equation.

Artificial Intelligence

1958: Yehoshua Bar-Hillel's "proof" that machine translation is impossible without common-sense knowledge

[From: *Advances in Computers*, vol. 1 (1960), p.91-163]

The Present Status of Automatic Translation of Languages'

YEHOShUA BAR-HILLEL


Hebrew University, Jerusalem, Israel



1. Aims and Methods, Survey and Critique
 - 1.1 Introduction
 - 1.2 Unreasonableness of Aiming at Fully Automatic, High Quality Translation
 - 1.3 Commercial Partly Mechanized, High Quality Translation Attainable in the Near Future
 - 1.4 Compromising in the Wrong Direction
 - 1.5 A Critique of the Overestimation of Statistics and the "Empirical Approach"

Artificial Intelligence


- 1959: John McCarthy's "Programs with Common Sense" focuses on knowledge representation
- 1959: Arthur Samuel's Checkers, the world's first self-learning program
- 1960: Hilary Putnam's Computational Functionalism
- 1962: Joseph Engelberger deploys the industrial robot Unimate at General Motors



Artificial Intelligence

The backpropagation algorithm

Arthur Bryson 1961
Henry Kelley 1960
Stuart Dreyfus 1962
Seppo Linnainmaa 1970
Paul Werbos 1974



Bryson 1961; Kelley 60; Dreyfus 62; Linnainmaa 70; Werbos 74

Artificial Intelligence

The backpropagation algorithm

1. Perform a feedforward pass, computing the activations for layers L_2 , L_3 , and so on up to the output layer L_{n_l} .

2. For each output unit i in layer n_l (the output layer), set

$$\delta_i^{(n_l)} = \frac{\partial}{\partial z_i^{(n_l)}} \frac{1}{2} \|y - h_{W,b}(x)\|^2 = -(y_i - a_i^{(n_l)}) \cdot f'(z_i^{(n_l)})$$

3. For $l = n_l - 1, n_l - 2, n_l - 3, \dots, 2$


For each node i in layer l , set

$$\delta_i^{(l)} = \left(\sum_{j=1}^{s_{l+1}} W_{ji}^{(l)} \delta_j^{(l+1)} \right) f'(z_i^{(l)})$$

4. Compute the desired partial derivatives, which are given as:

$$\begin{aligned} \frac{\partial}{\partial W_{ij}^{(l)}} J(W, b; x, y) &= a_j^{(l)} \delta_i^{(l+1)} \\ \frac{\partial}{\partial b_i^{(l)}} J(W, b; x, y) &= \delta_i^{(l+1)}. \end{aligned}$$

x	Input features for a training example, $x \in \mathbb{R}^n$.
y	Output/target values. Here, y can be vector valued. In the case of an autoencoder, $y = x$.
$(x^{(0)}, y^{(0)})$	The i -th training example
$h_{W,b}(x)$	Output of our hypothesis on input x , using parameters W, b . This should be a vector of the same dimension as the target value y .
$W_{ij}^{(l)}$	The parameter associated with the connection between unit j in layer l , and unit i in layer $l + 1$.
$b_i^{(l)}$	The bias term associated with unit i in layer $l + 1$. Can also be thought of as the parameter associated with the connection between the bias unit in layer l and unit i in layer $l + 1$.
θ	Our parameter vector. It is useful to think of this as the result of taking the parameters W, b and "unrolling" them into a long column vector.
$a_i^{(l)}$	Activation (output) of unit i in layer l of the network. In addition, since layer L_1 is the input layer, we also have $a_i^{(1)} = x_i$.
$f(\cdot)$	The activation function. Throughout these notes, we used $f(z) = \tanh(z)$.
$z_i^{(l)}$	Total weighted sum of inputs to unit i in layer l . Thus, $a_i^{(l)} = f(z_i^{(l)})$.



Equivalently in matrix-vector language:

1. Perform a feedforward pass, computing the activations for layers L_2 , L_3 , up to the output layer L_{n_l} , using Equations (6-7).

2. For the output layer (layer n_l), set

$$z^{(l+1)} = W^{(l)} a^{(l)} + b^{(l)} \quad (6)$$

$$a^{(l+1)} = f(z^{(l+1)}) \quad (7)$$

$$\delta^{(n_l)} = -(y - a^{(n_l)}) \bullet f'(z^{(n_l)})$$

3. For $l = n_l - 1, n_l - 2, n_l - 3, \dots, 2$

Set

$$\delta^{(l)} = ((W^{(l)})^T \delta^{(l+1)}) \bullet f'(z^{(l)})$$

4. Compute the desired partial derivatives:

$$\nabla_{W^{(l)}} J(W, b; x, y) = \delta^{(l+1)} (a^{(l)})^T,$$

$$\nabla_{b^{(l)}} J(W, b; x, y) = \delta^{(l+1)}.$$

Source: Andrew Ng

Artificial Intelligence

1963: The birth of computer vision (Lawrence Roberts)


Machine Perception Of Three-Dimensional Solids

Massachusetts Institued of Technology

Lincoln Laboratory

22 May 1963

Lawrence G. Roberts



Artificial Intelligence

1964: Peter Toma demonstrates the machine-translation system Systran

1964: IBM's "Shoebox" for speech recognition



Artificial Intelligence

1965: Jack Good speculates about
"ultraintelligent machines" (the
"singularity")



1965: Herbert Simon predicts that
"Machines will be capable, within 20
years, of doing any work a man can do"

1965: Hubert Dreyfus's "Alchemy and
Artificial Intelligence"

ALCHEMY AND ARTIFICIAL INTELLIGENCE


Hubert L. Dreyfus *

The RAND Corporation, Santa Monica, California




Artificial Intelligence

1965: Ed Feigenbaum's Dendral expert system: domain-specific knowledge




Artificial Intelligence

1965: Lofti Zadeh's Fuzzy Logic



1966: Ross Quillian's semantic networks




Extracting Semantic Networks from Text via Relational Clustering
(Stanley Kok and Pedro Domingos, 2008)

1966: ALPAC report on Machine Translation

Artificial Intelligence


1966: Joe Weizenbaum's chatbot Eliza





```
=====
EEEEEEEE L          IIIIII ZZZZZZZ AAA
E         L          I      Z      A  A
E         L          I      Z      A  A
EEEEEE   L          I      Z      A  A
E         L          I      Z      A  A
E         L          I      Z      A  A
EEEEEEEE LLLLLLLL IIIIII ZZZZZZ  A  A
=====
ELIZA > (PLEASE TYPE IN ALL CAPS) WHAT'S YOUR NAME DEAR ?
AMIT
ELIZA > HI, I'M ELIZA. WHAT DO YOU WANT TO TALK ABOUT ?
AMIT > CAN I TALK ABOUT MY PROBLEM ?
ELIZA > SURE... !
AMIT > DUE TO ON GOING EXAMS IN COLLEGE, I HAVE A LOT OF STRESS.
ELIZA > PLEASE RELAX A LITTLE. SLEEP WELL.
AMIT > THANKS FOR YOUR ADVICE.
ELIZA > NO MENTION.
AMIT > BYE
ELIZA > BYE AND KEEP IN TOUCH...
=====
```

Artificial Intelligence

1965: Alexey Ivakhnenko publishes the first learning algorithms for multi-layered networks





1966: Leonard Baum popularizes the Hidden Markov Model ("Statistical Inference for Probabilistic Functions of Finite State Markov Chains")



Artificial Intelligence

1969: Marvin Minsky & Samuel Papert's
"Perceptrons" kills neural networks

1971: Noam Chomsky's article against
Skinner's behaviorism



Artificial Intelligence


Why nobody argued

- Pitts died in May 1969
- McCullouch died in September 1969
- Rosenblatt died in 1971




The two schools of AI

1956: Allen Newell and Herbert
Simon's "Logic Theorist"
1959: John McCarthy's "Programs
with Common Sense"
1965: Ed Feigenbaum's Dendral
1965: Lofti Zadeh's Fuzzy Logic
1966: Ross Quillian's Semantic
Networks
1969: SRI's Shakey the Robot
1969: Roger Schank's Conceptual
Dependency Theory
1972: Bruce Buchanan's MYCIN
1972: Terry Winograd's SHRDLU
1974: Marvin Minsky's Frame



The first IJCAI: 1969

Only 2 papers on neural networks!



Proceedings of the First International Joint
Conference on Artificial Intelligence

IJCAI-69 Contents

INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE

May 7-9, 1969 Washington, D. C

Artificial Intelligence



1969: Cordell Green's automatic synthesis of programs

1969: Stanford Research Institute's Shakey the Robot

1969: Roger Schank's Conceptual Dependency Theory

A Conceptual Dependency Representation for a Computer-oriented Semantics



[Roger C. Schank](#)

Computer Science Department, School of Humanities and Sciences,
Stanford University, 1969

- 402 pages



Neuroscience


- 1949: Donald Hebb's cell assemblies (selective strengthening or inhibition of synapses causes the brain to organize itself into regions of self-reinforcing neurons - the strength of a connection depends on how often it is used)
- 1951: Wilder Penfield and Herbert Jasper prove that electrical stimulation of the temporal lobes can yield vivid recall of lost memories
- 1952: Paul Maclean discovers the "limbic system"
- 1952: Alan Hodgkin and Andrew Huxley discover how action potentials in neurons are propagated
- 1953: John Eccles describes excitatory and inhibitory potentials
- 1953: Roger Sperry studies the "split brain" in animals
- 1953: Eugene Aserinsky discovers "rapid eye movement" (REM) sleep that corresponds with periods of dreaming

Neuroscience

- 1957: Vernon Mountcastle discovers the modular organization of the brain (vertical columns)
- 1962: Roger Sperry and Michael Gazzaniga discover that the two hemispheres of the human brain are specialized in different functions
- 1962: David Hubel and Torsten Wiesel study the visual cortex of the cat
- 1964: John Young proposes a "selectionist" theory of the brain
- 1964: Paul Maclean's triune brain: three layers, each layer corresponding to a different stage of evolution
- 1964: Benjamin Libet discovers that the readiness potential precedes conscious awareness by about half a second
- 1968: Niels Jerne's selectionist model of the brain

Cognitive Science

- 1952: Jerome Bruner's Cognition Project (Harvard)
- 1956: Herb Simon and Al Newell's Logic Theorist (CMU)
- 1957: Noam Chomsky's "Syntactic Structures" (MIT)
- 1960: George Miller's and Jerome Bruner's Center for Cognitive Studies (Harvard)
- 1960: "Plans and the Structure of Behavior" by George Miller, Eugene Galanter, and Karl Pribram (Harvard)
- 1963: Ed Feigenbaum and Julian Feldman's Computers and Thought" (UC Berkeley)
- 1965: Donald Michie's Experimental Programming Unit (Edinburgh)
- 1965: George Mandler's Center for Human Information Processing (UCSD)





Cognitive Science



1967: Ulric Neisser's "Cognitive Psychology" (University of Pennsylvania)

1972: Don Norman's and Peter Lindsay's "Human Information Processing" (UCSD)

1972: Newell & Simon's "Human Problem Solving" (CMU)



Artificial Intelligence

1971: Ingo Rechenberg publishes his thesis
"Evolution Strategies", a set of optimization
methods for evolutionary computation



1972: Harry Klopff's "Brain Function and
Adaptive Systems"



APCRL-72-0164
3 MARCH 1972
SPECIAL REPORTS, NO. 133



AIR FORCE CAMBRIDGE RESEARCH LABORATORIES
L. G. HANSCOM FIELD, BEDFORD, MASSACHUSETTS

Brain Function and Adaptive Systems -

A. HARRY KLOPF


NEURON $\equiv$ MAN .

MAN $\equiv$ HEDONIST .

NEURON $\equiv$ HEDONIST .

Artificial Intelligence

1972: Bruce Buchanan's MYCIN



- a knowledge base
- a patient database
- a consultation/explanation program
- a knowledge acquisition program

Knowledge is organised as a series of IF THEN rules


Artificial Intelligence

1972: Terry Winograd's Shrdlu



Person: PICK UP A BIG RED BLOCK.
Computer: OK. (does it)
Person: GRASP THE PYRAMID.
Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.
Person: FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE HOLDING AND PUT IT INTO THE BOX.
Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.
Computer: OK. (does it)
Person: WHAT DOES THE BOX CONTAIN?
Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.
Person: WHAT IS THE PYRAMID SUPPORTED BY?
Computer: THE BOX.
Person: HOW MANY BLOCKS ARE NOT IN THE BOX?
Computer: FOUR OF THEM.
Person: IS AT LEAST ONE OF THEM NARROWER THAN THE ONE WHICH I TOLD YOU TO PICK UP?
Computer: YES, THE RED CUBE.
Person: IS IT SUPPORTED?
Computer: YES, BY THE TABLE.
Person: CAN THE TABLE PICK UP BLOCKS?
Computer: NO.
Person: CAN A PYRAMID BE SUPPORTED BY A BLOCK?
Computer: YES.
Person: CAN A PYRAMID SUPPORT A PYRAMID?
Computer: I DON'T KNOW.
Person: STACK UP TWO PYRAMIDS.
Computer: I CAN'T.
Person: WILL YOU PLEASE STACK UP BOTH OF THE RED BLOCKS AND EITHER A GREEN CUBE OR A PYRAMID?
Computer: OK. (does it)
Person: WHICH CUBE IS SITTING ON THE TABLE?

Artificial Intelligence



1972: Kenneth Colby's chatbot Parry

The first chatbot to chatbot conversation
(Stanford ran Parry and MIT ran Eliza
over the Arpanet)

ELIZA: Tell me your problems. Please terminate input with a period or a question mark.

PARRY: People get on my nerves sometimes.

ELIZA: I am not sure I understand you fully.

PARRY: You should pay more attention.

ELIZA: Suppose you should pay more attention.

PARRY: You're entitled to your own opinion.

ELIZA: What makes you think I am entitled to my own opinion?

PARRY: It bothers me just to be around people in general.

ELIZA: Please go on.

PARRY: I went to the track at Bay Meadows a while back.

ELIZA: What does that suggest to you?

PARRY: I went to the races.

Artificial Intelligence

1972: Hubert Dreyfus's "What Computers Can't Do"




1973: Jim Baker applies the Hidden Markov Model to speech recognition

1973: The Lighthill Report kills AI in Britain

1974: Marvin Minsky's Frame (see chapter on "Cognition")

1975: Roger Schank's Script (see chapter on "Cognition")

1975: John Holland's Genetic Algorithms



Restaurant Script (Schank & Abelson, 1975)

```
Scene 1 Entering
PTRANS self into restaurant
ATTEND eyes to empty tables
MBUILD place to sit
PTRANS self to table
MOVE sit down

Scene 2 Ordering
ATRANS Get menu
MTRANS Read menu
MBUILD Decide what self wants
MTRANS Order to waitress/waiter

Scene 3 Eating
ATRANS Get food
INGEST Eat food

Scene 4 Exiting
MTRANS Ask for check
ATRANS Get check
ATRANS Tip waitress/waiter
PTRANS Self to cashier
ATRANS Money to cashier
PTRANS Self out of restaurant
```

Go Back

Artificial Intelligence

Fuzzy Inference Systems


- Zadeh's "Outline of a new approach to the analysis of complex systems and decision processes" (1973)
- Ebrahim "Abe" Mamdani 1975
- Yahachiro Tsukamoto 1979
- Michio Sugeno 1985

Artificial Intelligence

1978: John McDermott's expert system R1/XCON

1978: Ryszard Michalski builds the first practical system that learns from examples, AQ11

1978: David Marr's theory of vision



http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/GOMES1/marr.html

Artificial Intelligence

1979: Johan DeKleer's qualitative reasoning

1979: Hans Berliner's BKG 9.8 at Carnegie-Mellon University (connected by satellite to the robot Gammonoid) beats the world champion of backgammon in Monte Carlo

Backgammon Computer Program Beats World Champion

Artificial Intelligence, vol. 14 (1980), pp. 205–220.



Speech Recognition

1968: Taras Vintsiuk's dynamic time warping to recognize words spoken at different speeds


1969 Raj Reddy's speech-recognition group at CMU:

- Harpy (Bruce Lowerre 1976),
- Hearsay-II (Rick Hayes-Roth, Lee Erman, Victor Lesser and Richard Fennell, 1975);
- Dragon (Jim Baker, 1975)



Speech Recognition

1976: Fred Jelinek's statistical method predicts statistically the next word




$$\hat{W} = \underset{W}{\operatorname{argmax}} p(W|A) = \underset{W}{\operatorname{argmax}} p(A|W)p(W) = \underset{W}{\operatorname{argmax}} \{\ln p(A|W) + \ln p(W)\}$$

$$\hat{W} = \underset{W}{\operatorname{argmax}} \{ \lambda \ln p(A|W) + \ln p(W) \}$$

Acoustic Model
Language Model

Acoustic Model (Hidden Markov Models)



Language Model

$$p(W) = p(w_1 w_2 w_3 w_4 w_5) = p(w_1) p(w_2|w_1) p(w_3|w_1 w_2) p(w_4|w_1 w_2 w_3) p(w_5|w_1 w_2 w_3 w_4)$$

$$p(W) \cong p(w_1) p(w_2|w_1) p(w_3|w_1 w_2) p(w_4|w_2 w_3) p(w_5|w_3 w_4)$$

Trigram

$$p(w_i|w_{i-1}) = \frac{c(w_{i-1} w_i)}{\sum_{w_i} c(w_{i-1} w_i)}$$

ML estimate

$$r^* = (r + 1) \frac{n_{r+1}}{n_r}$$

Good Turing Smoothing


Speech Recognition

1980: Jack Ferguson's "Blue Book" popularizes statistical methods based on the Hidden Markov Model

IBM (Fred Jelinek) vs Bell Labs (Lawrence Rabiner)
speaker-dependent vs speaker-independent

Neural Networks

- 1969: Bryson and Yu-Chi Ho's gradient method
- 1970: Seppo Linnainmaa
- 1970: Pete Uttley's Informon
- 1972: James Anderson
- 1972: Teuvo Kohonen
- 1973: Christoph von der Malsburg
- 1974: Paul Werbos' backpropagation
- 1975: Kunihiro Fukushima's Cognitron
- 1975: Stephen Grossberg's adaptive resonance theory
- 1978: Shunichi Amari



Artificial Intelligence




1980: John Searle's "Chinese Room"

1980: Intellicorp, the first major start-up for Artificial Intelligence

1982: Japan's Fifth Generation Computer Systems project

Computer generations:

- 0th generation: 500 B.C., mechanical gears.
- 1st generation: 1940's, vacuum tubes.
- 2nd generation: 1950's, transistors.
- 3rd generation: 1960's, integrated circuits (ICs).
- 4th generation: Microprocessors.
- FGCS, Fifth Generation Computer Systems project.
 - Japan's Ministry of International Trade and Industry.



Artificial Intelligence

1980s: Second A.I. bubble

1982: Japan's Fifth Generation Computer Systems project

1984: Doug Lenat's Cyc to catalog common sense

1988: Hans Moravec in "Mind Children"
(1988): "robots will eventually succeed us: humans clearly face extinction".



Computers *versus* Common Sense

- **CYC:** Software that partially understands
- The impact of that on the AIM dream




Dr. Douglas B. Lenat

7718 Wood Hollow Drive, Austin, TX 78731

Email: Lenat@cyc.com

Phone: 512-342-4001

Nov 6, 2000



A.I. Winters

1957: Herbert Simon declares that *"there are now in the world machines that think, that learn, and that create"*

1958: A New York Times article (8 July 1958) reporting a press conference by Rosenblatt that *"the Perceptron is the embryo of an electronic computer that (the US Navy) expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence"*

1958: Bar-Hillel publishes a "proof" that machine translation is impossible

1965: Herbert Simon predicts that *"machines will be capable within 20 years of doing any work a man can do"*

1966: The ALPAC Report causes reduction in funding for machine translation research



A.I. Winters


1970: Marvin Minsky to Life Magazine:

*“In from three to eight years we will have a machine with **the general intelligence of an average human being**”*

1972: Richard Karp shows there are many problems that can probably only be solved in exponential time

1973: The Lighthill Report kills A.I. in the UK

1980s: Fifth Generation illusion



A.I. Winters

Limited knowledge of the world:

- Restricted to micro-worlds (e.g. Blocks World)

- Restricted to pattern-matching (e.g. Eliza)

Inherent limitations of computability:

- Intractability, combinatorial explosion

- Undecidability: the halting problem

What saved A.I.

Neuroscience: Fukushima's convolutional nets



Physics: Hopfield's recurrent neural networks

Canada: CIFAR (Canadian Institute for Advanced Research)



Vannoccio Biringuccio's
"De la Pirotechnia"
(1540, first printed book
on metallurgy published
in Europe)

Next...

- See <http://www.scaruffi.com/singular> for the index of this Powerpoint presentation and links to the other parts

