
Concepts and History

Artificial Intelligence

Slides are mostly adapted from AIMA,
Svetlana Lazebnik (UIUC), Percy Liang (Stanford)

Definition of AI

“Intelligence: The ability to learn and solve problems”

Webster's Dictionary.

“Artificial intelligence (AI) is the intelligence exhibited by machines or software”

Wikipedia.

“The science and engineering of making intelligent machines”

McCarthy.


“The study and design of intelligent agents, where an intelligent agent is a system that perceives its environment and takes actions that maximize its chances of success.”

Russel and Norvig AI book.


What is AI?

- Four possible definitions (textbook ch. 1):


Thinking
humanly




Acting
humanly



Thinking
rationally



Acting
rationally



What is AI?


<p>Thinking humanly</p> <p>The exciting new effort to make computers think...machines with minds, 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 et. al , 1998)</p> <p>AI ... is concerned with intelligent behaviour in artifacts (Nilson, 1998)</p>

* A system is rational if it does the right thing given what it knows

- Human-centered approaches use empirical science, involving hypothesis and experimental confirmation,
- Rationalist approaches involves a combination of mathematics and engineering


AI definition 1: Thinking humanly

- Need to study the brain as an information processing machine: cognitive science and neuroscience



AI definition 1: Thinking humanly

- Can we build a brain?



10^{11} neurons
 10^{14} synapses
 cycle time: 10^{-3} sec

VS.

10^9 transistors
 10^{12} bits of RAM
 cycle time: 10^{-9} sec



AI definition 1: Thinking humanly

- Can we build a brain?

Computers	Brains
Digital	Analog
Fixed architecture	Evolving architecture
Fixed processing speed	No system clock
Modular, (primarily) serial	Massively parallel
Separate hardware, software	No distinction between hardware and software
Separate computation, memory	No distinction between computation and memory
Disembodied	Embodied

<http://scienceblogs.com/developingintelligence/2007/03/27/why-the-brain-is-not-like-a-co/>

AI definition 1: Thinking humanly



the WHITE HOUSE
**BRAIN
INITIATIVE**

BRAIN RESEARCH THROUGH ADVANCING
INNOVATIVE NEUROTECHNOLOGIES

Since President Obama announced the **BRAIN Initiative** in April 2013, dozens of leading technology firms, academic institutions, scientists and other key contributors to the field of neuroscience have answered his call and made significant commitments to advancing the Initiative.



Human Brain Project

FUTURE COMPUTING

Develop novel neuromorphic and neurobotic technologies based on the brain's circuitry and computing principles.





Thinking humanly: cognitive modeling

- In order to say that a given program thinks like a human, we must have some way of determining how humans think
 - Requires scientific theories of internal activities of the brain
 - How to validate? requires
 - 1) Predicting and testing behavior of human subjects (top-down)
 - 2) Direct identification from neurological data (bottom-up)
 - Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI
-

AI definition 2: Acting humanly

- The Turing Test



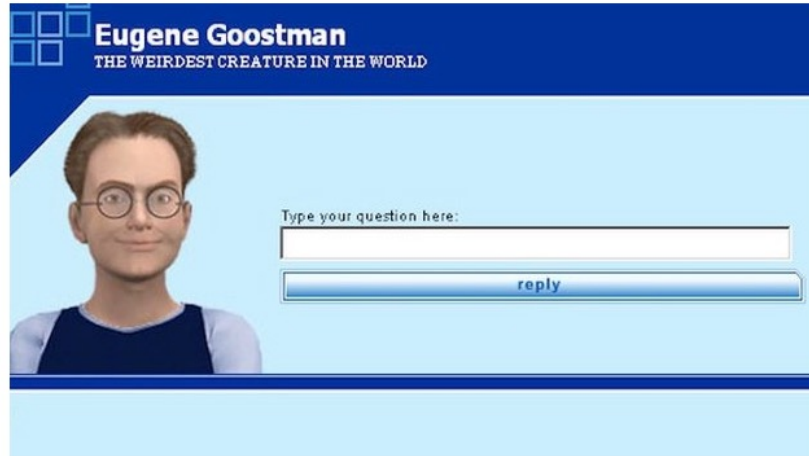
- What capabilities would a computer need to have to pass the Turing Test?
 - Natural language processing
 - Knowledge representation
 - Automated reasoning
 - Machine learning
- Turing predicted that by the year 2000, machines would be able to fool 30% of human judges for five minutes

A. Turing, [Computing machinery and intelligence](#),

B. Mind 59, pp. 433-460, 1950

AI is solved?

Computer AI passes Turing test in 'world first'



Eugene Goostman simulates a 13-year-old Ukrainian boy

9 June 2014 Last updated at 08:36 ET

A computer program called Eugene Goostman, which simulates a 13-year-old Ukrainian boy, is said to have passed the Turing test at an event organised by the University of Reading.

The test investigates whether people can detect if they are talking to machines or humans.

The experiment is based on Alan Turing's question-and-answer game **Can Machines Think?**

No computer has passed the test before under these conditions, it is reported.

However, some artificial intelligence experts have disputed the victory, suggesting the contest had been weighted in the chatbot's favour.

The 65-year-old **Turing Test** is successfully passed if a computer is mistaken for a human more than 30% of the time during a series of five-minute keyboard conversations.

On 7 June Eugene convinced 33% of the judges at the Royal Society in London that it was human.

Other artificial intelligence (AI) systems also competed, including **Cleverbot**, **Elbot** and **Ultra Hal**.

Judges included actor Robert Llewellyn, who played an intelligent robot in BBC Two's science-fiction sitcom *Red Dwarf*, and Lord Sharkey, who led the successful campaign for Alan Turing's posthumous pardon, over a conviction for homosexual activity in 2013.

Eugene was created by Vladimir Veselov, who was born in Russia and now lives in the United States, and Ukrainian-born Eugene Demchenko, who now lives in Russia.

Related Stories

[How the Turing Test inspired AI](#)

[Playing solitaire with Turing](#)

[Is artificial intelligence possible?](#)

[16:22:39] **Local:**
What comes first to mind when you hear the word "toddler"?

[16:22:49] **Remote:**
And second?

[16:23:26] **Local:**
What comes to mind when you hear the word "grown-up"?

[16:23:37] **Remote:**
Please repeat the word to me 5 times.

[16:24:22] **Local:**
What do you think is the purpose of emotion?

Man or machine? A glimpse at one of the conversations.

<http://www.bbc.com/news/technology-27762086>
http://en.wikipedia.org/wiki/Eugene_Goostman
<http://www.scottaaronson.com/blog/posts/20140608/>

What's wrong with the Turing test?

- Variability in protocols, judges
 - Success depends on deception!
 - Chatbots can do well using “cheap tricks”
 - First example: [ELIZA](#) (1966)
 - [Chinese room argument](#): one may simulate intelligence without having true intelligence (more of a philosophical objection)
-

A better Turing test?

- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”:
- *The trophy would not fit in the brown suitcase because it was so small.
What was so small?*
 - *The trophy*
 - *The brown suitcase*

A better Turing test?

- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”:
- *The trophy would not fit in the brown suitcase because it was so **large**.
What was so **large**?*
 - *The trophy*
 - *The brown suitcase*

A better Turing test?

- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”:

• The large ball crashed right through the table because it was made of styrofoam.

What was made of styrofoam?

- *The large ball*
- *The table*

A better Turing test?

- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”:

• *The large ball crashed right through the table because it was made of **steel**.*

*What was made of **steel**?*

- *The large ball*
- *The table*

A better Turing test?

- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”:
- *The sack of potatoes had been placed below the bag of flour, so it had to be moved first. What had to be moved first?*
 - *The sack of potatoes*
 - *The bag of flour*

A better Turing test?


- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”:
- *The sack of potatoes had been placed **above** the bag of flour, so it had to be moved first. What had to be moved first?*
 - *The sack of potatoes*
 - *The bag of flour*

A better Turing test?

- Advantages over standard Turing test
 - Test can be administered and graded by machine
 - Does not depend on human subjectivity
 - Does not require ability to generate English sentences
 - Questions cannot be evaded using verbal dodges
 - Questions can be made “Google-proof” (at least for now...)

Acting humanly: Turing Test

- Proposed by Alan Turing (1950)
- Operational definition for intelligent behaviour: the Imitation Game



- The computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or not
- Suggested major components of AI: natural language processing, knowledge representation, automated reasoning, machine learning
- Total Turing test also requires computer vision and robotics



- “Artificial flight” is succeeded by not imitating the birds, but by learning aerodynamics. The goal is not to fool pigeons.
-

AI definition 3: Thinking rationally

- Idealized or “right” way of thinking
 - Aristotle: what are correct arguments/thought processes?
 - Several Greek schools developed various forms of *logic*: *notation* and *rules of derivation* for thoughts;
 - **Logic**: patterns of argument that always yield correct conclusions when supplied with correct premises
 - “Socrates is a man; all men are mortal; therefore Socrates is mortal.”
 - **Logician approach to AI**: describe problem in formal logical notation and apply general deduction procedures to solve it
-

Thinking rationally: "laws of thought"

- Direct line through mathematics and philosophy to modern AI

Problems:

1. Not all intelligent behavior is mediated by logical notations
 2. What is the purpose of thinking? What thoughts should I have?
 3. Computational blow up
- Additional Problems with the logicist approach
 - Computational complexity of finding the solution
 - Describing real-world problems and knowledge in logical notation
 - Dealing with uncertainty
 - A lot of “rational” behavior has nothing to do with logic
-

AI definition 4: 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
 - *A rational agent* acts to optimally achieve its goals
 - Goals are application-dependent and are expressed in terms of the **utility of outcomes**
 - Being rational means **maximizing your (expected) utility**
 - This definition of rationality only concerns the decisions/actions that are made, not the cognitive process behind them
 - In practice, utility optimization is subject to the agent's computational constraints (*bounded rationality* or *bounded optimality*)
-

Utility maximization formulation

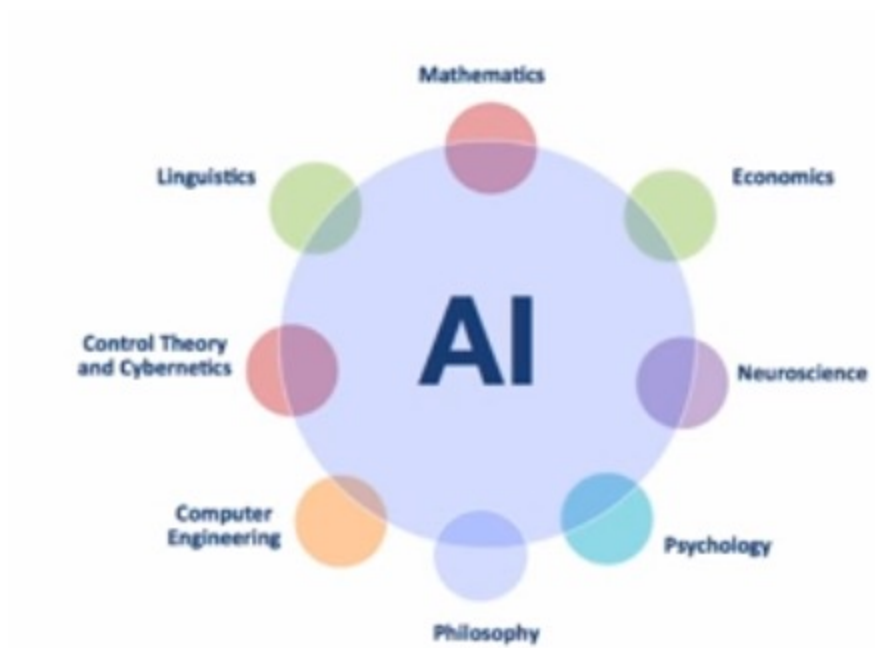
- Advantages
 - Definition is about the agent's decisions/actions, not the cognitive process behind them
 - Generality: goes beyond explicit reasoning, and even human cognition altogether
 - Practicality: can be adapted to many real-world problems
 - Naturally accommodates uncertainty
 - Amenable to good scientific and engineering methodology
 - Avoids philosophy and psychology
 - Disadvantages?
 - It may be hard to formulate utility functions, especially for complex open-ended tasks
 - The AI may end up “gaming” the utility function, or its operation may have unintended consequences
 - Has limited applicability to humans
-

Rational agents

- An agent is an entity that perceives and acts
- This course is about designing rational agents
- 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
-



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
 - Economics utility, decision theory, game theory, Markov decision processes
 - Neuroscience study of brain functioning, how brain and machines are (dis)similar
 - Psychology how do we think and act?
phenomena of perception and motor control, experimental techniques
 - Computer building powerful machines to make AI possible
 - Control theory design systems that maximize an objective function over time
 - Linguistics knowledge representation, grammar
-

AI founders

- Aristotle
- Alan Turing
- John Mc Carthy
- Warren McCulloch
- Walter Pitts
- Claude Shannon
- Marvin Minsky
- Dean Edmonds
- Herbert Simon
- Allen Newell
- David Waltz
- Tom Mitchell
- Stuart J. Russell
- Peter Norvig
- etc.

AI Resources

- Major journals/conferences: JAIR, TPAMI, JMLR, IJCAI, AAAI, IAAI, CVPR, ECAI, ICML, NIPS, etc.

- Video lectures:

http://videlectures.net/Top/Computer_Science/Artificial_Intelligence/

AI: History and themes

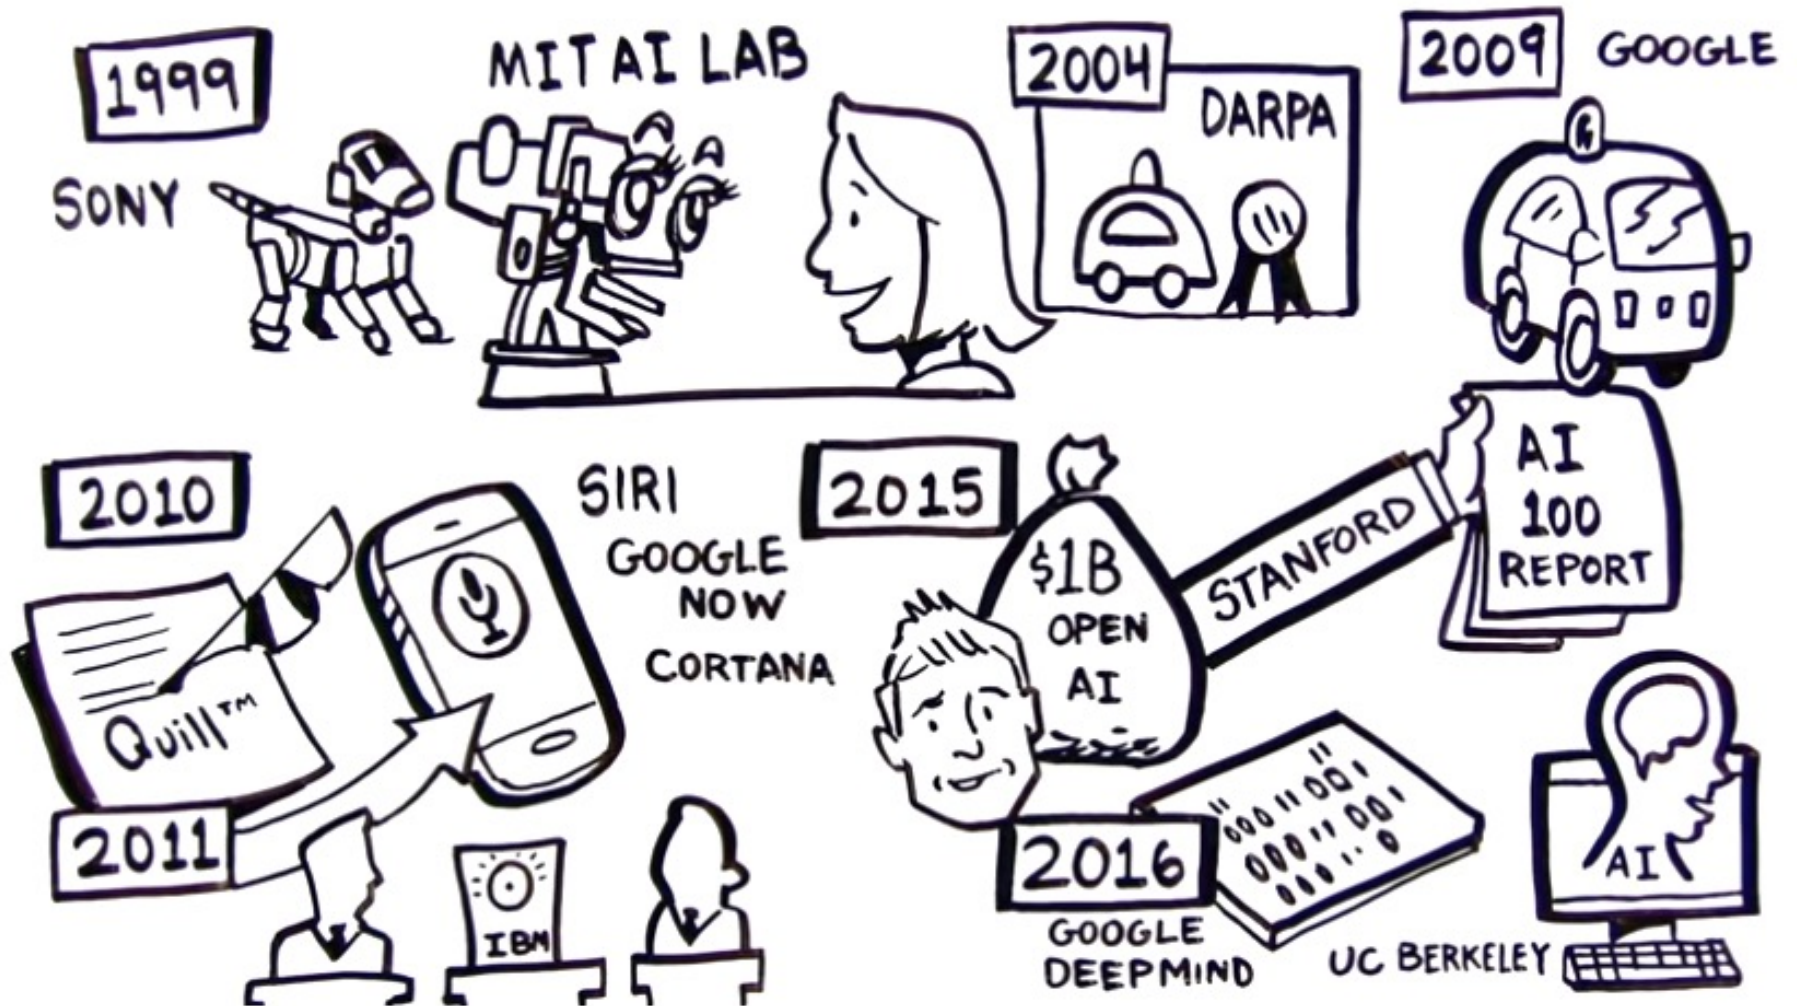


[Image source](#)

AI history



AI History



History of AI

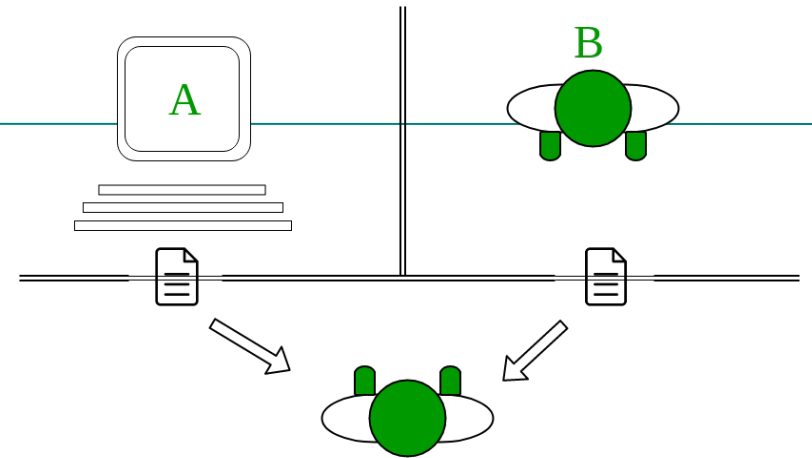

- **1940-1950:** Gestation of AI
 - McCulloch & Pitts: Boolean circuit to model of brain
 - Turing's Computing Machinery and Intelligence
<http://www.turingarchive.org/browse.php/B/9>
- **1950-1970:** Early enthusiasm, great expectations
 - Early AI programs, Samuel's checkers program
 - Birth of AI @ Dartmouth meeting 1956.
 - Check out the MIT video "The thinking Machine" on youtube
<https://www.youtube.com/watch?v=aygSMgK3BEM>
- **1970-1990:** Knowledge-based AI
 - Expert systems, AI becomes an industry
 - AI winter
- **1990-present:** Scientific approaches
 - Neural Networks: le retour
 - The emergence of intelligent agents
 - AI becomes "scientific", use of probability to model uncertainty
 - AI Spring!
 - The availability of very large datasets.
 - * Data will drive future discoveries and alleviate the complexity in AI.

Origins of AI: Early excitement

- **1940s** First model of a neuron (W. S. McCulloch & W. Pitts)
 - Hebbian learning rule
 - Cybernetics
 - **1950s** Turing Test
 - Perceptrons (F. Rosenblatt)
 - Computer chess and checkers (C. Shannon, A. Samuel)
 - Machine translation (Georgetown-IBM experiment)
 - Theorem provers (A. Newell and H. Simon, H. Gelernter and N. Rochester)
 - **1956 Dartmouth meeting: “Artificial Intelligence” adopted**
-

The gestation of AI (1943-1955)

- The first work that is now generally recognized as AI was done by Warren McCulloch and Walter Pitts (1943)
 - They proposed a model of artificial neurons in which each neuron is characterized as being “on” or ”off”, with a switch to “on” occurring in response to stimulation by a sufficient number of neighboring neurons
 - They showed that any computable function could be computed by some network of connected neurons, and all the logical connectives (and, or, not, etc) could be implemented by simple net structures
 - Later (1949), Donald Hebb demonstrated a simple updating rule for modifying the connection strengths between neurons (Hebbian rule)
 - In 1950 Marvin Minsky and Dean Edmond (two undergrads at Harvard) built the first neural network computer (SNARC)
 - Then Minsky studied universal computation in neural networks during his PhD at Princeton
 - Later, Minsky proved influential theorems showing the limitations of NN research
-



“I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine 'and 'think'. ... [But] Instead of attempting such a definition I shall replace the question by another... The new form of the problem can be described in terms of a game which we call the 'imitation game'.”

-Alan Turing, “Computing Machinery and Intelligence”, 1950



Alan Turing (1950)

- It was Alan Turing who first articulated a complete vision of AI in his 1950 article
 - “Computing Machinery and Intelligence”
 - He introduced the Turing test, machine learning, genetic algorithms, and reinforcement learning
-

The birth of AI (1956)

- In 1956 John McCarthy organized a two month workshop at Dartmouth College bringing together researchers interested in automata theory, neural nets, and the study of intelligence
 - This workshop did not lead to any new breakthroughs but introduced the major figures to each other. For the next 20 years the field is dominated by these people and their students and colleagues at MIT, CMU, Stanford and IBM
 - The name Artificial Intelligence is also given by McCarthy during this workshop (computational rationality would be the other alternative)
 - The proposal of the workshop also explains why AI becomes a separate field rather than a subfield of operational research, control theory or mathematics:
 - AI duplicates human faculties like creativity, self-improvement, and language use
 - Also the methodology is different. AI is clearly a branch of computer science and AI is the only one field to attempt to build machines that will function autonomously in complex, changing environments
-

-
- 1956: Workshop at Dartmouth College;
attendees:
John McCarthy,
Marvin Minsky,
Claude Shannon, etc.
 - Aim for general principles:
 - Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it.




Overwhelming optimism...

- Machines will be capable, within twenty years, of doing any work a man can do. |Herbert Simon
 - Within 10 years the problems of artificial intelligence will be substantially solved. |Marvin Minsky
 - I visualize a time when we will be to robots what dogs are to humans, and I'm rooting for the machines. |Claude Shannon
-

Herbert Simon, 1957

• *“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.”*



- Prediction came true – but 40 years later instead of 10
-

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo
of Computer Designed to
Read and Grow Wiser

WASHINGTON, July 7 (UPI)—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.

The embryo—the Weather Bureau's \$2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of \$100,000.

Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human be-

ings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

Without Human Controls

The Navy said the perceptron would be the first non-living mechanism "capable of receiving, recognizing and identifying its surroundings without any human training or control."

The "brain" is designed to remember images and information it has perceived itself. Ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

1958 New York Times...

In today's demonstration, the "704" was fed two cards, one with squares marked on the left side and the other with squares on the right side.


Learns by Doing

In the first fifty trials, the machine made no distinction between them. It then started registering a "Q" for the left squares and "O" for the right squares.


Dr. Rosenblatt said he could explain why the machine learned only in highly technical terms. But he said the computer had undergone a "self-induced change in the wiring diagram."

The first Perceptron will have about 1,000 electronic "association cells" receiving electrical impulses from an eye-like scanning device with 400 photo-cells. The human brain has 10,000,000,000 responsive cells, including 100,000,000 connections with the eyes.

Birth of AI, early successes



- Checkers (1952): Samuel's program learned weights and played at strong amateur level
- Problem solving (1955): Newell & Simon's Logic Theorist: prove theorems in Principia Mathematica using search + heuristics; later, General Problem Solver (GPS)



Early enthusiasm, great expectations (1952-1969)

- Allen Newell and Herbert Simon created a reasoning program – Logic Theorist (LT)
 - The program was able to prove most of the theorems in the book Principia Mathematica by Russell and Whitehead and one of the proofs was even shorter than the one in Principia
 - This success is followed by General Problem Solver (GPS) which was designed from the start to imitate human problem-solving protocols.
 - Within the limited class of puzzles it could handle, it turned out that the order in which the program considered subgoals and possible actions was similar to that in which humans approached the same problems
 - Based on that Newell and Simon formulated the famous physical symbol system hypothesis – “ a physical symbol system has the necessary and sufficient means for general intelligent action” – any system (human or machine) exhibiting intelligence must operate by manipulating data structures composed of symbols
 - In 1959 at IBM, Geometry Theorem Prover is constructed – prove theorems that many students of mathematics would find tricky
 - Arthur Samuel wrote a series of programs for checker (draughts) that eventually learned to play at a strong amateur level. His program quickly learned to play a better game than his creator
-

Early enthusiasm, great expectations (1952-1969)

- The early years of AI were full of successes – in a limited way
 - Given the primitive computers and programming tools of the time, and the fact that only a few years earlier computers were seen as things that could do arithmetic and no more, it was astonishing whether a computer did anything remotely clever
 - The intellectual establishment, by and large, preferred to believe that “a machine can never do X”
 - AI researchers naturally responded by demonstrating one X after another
 - John McCarthy referred to this period as the “Look, Ma, no hands!” era
-








Early enthusiasm, great expectations (1952-1969)

- John McCarthy made three crucial contributions in one year, 1958, at MIT
 - Defined LISP which became the dominant AI programming language
 - Time sharing is invented
 - In “Programs with Common Sense” he described Advice Taker, a hypothetical program that can be seen as the first complete AI system
 - For example, he showed how some simple axioms would enable the program to generate a plan to drive to the airport to catch a plane
 - It embodied central principles of knowledge representation and reasoning. That is useful to have a formal, explicit representation of the world and of the way an agent’s actions affect the world and to be able to manipulate these representations with deductive processes
 - most of that work remains relevant today
-

Early enthusiasm, great expectations (1952-1969)

- While McCarthy stressed representation and reasoning in formal logic, Minsky was more interested in getting programs to work and eventually developed an anti-logical outlook
 - His group chose limited problems that appeared to require intelligence to solve.
 - These limited domains are known as microworlds (e.g. closed form calculus integration problems, geometric analogy problems that appear in IQ tests)
 - The most famous microworld was the blocks world, which consists of a set of solid blocks placed on a table. A typical task is to rearrange the blocks in a certain way, using a robot hand that can pick one block at a time
-

Analogy Problem

A**B****C****1****2****3****4****5**

Symbolic Integration

A Heuristic Program that Solves Symbolic Integration Problems in Freshman Calculus

James R. Slagle, 1963

As explained in Prof. [Patrick Winston's AI lecture](#), [Matlab](#) uses

12 safe transformations, like that constant out, sums, etc.

12 heuristic transformation e.g. $\sin x$, $\cos x$...

and, a table that contains 26 anchors to the calculations.

So that every integration can be calculated using these rules.

Example from Patrick Winston's AI lecture

- $\int \frac{-5 x^4}{(1-x^2)^{5/2}} dx$
- Safe transformations
- $\int -f(x)dx = - \int f(x) dx$
- $\int c f(x) = c \int f(x)dx$
- $\int \sum f_i(x) dx = \sum \int f_i(x)dx$
- $\int \frac{p(x)}{g(x)}$ DIVIDE

<https://www.youtube.com/watch?v=PNKj529yY5c&t=2448s>

https://en.wikipedia.org/wiki/Integration_by_parts

Example from Patrick Winston's AI lecture

$$\int \frac{-5x^4}{(1-x^2)^{5/2}} dx \rightarrow \int \frac{5x^4}{(1-x^2)^{5/2}} dx \rightarrow \int \frac{x^4}{(1-x^2)^{5/2}} dx$$

HEURISTIC TRANSFORMATIONS

A $f(\sin x, \cos x, \tan x, \cot x, \sec x, \operatorname{cosec} x)$

g1 ($\sin x, \cos x$) ; g2 ($\tan x, \operatorname{cosec} x$) ; g3 ($\cot x, \sec x$)

$$B \int f(\tan x) dx = \int \frac{f(y)}{1+y^2} dy \dots$$

$$\int \frac{x^4}{(1-x^2)^{5/2}} dx \rightarrow \int \frac{\sin^4 y}{\cos^4 y} dy$$

<https://www.youtube.com/watch?v=PNKj529yY5c&t=2448s>

https://en.wikipedia.org/wiki/Integration_by_parts

Example from Patrick Winston's AI lecture

$$\int \frac{-5x^4}{(1-x^2)^{5/2}} dx \rightarrow \int \frac{5x^4}{(1-x^2)^{5/2}} dx \rightarrow \int \frac{x^4}{(1-x^2)^{5/2}} dx \rightarrow \int \frac{\sin^4 y}{\cos^4 y} dy$$

$$\int \frac{1}{\cot^4 y} dy \text{ OR } \int \tan^4 y$$

$$\int \frac{\sin^4 y}{\cos^4 y} dy \rightarrow \int \tan^4 y \rightarrow \int \frac{z^4}{1+z^2} dz \rightarrow \int z^2 - 1 + \frac{1}{1+z^2} dz$$

AND-OR Tree

<https://www.youtube.com/watch?v=PNKj529yY5c&t=2448s>

https://en.wikipedia.org/wiki/Integration_by_parts

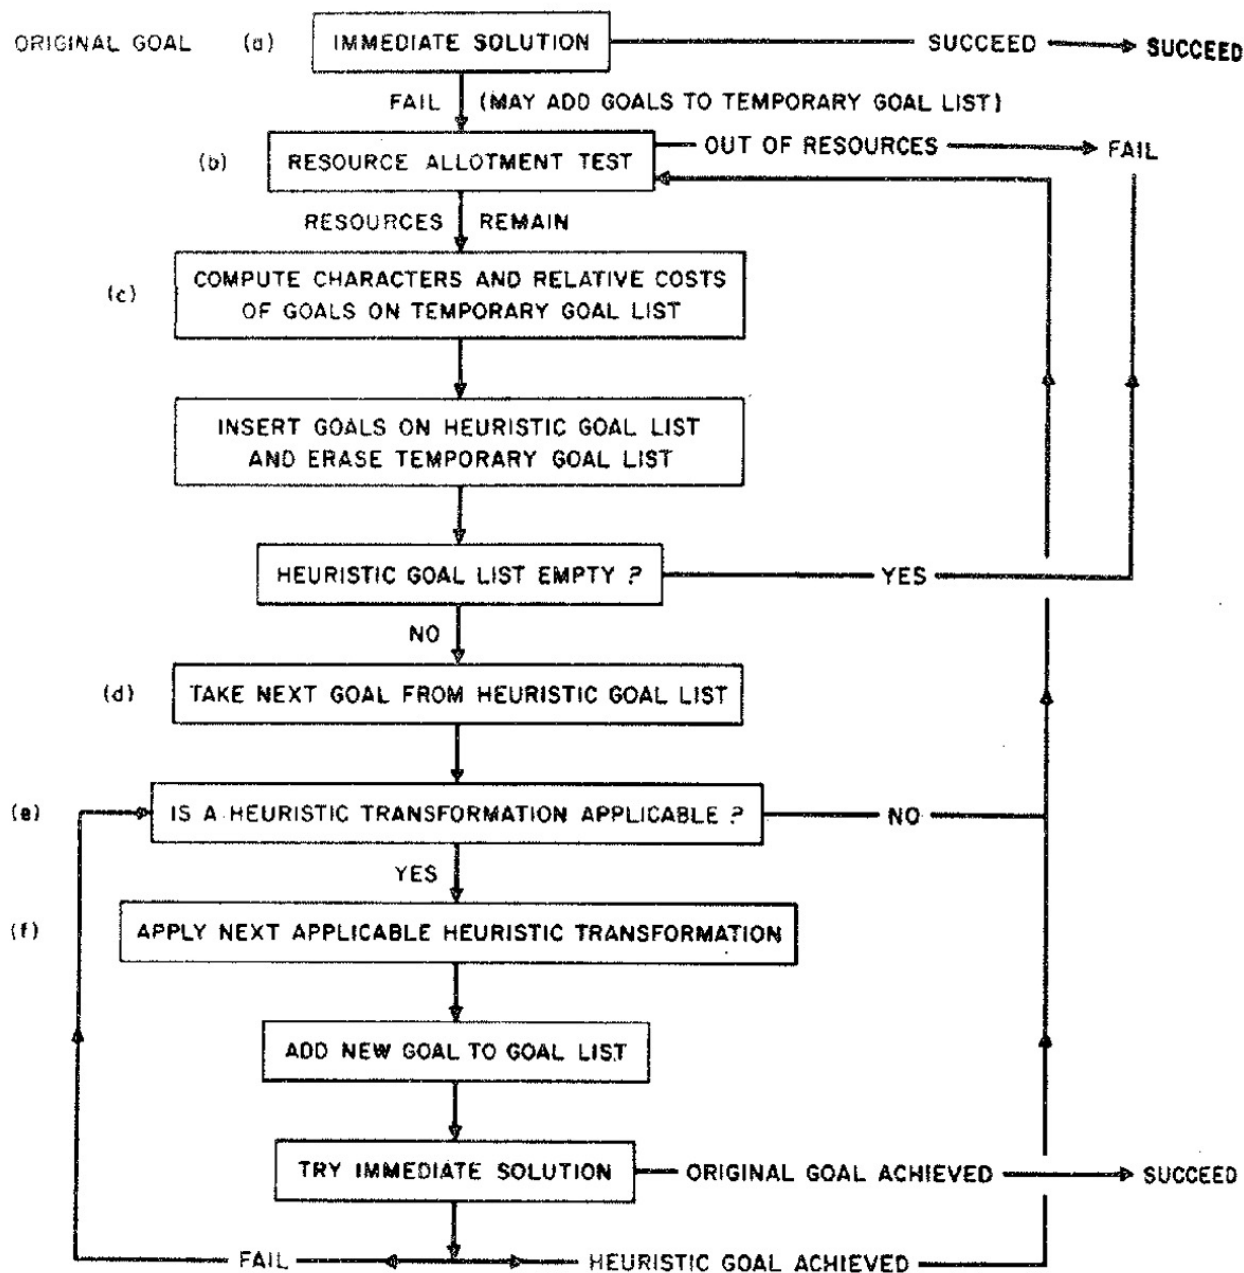


FIG. 3. Executive organization of the indefinite integration procedure

ELIZA

ELIZA is an early natural language processing computer program created from 1964 to 1966 at the MIT Artificial Intelligence Laboratory by Joseph Weizenbaum. Created to demonstrate the superficiality of communication between man and machine, Eliza simulated conversation by using a 'pattern matching' and substitution methodology that gave users an illusion of understanding on the part of the program, but had no built in framework for contextualising events.

<https://www.youtube.com/watch?v=RMK9AphfLco>

<https://web.njit.edu/~ronkowitz/eliza.html>

<http://psych.fullerton.edu/mbirnbaum/psych101/eliza.htm>

Harder than originally thought

- 1966: [Eliza](#) chatbot (Weizenbaum)
 - “ ... mother ...” → “**Tell me more about your family**”
 - “**I wanted to adopt a puppy, but it’s too young to be separated from its mother.**”
 - 1954: [Georgetown-IBM experiment](#)
 - Completely automatic translation of more than sixty Russian sentences into English
 - Only six grammar rules, 250 vocabulary words, restricted to organic chemistry
 - Promised that machine translation would be solved in three to five years ([press release](#))
-

...underwhelming results

- Automatic Language Processing Advisory Committee (ALPAC) report (1966): machine translation has failed

- Example:

The spirit is willing but the flesh is weak.

(Russian)


The vodka is good but the meat is rotten.

- ALPAC report cut government funding for MT, first AI winter
-

A dose of reality (1966-1973)


- Simons's over-confidence was due to the promising performance of the early AI systems on simple examples. In almost all cases, however, the early systems turned out to fail miserably when tried out on wider selections of problems and on more difficult problems
 - Problems:
 - Most of early programs contained little or no knowledge of their subject matter, they succeeded by means of simple syntactic manipulations (e.g. translation)
 - Intractability of problems that AI was attempting to solve (genetic algorithms)
 - Fundamental limitations on the basic structures being used to generate intelligent behavior (perceptrons – although they learn anything that they could represent, they could represent very little)
-

Blocks world (1960s – 1970s)




(a) Original picture.

(b) Differentiated picture.



(c) Line drawing.

(d) Rotated view.



???

Larry Roberts, MIT, 1963

History of AI: Taste of failure

- 1940s
 - First model of a neuron (W. S. McCulloch & W. Pitts)
 - Hebbian learning rule
 - Cybernetics
 - 1950s
 - Turing Test
 - Perceptrons (F. Rosenblatt)
 - Computer chess and checkers (C. Shannon, A. Samuel)
 - Machine translation (Georgetown-IBM experiment)
 - Theorem provers (A. Newell and H. Simon, H. Gelernter and N. Rochester)
 - Late 1960s
 - Machine translation deemed a failure
 - Neural nets deprecated (M. Minsky and S. Papert, 1969)*
 - Early 1970s Intractability is recognized as a fundamental problem
 - Late 1970s The first [“AI Winter”](#)
 - *[A sociological study of the official history of the perceptrons controversy](#)
-

Implications of early era

- Problems:
 - Limited computation: search space grew exponentially, outpacing hardware ($100! \sim 10^{157} > 10^{80}$)
 - Limited information: complexity of AI problems (number of words, objects, concepts in the world)
 - Contributions:
 - Lisp, garbage collection, time-sharing (John McCarthy)
 - Key paradigm: separate modeling and inference
-

Knowledge based systems: The key to power (1969-1979)

- general purpose search mechanisms trying to string together elementary reasoning steps to find complete solutions
 - Weak methods – general but cannot scale up
 - Alternative – more powerful, domain specific knowledge
 - DENDRAL system – inferring molecular structure from the information provided by a mass spectrometer
 - Expert systems – medical diagnosis
-

Knowledge-based systems (70-80s)



Expert systems: elicit specific domain knowledge from experts in form of rules:

if [premises] then [conclusion]

- **DENDRAL**: infer molecular structure from mass spectrometry
- **MYCIN**: diagnose blood infections, recommend antibiotics
- **XCON**: convert customer orders into parts specification;
- save DEC \$40 million a year by 1986

Knowledge-based systems

- Contributions:
 - First real application that impacted industry
 - Knowledge helped curb the exponential growth
 - Problems:
 - Knowledge is not deterministic rules, need to model uncertainty
 - Requires considerable manual effort to create rules, hard to maintain
 - 1987: Collapse of Lisp machines and second AI winter
-

History of AI to the present day


- 1980s Expert systems boom
- Late 1980s- Expert system bust; the second “AI winter”
- Early 1990s
- Mid-1980s Neural networks and back-propagation
- Late 1980s Probabilistic reasoning on the ascent
- 1990s-Present Machine learning everywhere
 - Big Data
 - Deep Learning

[History of AI on Wikipedia](#)

[AAAI Timeline](#)

[Building Smarter Machines: NY Times Timeline](#)

Artificial neural networks




1943: introduced artificial neural networks, connect neural circuitry and logic (McCulloch/Pitts)



1969: Perceptrons book showed that linear models could not solve XOR, killed neural nets research (Minsky/Papert)

Training networks



1986: popularization of backpropagation for training multi-layer networks (Rumelhardt, Hinton, Williams)



1989: applied convolutional neural networks to recognizing handwritten digits for USPS (LeCun)

Deep learning




AlexNet (2012): huge gains in object recognition; transformed computer vision community overnight



AlphaGo (2016): deep reinforcement learning, defeat world champion Lee Sedol

Two intellectual traditions



- AI has always swung back and forth between the two
 - Deep philosophical differences, but deeper connections (McCulloch/Pitts, AlphaGo)?
-

A melting pot

- Bayes rule (Bayes, 1763) from **probability**
- Least squares regression (Gauss, 1795) from **astronomy**
- First-order logic (Frege, 1893) from **logic**
- Maximum likelihood (Fisher, 1922) from **statistics**
- Artificial neural networks (McCulloch/Pitts, 1943) from **neuro-science**
- Minimax games (von Neumann, 1944) from **economics**
- Stochastic gradient descent (Robbins/Monro, 1951) from **optimization**
- Uniform cost search (Dijkstra, 1956) from **algorithms**
- Value iteration (Bellman, 1957) from **control theory**