Artificial Intelligence CS365

Amitabha Mukerjee

What is intelligence

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
  - "Can machines think?"

- Imitation Game

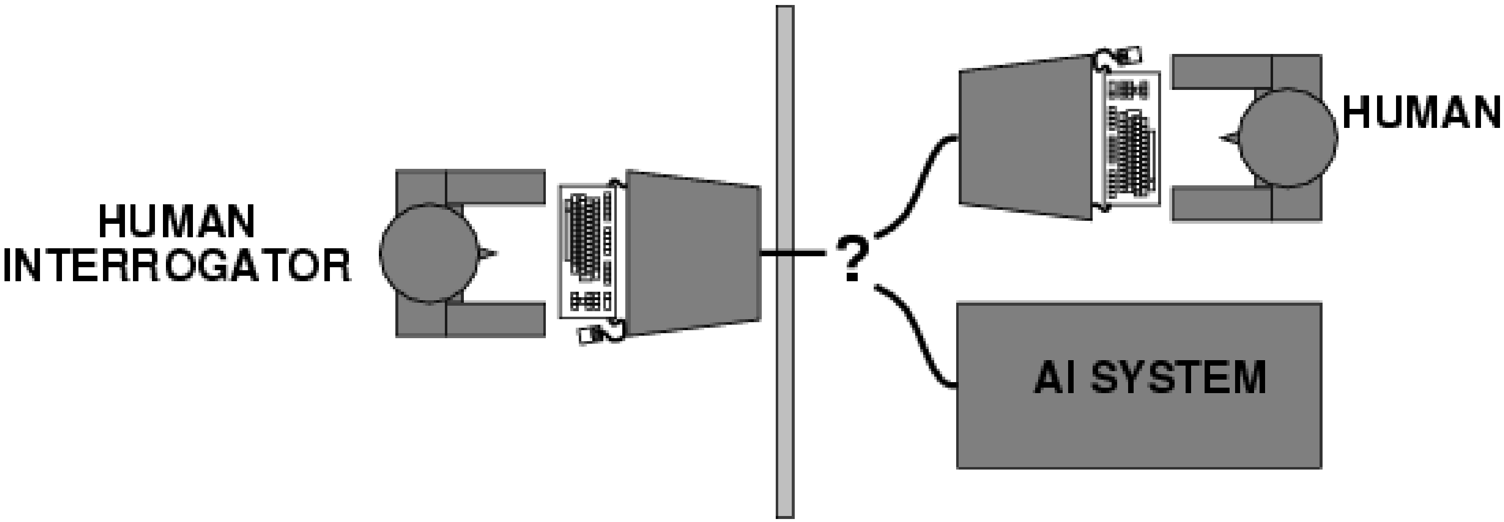

Acting humanly: Turing Test

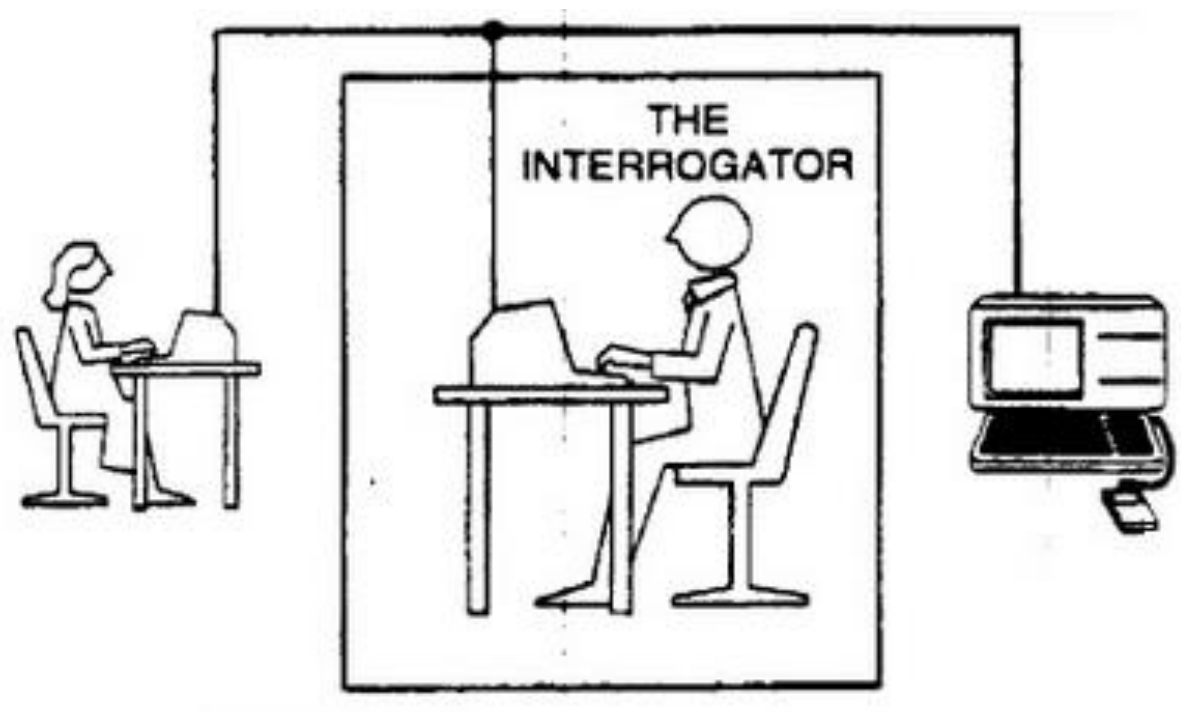

What is AI?

four views:

Think like a human	Think rationally
Act like a human	Act rationally

Subject matters in AI

Get machines to do what humans do but machines can't

Unlike any classical subject, the frontier of what is AI is not static.

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- Greek philosphers: forms of *logic*:
  3-step *syllogism*
- Indian philosophy: 5-step inference
- Problem:
  - Most intelligent behavior does not rely on logical deliberation

Perception

Kanisza triangle

AI history

Timeline : Prehistory / Early AI

- Pre-history: Pascal, Leibniz

  hoaxes

  Babbage

- 1943 McCulloch & Pitts: Boolean circuit model of neuron

- 1950 Turing's "Computing Machinery and Intelligence“

- 1956 Dartmouth meeting: “Artificial Intelligence“ name

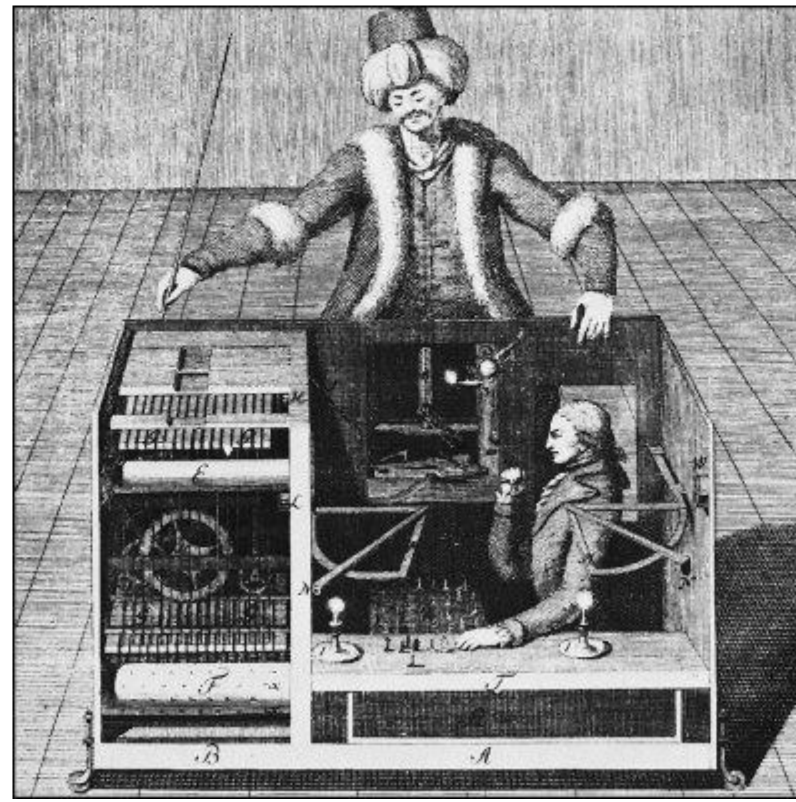

von kempelen’s chess-playing turk, 1769 (hoax)

1955: coining the name “Artificial Intelligence”

John McCarthy,
Marvin Minsky,
N Rochester, and
Claude Shannon:
(1955) :

“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.”

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.

Timeline : AI – Logical Models

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" name
- 1956 Newell & Simon's Logic Theorist,
- 1959 Samuel's checkers program: learned by playing itself

1956 : Logic Theorist

Herbert Simon
&
Alan Newell:

The Logic Theorist 1956

proved 38 of 52 theorems
in ch. 2
Principia Mathematica.
co-author of journal
submission based on a
more elegant proof.
paper was rejected..

Timeline : AI – Logical Models

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" name
- 1956 Newell & Simon's Logic Theorist

- 1959 Samuel's checkers program: game search; later version learned by playing itself
- 1964-66 ELIZA (psychotherapist) by Joseph Weizenbaum

1966 : ELIZA (Social)

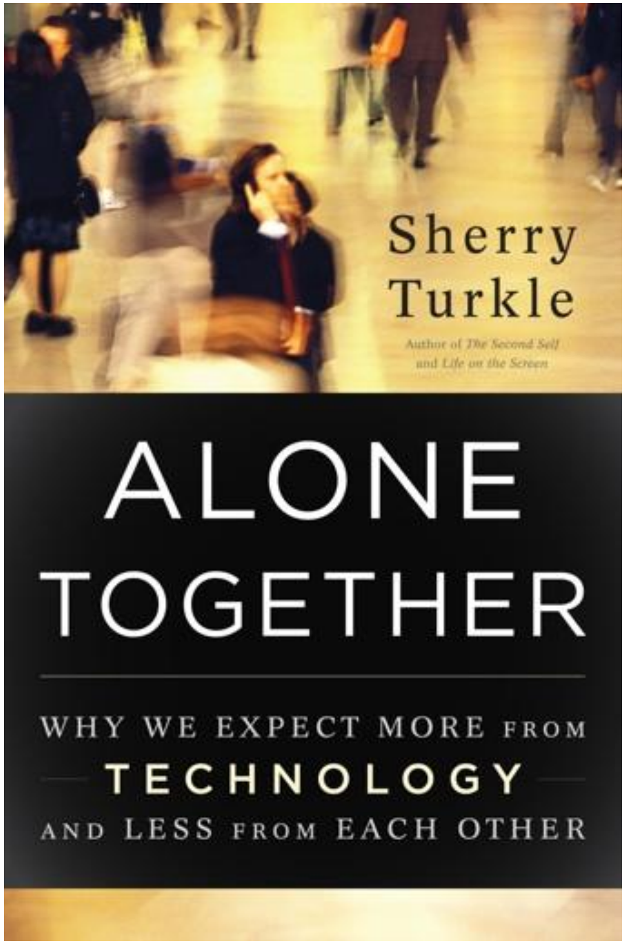

My first brush with a computer program that offered **companionship** was in the mid-1970s. I was among MIT students using Joseph Weizenbaum's ELIZA, a program that engaged in dialogue in the style of a psychotherapist …

Weizenbaum's students **knew that the program did not understand**; nevertheless, they wanted to chat with it. ... they wanted to be alone with it. They wanted to **tell it their secrets**.

- Sherry Turkle, MIT Sociologist

Timeline : AI – Logical Models

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" name
- 1956 Newell & Simon's Logic Theorist

- 1959 Samuel's checkers program: game search; later version learned by playing itself
- 1964-66 ELIZA (psychotherapist) by Joseph Weizenbaum

- 1965 Robinson's resolution algorithm for first order logic

- 1969 Minsky / Papert's *Perceptron*
- 1970-1975 Neural network research almost disappears; [*sociology of science* study]
- 1966-72 Shakey the robot
- 1969-79 Early knowledge-based systems (expert systems)

1958: Rosenblatt - Perceptrons

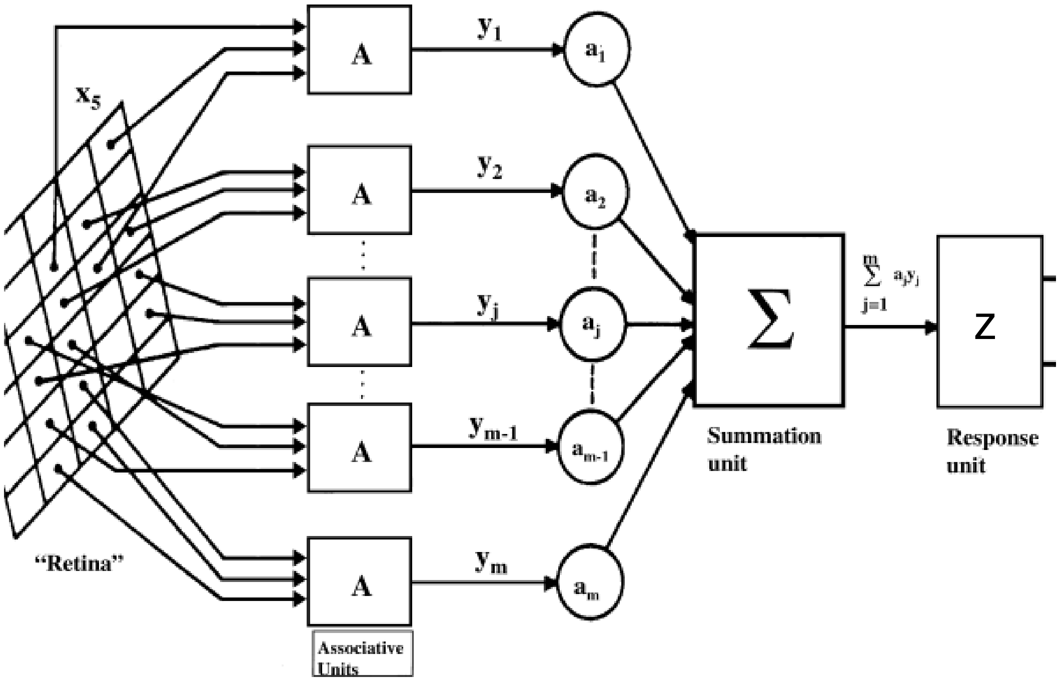

if ∑> *θ, response z = 1, else zero*

$\Delta\theta$ = - (t – z) [t = correct response]

Δw_i = - (t – z) y_i

if z=1 when t=0; then increase θ, and decrease w_i for all positive inputs y_i

The hype of AI

Rosenblatt's press conference 7 July 1958:

The perceptron, an electronic computer [was revealed today] that

z

- will be able to walk, talk, see, write, reproduce itself
- be conscious of its existence.

Later perceptrons will be able to

- recognize people and call out their names
- instantly translate speech in one language to speech and writing in another

1969: Minsky / Papert: Perceptrons

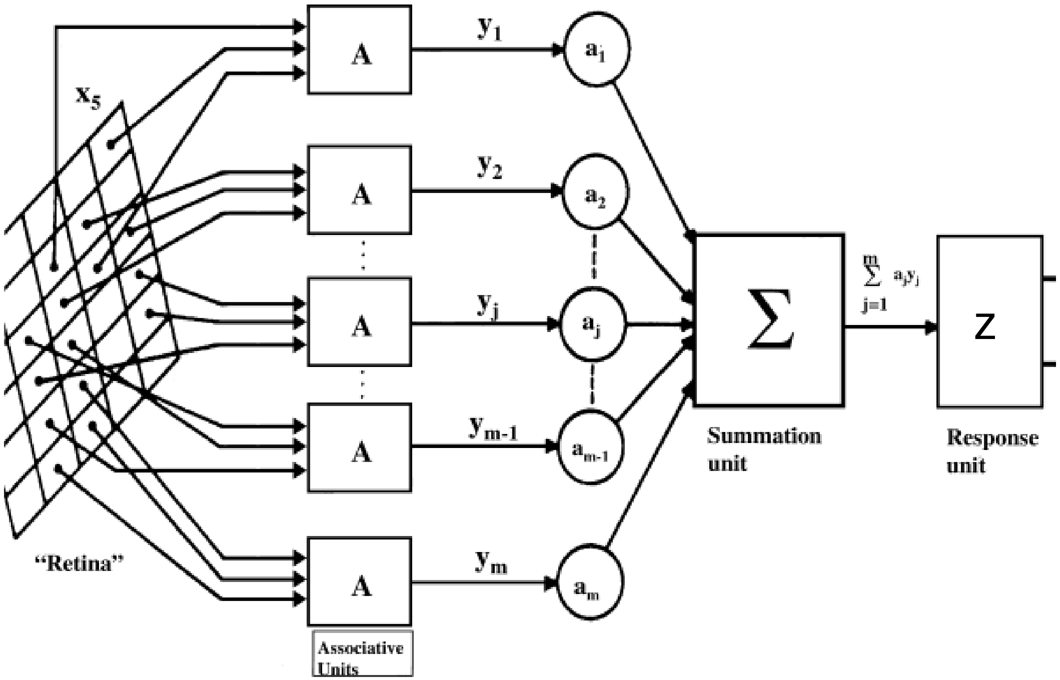

if ∑> *θ, response z = 1, else zero*

$\Delta\theta$ = - (t – z) [t = correct response]

Δw_i = - (t – z) y_i

if z=1 when t=0; then increase θ, and decrease w_i for all positive inputs y_i

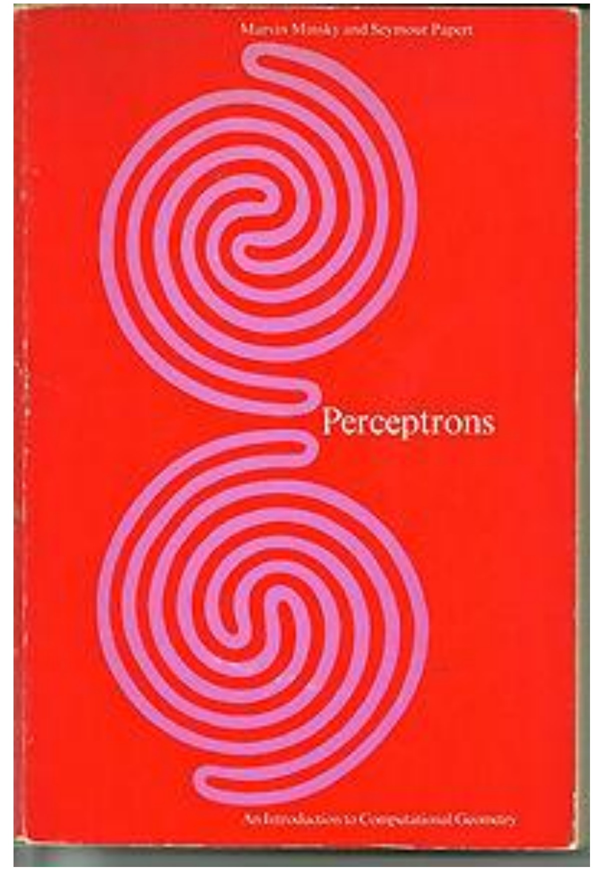

A single-layer perceptron can't learn XOR.

requires

$w_1 > 0$, $w_2 > 0$ but $w_1+w_2 < 0$

Shakey the Robot : 1972

Stanford SRI 1966-1972

STRIPS: planner

Richard Fikes

Nils Nilsson

States (propositions)

Actions (pre-condition, post-condition)

Initial / Goal states

Problem w post-conditions: which states are persistent?

→ **Frame Problem**

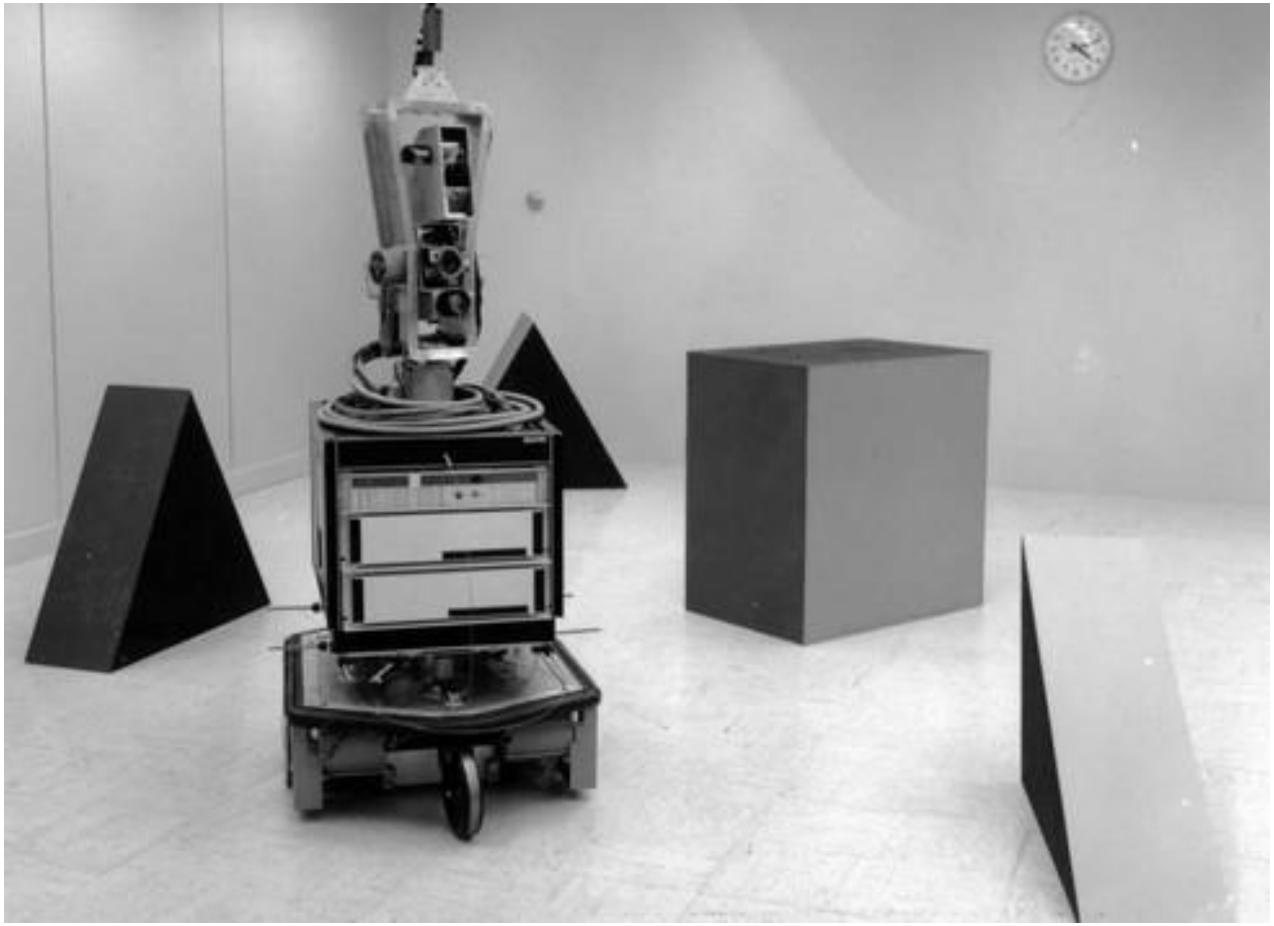

Timeline : AI – Logical Models

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" name
- 1956 Newell & Simon's Logic Theorist
- 1959 Samuel's checkers program: game search; later version learned by playing itself
- 1964-66 ELIZA (psychotherapist) by Joseph Weizenbaum

- 1965 Robinson's resolution algorithm for first order logic
- 1969 Minsky / Papert's *Perceptron*
- 1970-1975 Neural network research almost disappears; [*sociology of science* study]
- 1966-72 Shakey the robot

- 1964-82 Mathlab / Macsyma : symbolic mathematics
- 1969-79 Early knowledge-based systems (expert systems)

“Expert” systems

DENDRAL 1969:

Expert knowledge for chemical structure

Ed Feigenbaum, Bruce Buchanan Joshua Lederberg

Input:

Chemical formula + ion spectrum from mass spectrometer

Output:

Molecular structure

recognizing ketone (C=O) :

if there are two peaks at x1 and x2 s.t.

(a) x1 + x2 = M +28 (M = molecule mass)

(b) x1−28 is a high peak;

(c) x2−28 is a high peak;

(d) At least one of x1 and x2 is high.

then there is a ketone subgroup

Reduces search by identifying some constituent structures

Timeline : AI – Learning

- 1986 Backpropagation algorithm : Neural networks become popular
- 1990-- Statistical Machine Learning

- 1991 *Eigenfaces :* face recognition [Turk and Pentland]
- 1995 [Dickmanns]: 1600km driving, 95% autonomous
  CMU *Navlab*: 5000km 98% autonomous

- 1996 EQP theorem prover finds proof for Robbins' conjecture
- 1997 Deep Blue defeats Kasparov
- 1997 *Dragon Naturally Speaking* speech recognition
- 1999 SIFT local visual feature model
- 2001 [Viola & Jones] : real time face detection
- 2007 DARPA Urban challenge (autonomous driving in traffic)
- 2010 *Siri* speech recognition engine
- 2011 *Watson* wins quiz show *Jeopardy*

Agent Models

Models of Agency

- **Agent** : function from percept histories to actions:

  $$[f: \mathcal{P} \rightarrow \mathcal{A}]$$

  - Intermediate: Precepts → concept categories
  - Goal : measure of performance [utility]

- Rational agent: one that has best performance

  → utility maximization
  → within computational limitations

Task / Environment

- [$f: \mathcal{P} \rightarrow \mathcal{A}$]

- What are precepts / actions for
  - Bicycle riding
  - Writing notes
    - Language decisions
    - Motor actions
  - Solving a sudoku
  - Drawing a cartoon

Task / Environment

- deterministic
- stochastic
- non-deterministic

- static
- dynamic

- continuous
- discrete

Task / Environment

- fully-observable
- partly-observable
- unobservable

Unobservable Problems

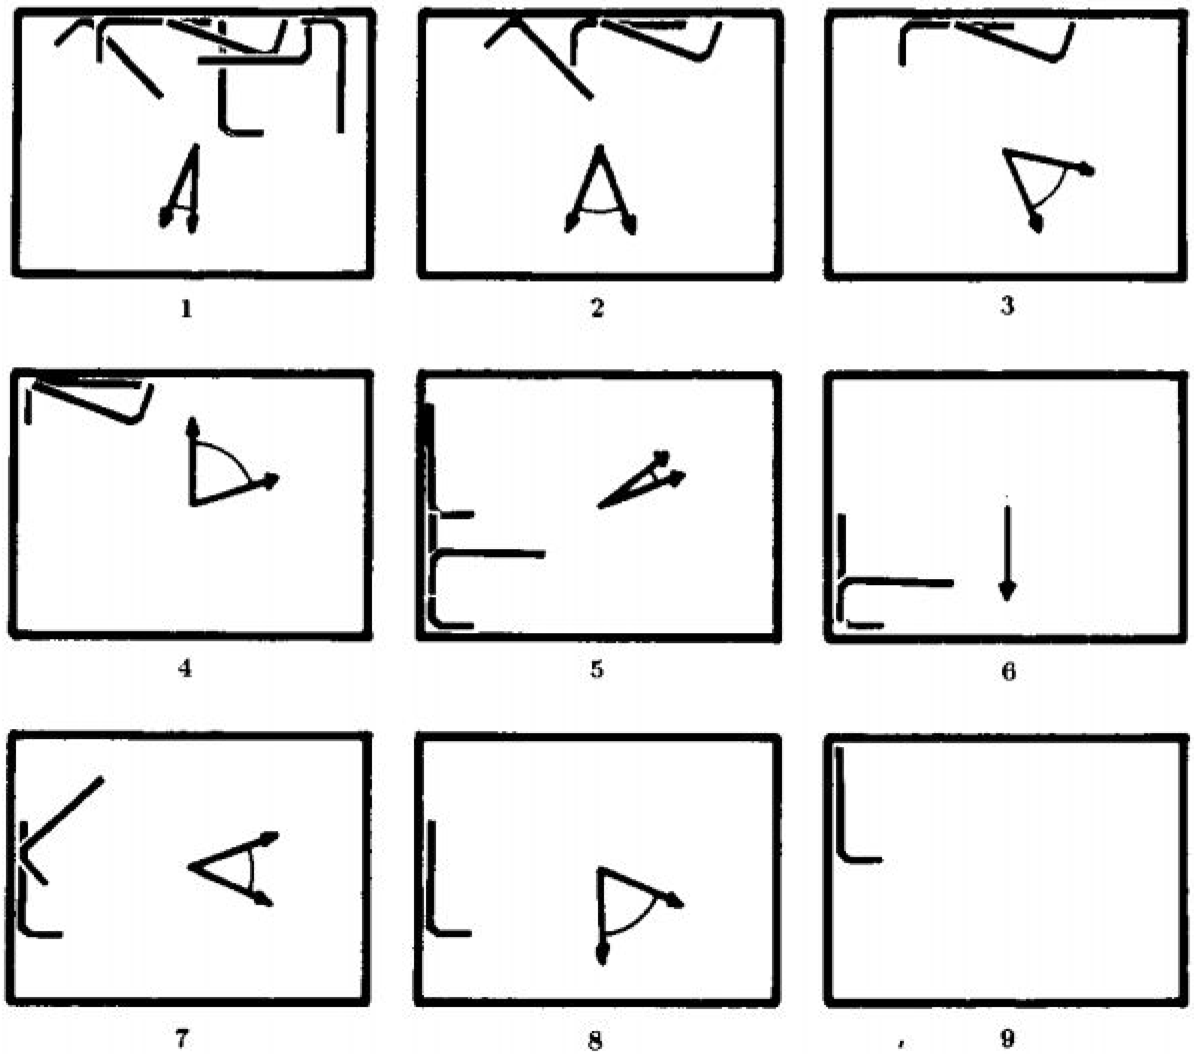

[erdmann / mason 1987]

Task / Environment

- [f: $\mathcal{P} \rightarrow \mathcal{A}$]

- What are precepts / actions for
  - Bicycle riding
  - Writing notes
    - Language decisions
    - Motor actions
  - Solving a sudoku
  - Drawing a cartoon

Learning

- [f: $\mathcal{P} \rightarrow \mathcal{A}$]

- Nature of $\mathcal{P}$ / $\mathcal{A}$:
  - continuous : regression
  - discrete : categorization

- Performance evaluation function?
- Intermediate "features"?

Learning vs Hand-coding

- [f: $\mathcal{P} \rightarrow \mathcal{A}$]

- Should we try to discover the function f, or use what we think will work?
  - Programming may be quicker in the short run
  - Learning : more robust and stable, but may require lots of data