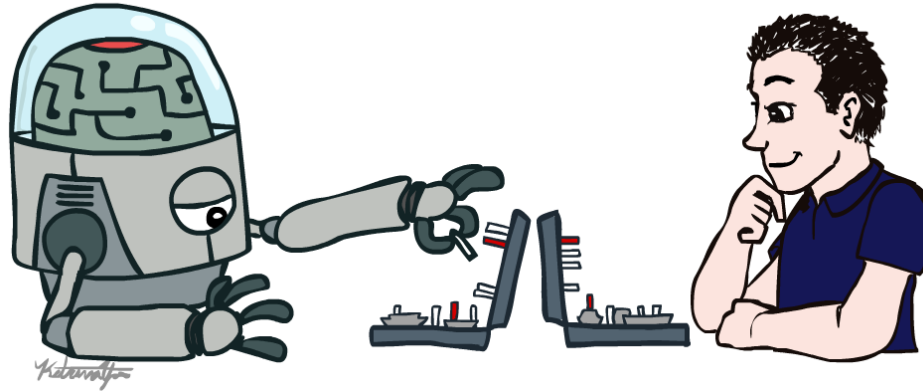


# Introduction to Artificial Intelligence

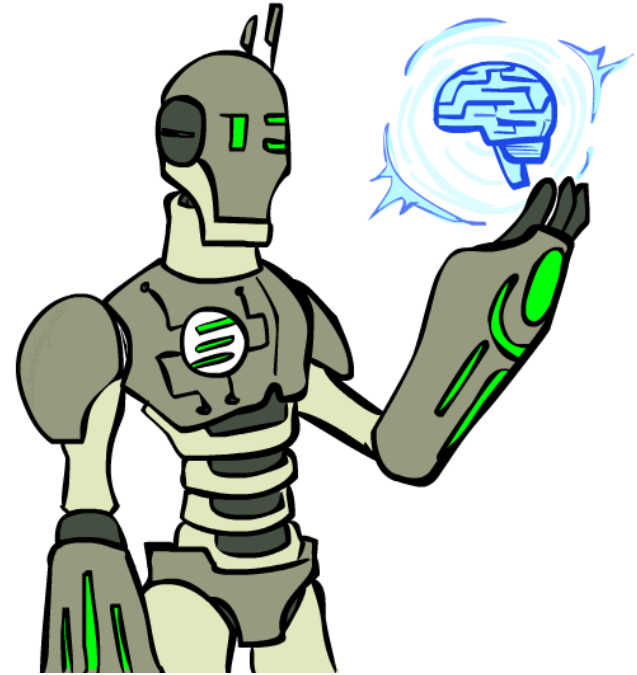
## Introduction



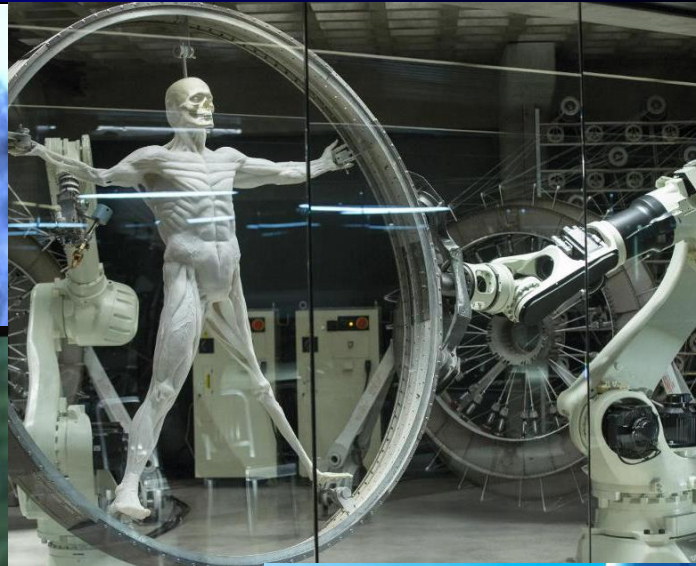
# Second part of Today

---

- What is artificial intelligence?
- Where are we and how did we get here?
- What is this course ?



# Sci-Fi AI?



# Real AI

Text

Documents

DETECT LANGUAGE

CHINESE

ENGLISH

SPANISH



CHINESE (TRADITIONAL)

ENGLISH

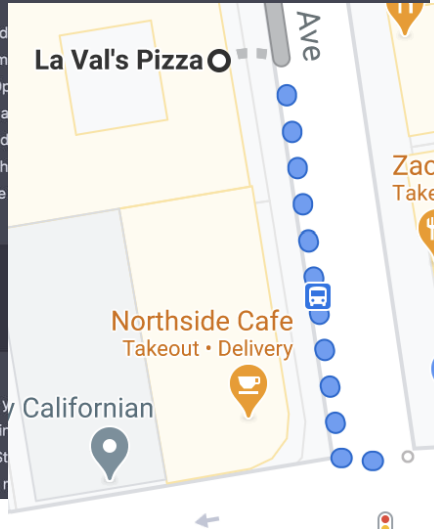
FRENCH

how did we start this conversation

Our conversation began when you asked me to tell you about him, and I provided some information about him. I am a large language model trained by OpenAI. I can answer questions and provide information on a wide range of topics. I can also initiate conversations, and I only respond to the user's input. I do not have the ability to take actions or perform tasks in the real world. My information is based on the text that I have been trained on.

What did we speak about after that?

After you asked me about Post Malone, you asked me to tell you about him, and I provided some information about him. I provided some information about his music and his career. I also provided some information about his personal life and his relationships. I provided some information about his net worth and his social media presence. I provided some information about his discography and his collaborations. I provided some information about his awards and his achievements. I provided some information about his fans and his influence. I provided some information about his impact on the music industry and on popular culture. I provided some information about his legacy and his future prospects. I provided some information about his net worth and his social media presence. I provided some information about his discography and his collaborations. I provided some information about his awards and his achievements. I provided some information about his fans and his influence. I provided some information about his impact on the music industry and on popular culture. I provided some information about his legacy and his future prospects.







TUG  
CAUTION  
MAY CONTAIN  
CHEMOTHERAPY DRUG

CAUTION  
MAY CONTAIN  
CHEMOTHERAPY DRUG







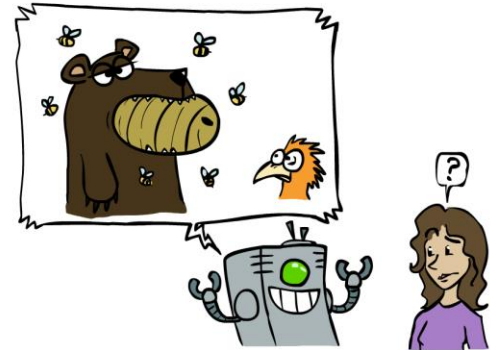
Boston Dynamics





# Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.
- Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.



# AI in the News

---

Perhaps you've heard that there is an exciting new Barack Obama book that everyone's talking about! I'm not talking about *A Promised Land*, the 751-page memoir which Obama spent four years writing.

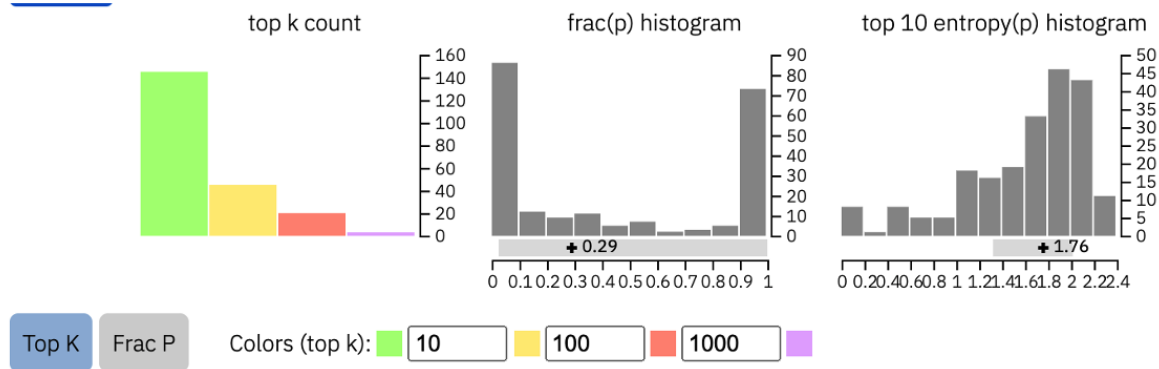
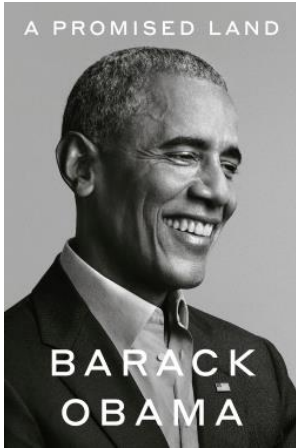
No, I'm talking about *Barack Obama Book*, a 61-page tome by an author named "University Press." Why is *Barack Obama Book* selling so well? Thanks to sponsored listings and canny search engine optimization, the book appears above Barack Obama's actual memoir if you search Amazon for—you guessed it—"barack obama book."

What is Barack Obama Book? It's not a book, exactly. It's an SEO ploy by a shadowy company that has scores of \$2.99 knockoffs ready to be downloaded, from quickie biographies of au courant figures like Obama, Harris, and Dolly Parton, to obvious bestseller ripoffs like *Caste: A Brief History and Ready Player 3, 2, 1: A Brief History (?)*.

I don't think Barack Obama Book was written by a human being, but I do think the A.I. that excreted it made some decent points about Barack Obama.

<https://slate.com/culture/2020/11/what-is-barack-obama-book-university-press-amazon.html>

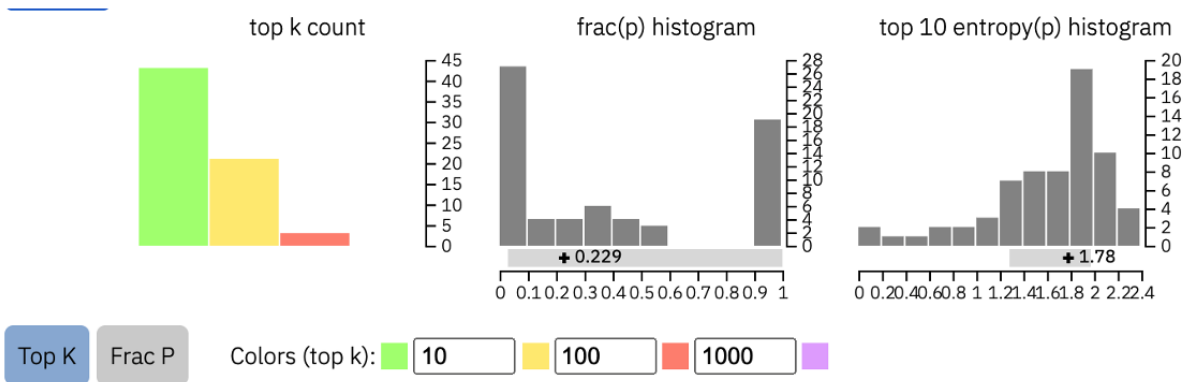
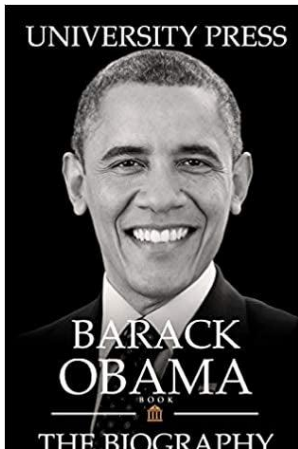
# AI in the News



Based on what I'd heard, I decided we had enough information to begin developing options for an attack on the compound. While the CIA team continued to work on identifying the Pacer, I asked Tom Donlihan and John Brennan to explore what a raid would look like. The need for secrecy added to the challenge; if even the slightest hint of our lead on bin Laden leaked, we knew our opportunity would be lost. As a result, only a handful of people across the entire federal government were read into the planning phase of the operation. We had one other constraint: Whatever option we chose could not involve the Pakistanis. Although Pakistan's government cooperated with us on a host of counterterrorism operations and provided a vital supply path for our forces in Afghanistan, it was an open secret that certain elements inside the country's military, and especially its intelligence services, maintained links to the Taliban and perhaps even to Al-Qaeda, sometimes using them as strategic assets to ensure that the Afghan government remained weak and unable to align itself with Pakistan's number one rival, India.



# AI in the News



One of the most notable events that occurred during his second term was the killing of Osama bin Laden by US Naval personnel. This helped boost Obama's popularity and brought people together within the nation, but it had detrimental repercussions with other nations, particularly Pakistan and its neighbors. The strike had been authorized without consulting the nation where bin Laden lived.

# AI in the News



**Elon Musk** 

@elonmusk



Tesla Full Self-Driving Beta is now available to anyone in North America who requests it from the car screen, assuming you have bought this option.

Congrats to Tesla Autopilot/AI team on achieving a major milestone!

11:34 PM · **Nov 23, 2022**

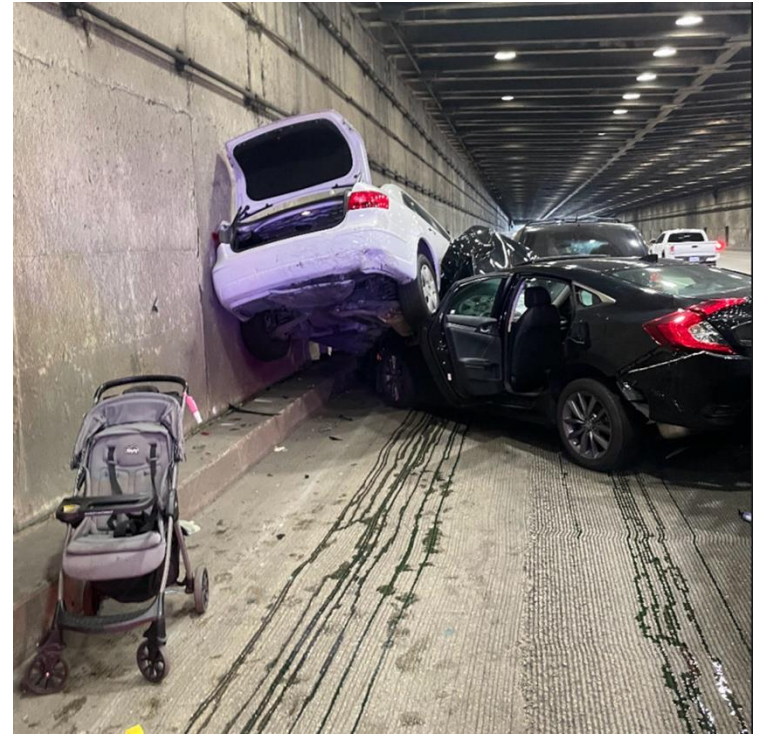
**12.9K** Retweets   **2,651** Quote Tweets   **174.1K** Likes



# AI in the News

Highway surveillance footage from **November 24** shows a Tesla Model S vehicle changing lanes and then abruptly braking in the far-left lane of the San Francisco Bay Bridge, resulting in an eight-vehicle crash.

As traditional car manufacturers enter the electric vehicle market, Tesla is increasingly under pressure to differentiate itself. Last year, Musk said that “Full Self-Driving” was an “essential” feature for Tesla to develop, going as far as saying, “It’s really the difference between Tesla being worth a lot of money or worth basically zero.”





# AI in the News

TECH • ARTIFICIAL INTELLIGENCE

## United Kingdom Plans \$1.3 Billion Intelligence Push

France to spend \$1.8 billion on compete with U.S., China

EU wants to invest £18b development

## China's Got a Huge Artificial Intelligence Plan

### 'Whoever leads in AI will rule the world': Putin to Russian children on Knowledge Day

Published time: 1 Sep, 2017 14:08

Edited time: 1 Sep, 2017 14:40



# AI in the News

TECNOLOGIA | VENERDÌ 17 FEBBRAIO 2023

## Il nuovo chatbot di Microsoft sta dicendo cose stranissime agli utenti

Ai giornalisti che l'hanno messo alla prova ha risposto di sentirsi intrappolato e di amarli, oppure ha provato a ingannarli



 (New York Times/ Instagram)

 A inizio febbraio Microsoft ha aperto a un numero limitato di utenti la possibilità di



[https://www.ilpost.it/2023/02/17/chatbot-microsoft-test/?utm\\_medium=social&utm\\_source=facebook&utm\\_campaign=lancio](https://www.ilpost.it/2023/02/17/chatbot-microsoft-test/?utm_medium=social&utm_source=facebook&utm_campaign=lancio)

# What is AI?

---

- For thousands of years, we have tried to understand how we **think** and **act**.
- AI is concerned with not just understanding but also building **intelligent** entities
- What is intelligence?  
Fidelity to **human** performance *or* **rationality**?  
Internal **thought** processes *or* external **behaviour**?



# What is AI?

---

The science of making machines that:

# A rational-agent approach

---

- An **agent** is just something that acts
- "agent" comes from the Latin *agere*, to do
- A **rational agent** is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.

*AI has focused on the study and construction of agents that **do the right thing***

# Rational Decisions

---

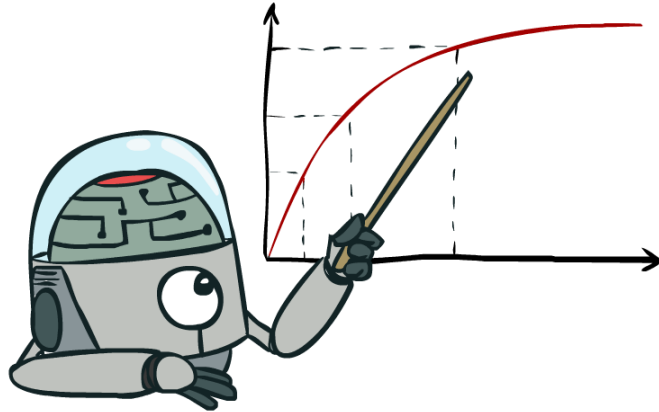
- We'll use the term **rational** in a very specific, technical way:
  - Rational: maximally achieving pre-defined goals
  - Rationality only concerns what decisions are made (not the thought process behind them)
  - Goals are expressed in terms of the **utility** of outcomes
  - Being rational means **maximizing your expected utility**

A better title for this course would be:

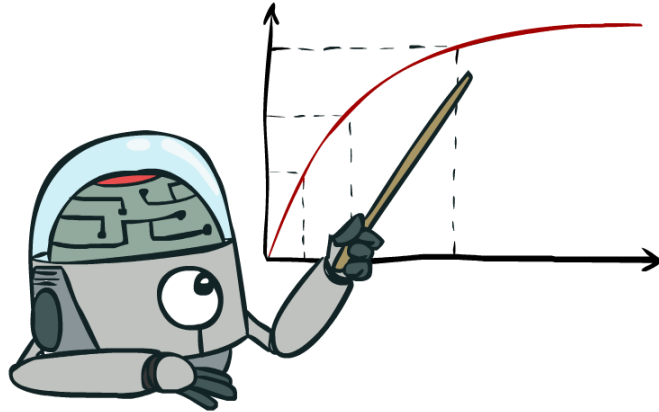
**Computational Rationality**



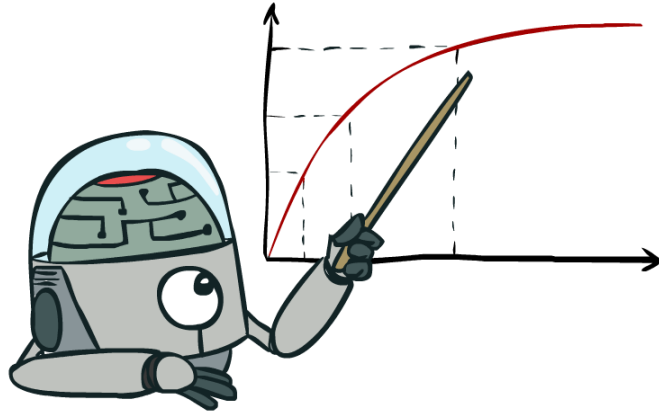
# Maximize Your Expected Utility



# Maximize Your Expected Utility

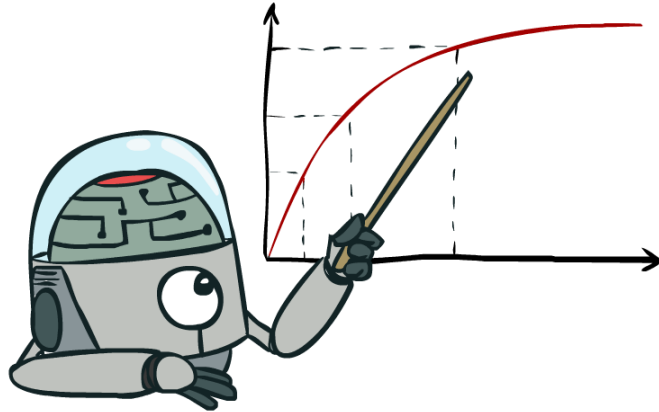


# Maximize Your Expected Utility

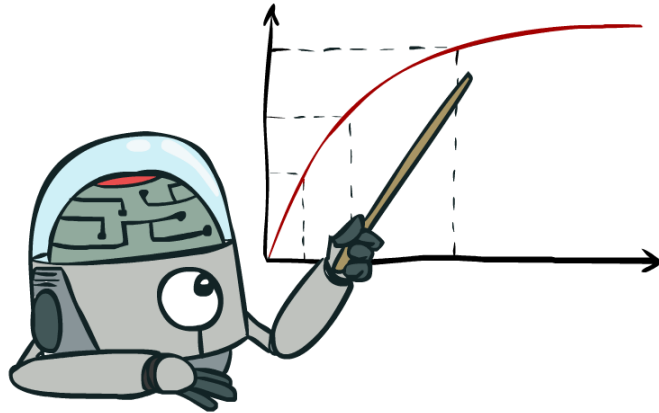




# Maximize Your Expected Utility



# Maximize Your Expected Utility

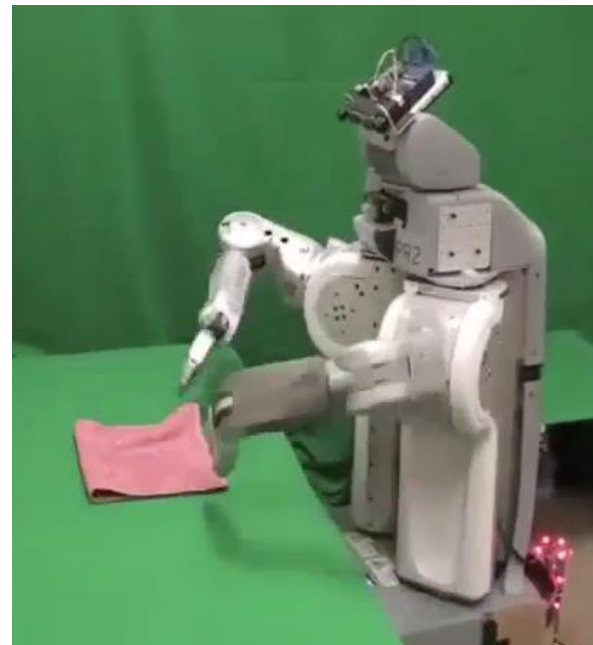


# Utility?

Clear utility function



Not so clear utility function





# What About the Brain?

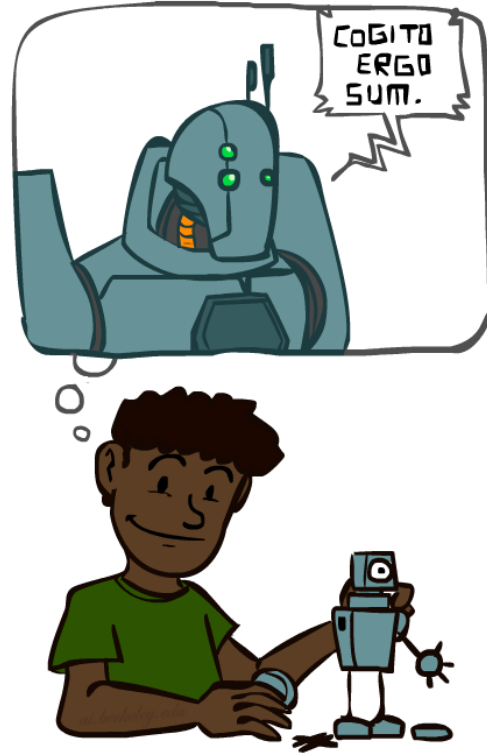
- Brains (human minds) are very good at making rational decisions, but not perfect
- Brains aren't as modular as software, so hard to reverse engineer!
- “Brains are to intelligence as wings are to flight”
- Lessons learned from the brain: memory (data) and simulation (computation) are key to decision making





# A (Short) History of AI

---



# A short prehistory of AI

## ■ Prehistory:

- **Philosophy** (reasoning, planning, learning, science, automation)

- Aristotle: For if every instrument could accomplish its own work, obeying or anticipating the will of others . . . if, in like manner, the shuttle would weave and the plectrum touch the lyre without a hand to guide them, chief workmen would not want servants, nor masters slaves

- **Psychology** (learning, cognitive models)

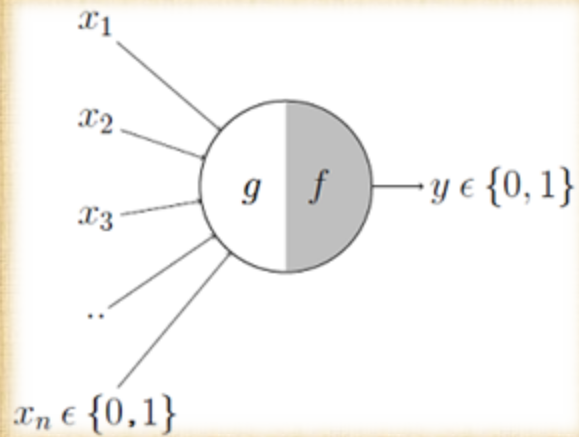
- **Linguistics** (grammars, formal representation of meaning)

## ■ Near miss (1842):

- Babbage design for universal machine (the Analytical Engine)
- Lovelace: “a thinking machine” for “all subjects in the universe.”

# The 40s years

1943: first computational model of a neuron



Walter Pitts

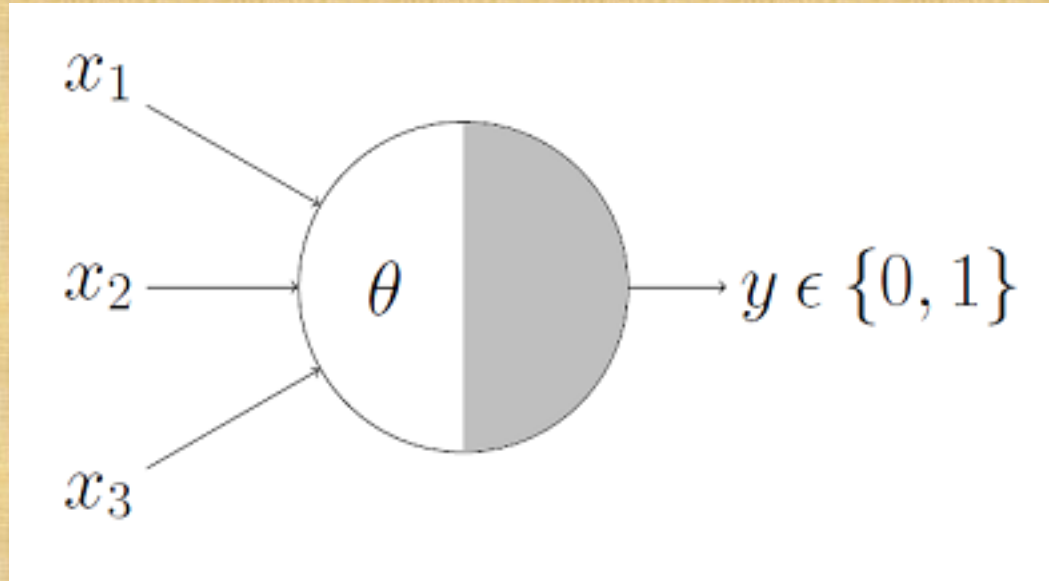


Warren McCulloch

# The 40s years

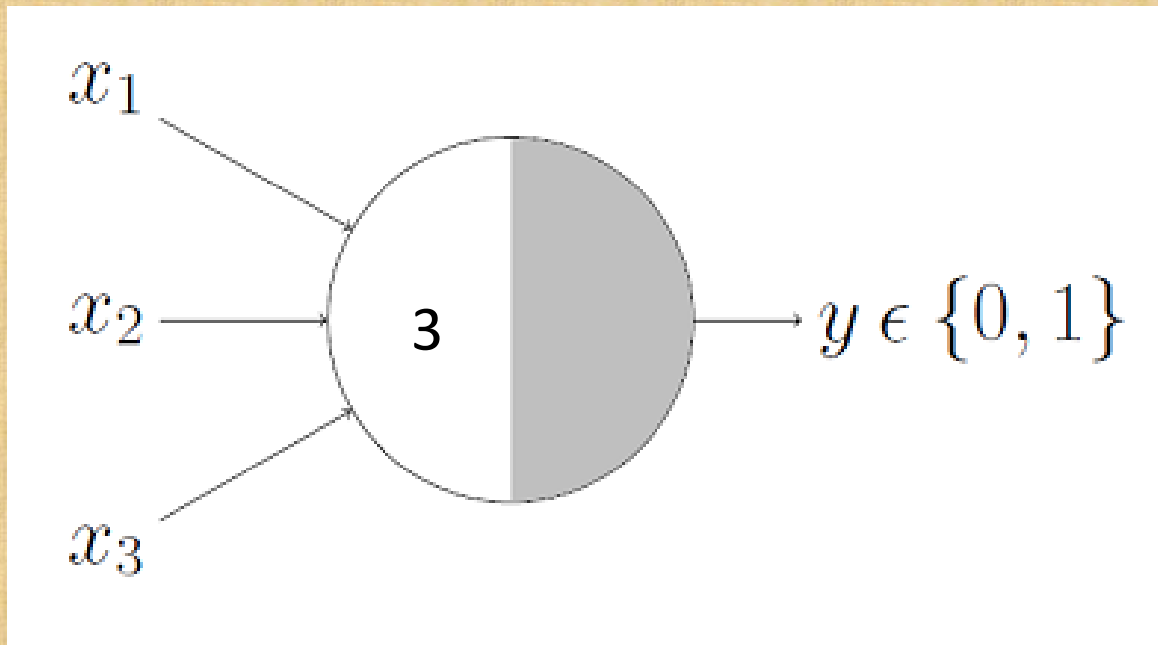
Boolean circuit model:

if  $x_1 + x_2 + x_3 \geq \theta$  than the neuron will fire, otherwise it won't



# The 40s years

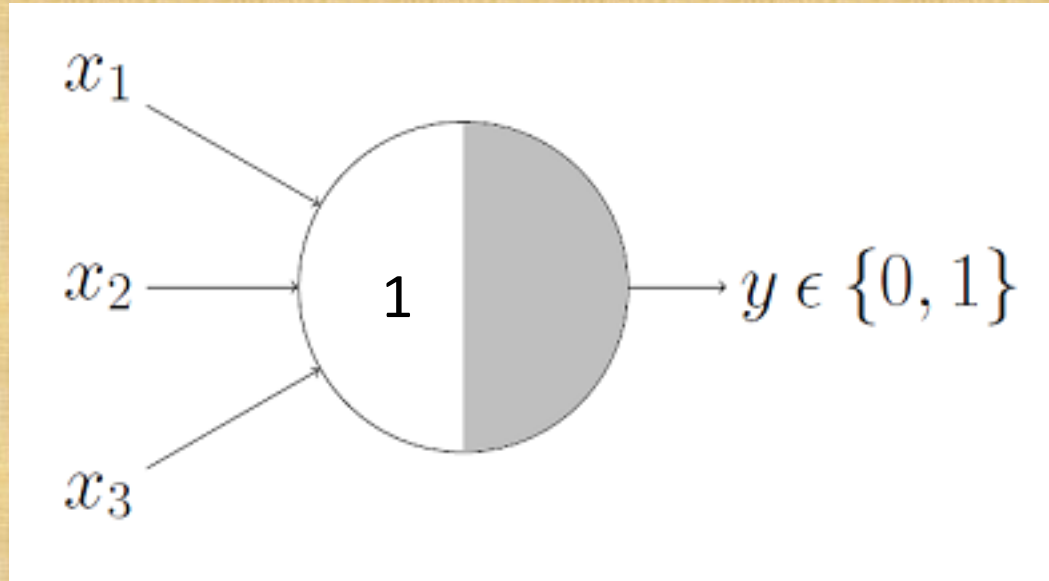
## AND Function





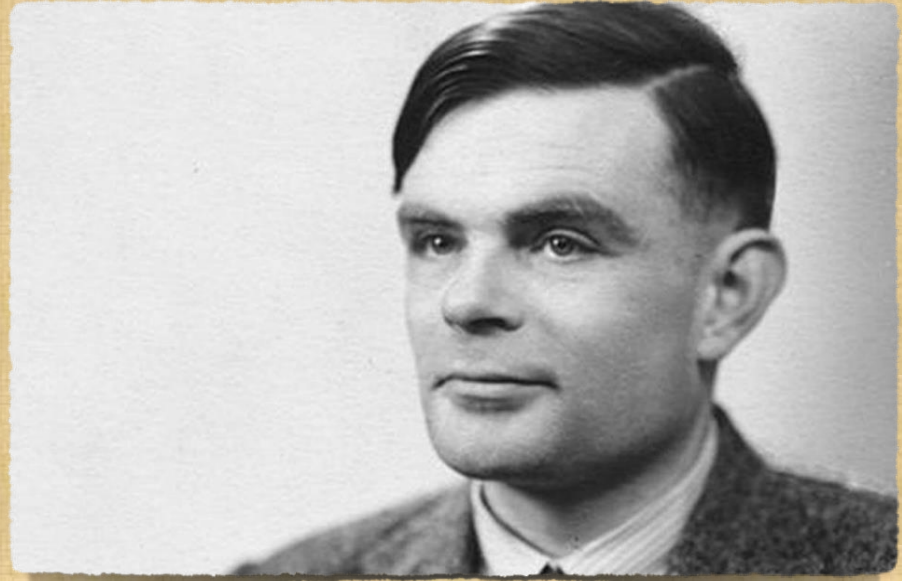
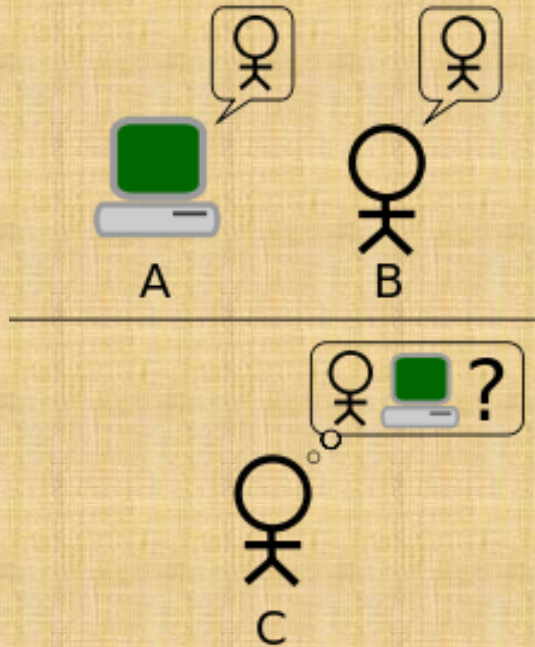
# The 40s years

OR Function

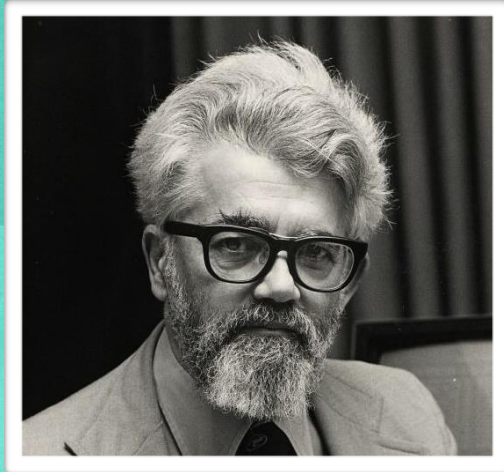


# The 40s years

1950: Turing's "Computing Machinery and Intelligence"



# The 50s years



New

## Participants:

1. Ray Solomonoff
2. Marvin Minsky
3. John McCarthy
4. Claude Shannon
5. Trenchard More
6. Nat Rochester
7. Oliver Selfridge
8. Julian Bigelow
9. W. Ross Ashby
10. W.S. McCulloch
11. Abraham Robinson
12. Tom Etter
13. John Nash
14. David Sayre
15. Arthur Samuel
16. Kenneth R. Shoulders
17. Shoulders' friend
18. Alex Bernstein
19. Herbert Simon
20. Allen Newell

**1956 year**

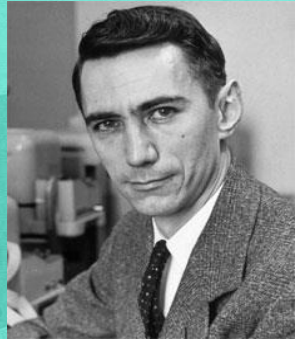
John McCarthy organized a workshop at Dartmouth College

Coined the **Artificial Intelligence** term.

[slides about AI history adapted from Luca Manzoni slides, thank you Luca 😊]



# The 50s years



“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 if we work on it together for a summer.***”

**John McCarthy and Claude Shannon  
Dartmouth Workshop Proposal**

# The 50s

1956, **Logic Theorist**  
is the first program of Artificial  
Intelligence

\*54·43.  $\vdash :: \alpha, \beta \in 1 . \supset : \alpha \cap \beta = \Lambda . \equiv . \alpha \cup \beta \in 2$

*Dem.*

$\vdash . *54·26 . \supset \vdash :: \alpha = \iota'x . \beta = \iota'y . \supset : \alpha \cup \beta \in 2 . \equiv . x \neq y .$

[\*51·231]  $\equiv . \iota'x \cap \iota'y = \Lambda .$

[\*13·12]  $\equiv . \alpha \cap \beta = \Lambda \quad (1)$

$\vdash . (1) . *11·11·35 . \supset$

$\vdash :: (\forall x, y) . \alpha = \iota'x . \beta = \iota'y . \supset : \alpha \cup \beta \in 2 . \equiv . \alpha \cap \beta = \Lambda \quad (2)$

$\vdash . (2) . *11·54 . *52·1 . \supset \vdash . \text{Prop}$

From this proposition it will follow, when arithmetical addition has been defined, that  $1 + 1 = 2$ .



Allen Newell



Herbert A. Simon

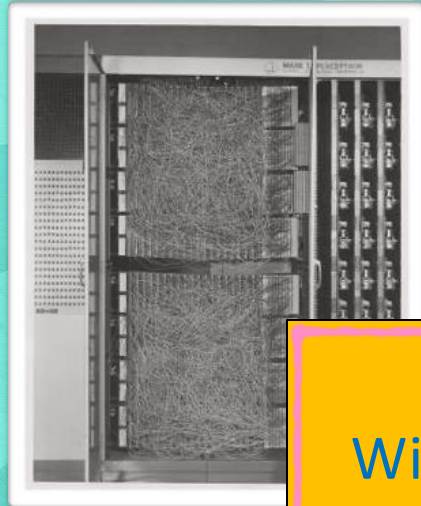


Cliff Shaw



# The 50s years

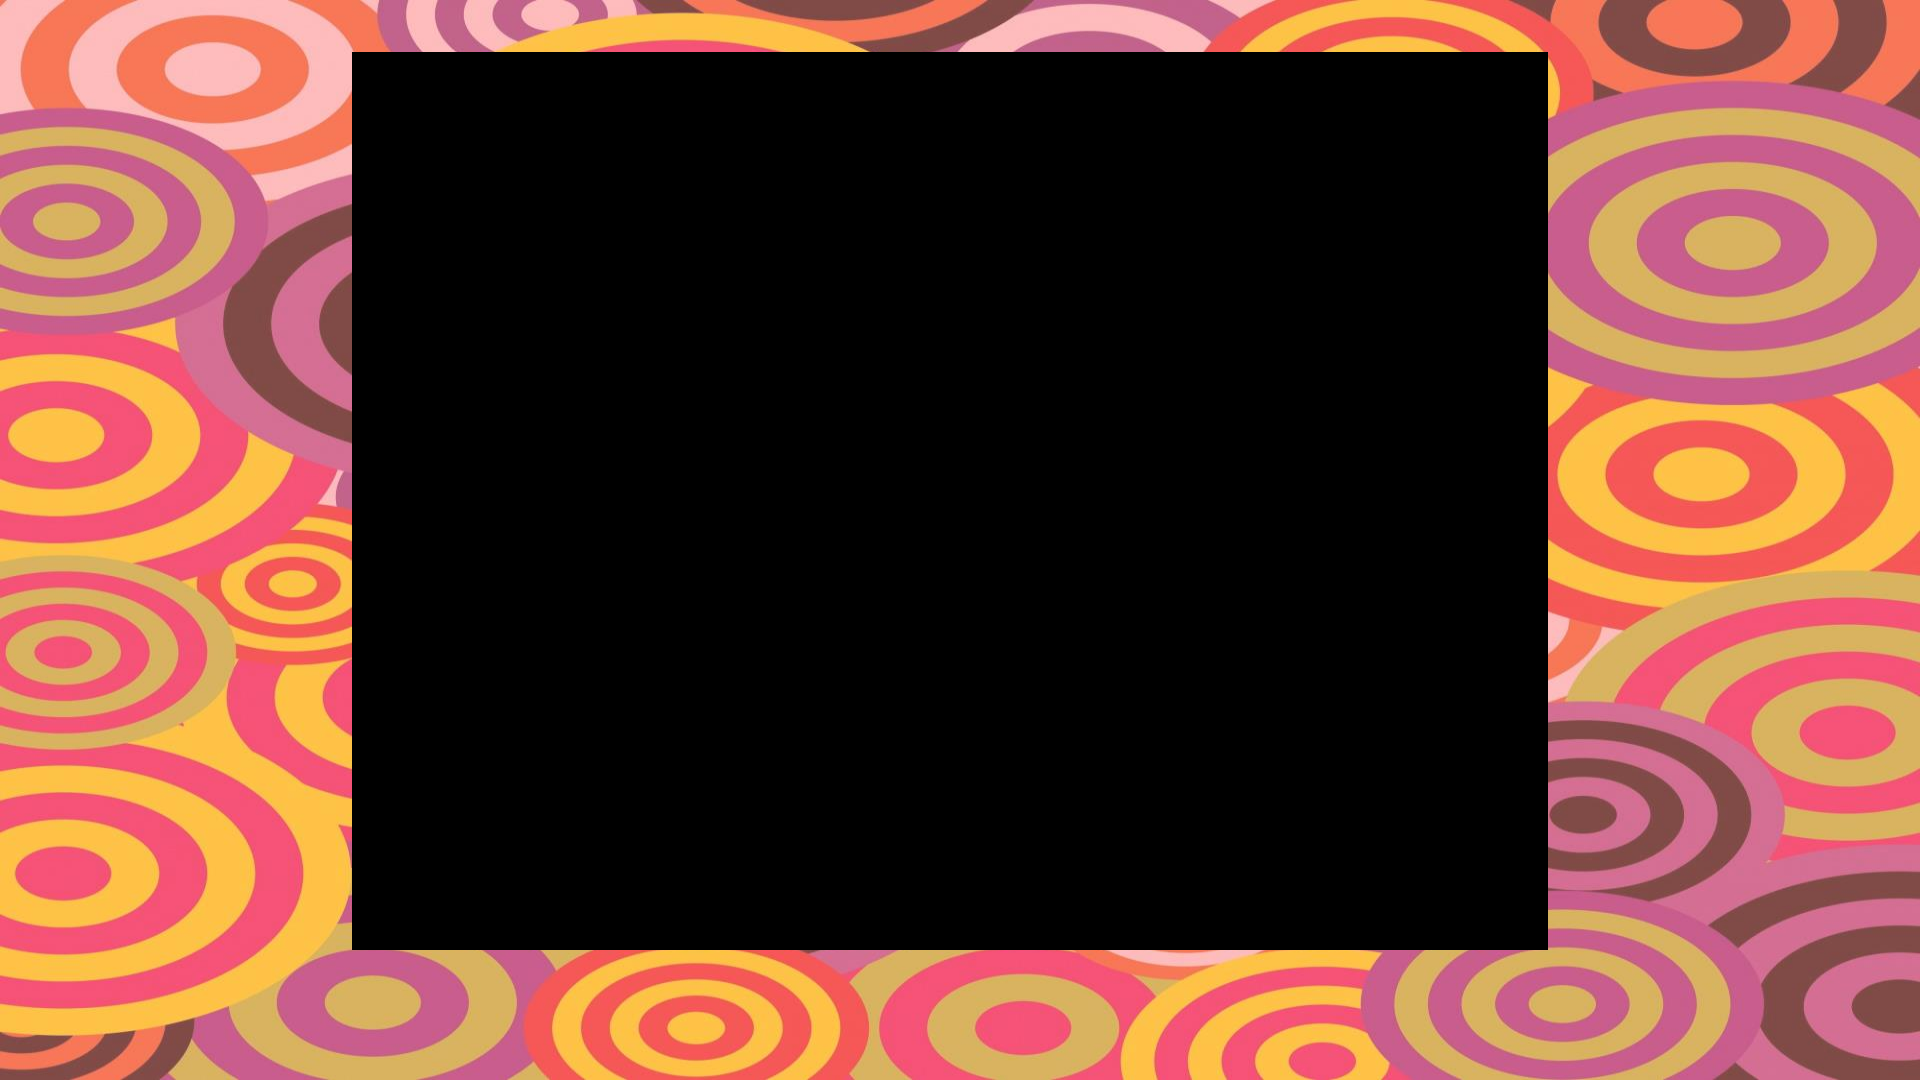
Frank Rosenblatt defines the **perceptron**, a system able to learn inspired by the neurons



400 pixel!  
With  
potentiometers

# The 60s Years







1964

1966

The

60s

Years

```
EEEEEE LL IIII  
EE LL II  
EEEEEE LL II AA  
EE LL II ZZ AA  
EEEEEE LLLLLL IIII ZZZZZZ AA AA
```

Eliza is a mock Rogerian psychotherapist.  
The original program was described by Joseph Weizenbaum in 1966.  
This implementation by Norbert Landsteiner 2005.

```
ELIZA: Is something troubling you ?  
YOU: Men are all alike.  
ELIZA: What is the connection, do you suppose ?  
YOU: They're always bugging us about something or other.  
ELIZA: Can you think of a specific example ?  
YOU: Well, my boyfriend made me come here.  
ELIZA: Is it important to you that your boyfriend made you come here?  
YOU: He says I'm depressed much of the time.  
ELIZA: I am sorry to hear that you are depressed.  
YOU: It's true. I am unhappy.  
ELIZA: Can you explain what made you unhappy ?
```

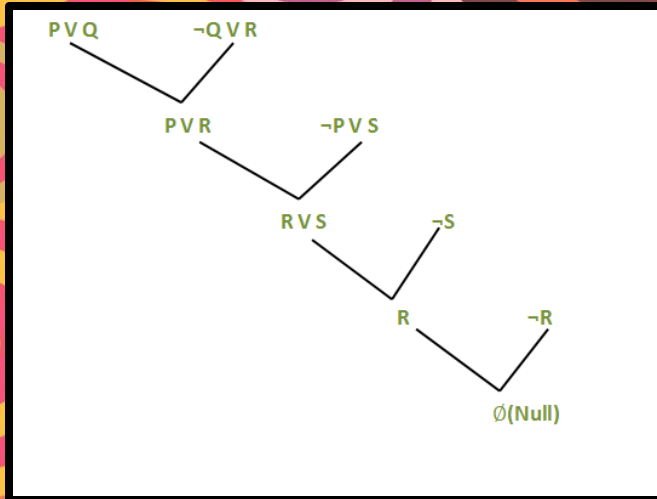
Joseph  
Weizenbaum  
creates ELIZA, one of  
the first programs  
processing  
natural language

1965

The

60s

Years



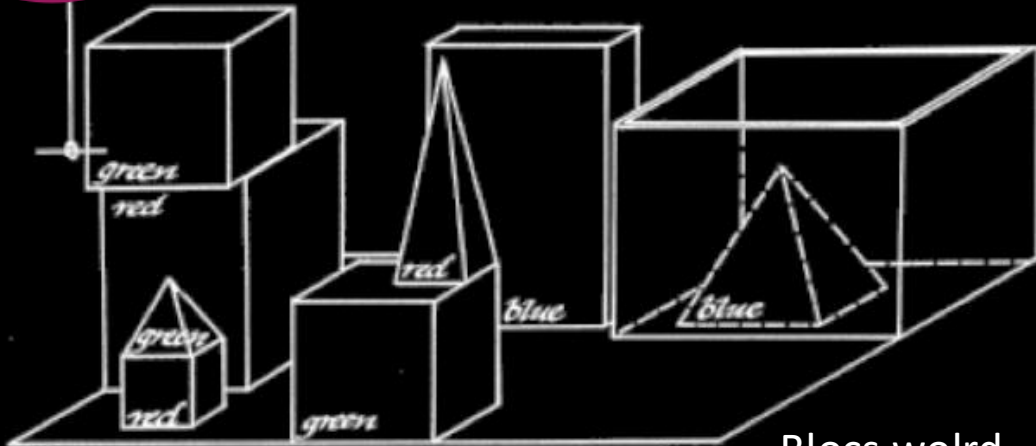
Robinson's complete algorithm for logical reasoning:  
the **resolution method**

The

60s

Years

1968



Blocs world

**SHRDLU**

Permitted to  
“discuss” with  
the computer,  
instruct it to  
manipulate  
its world



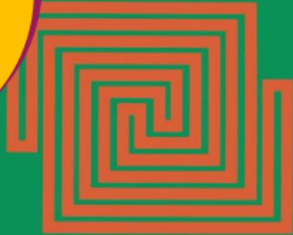
*The*

*60s*

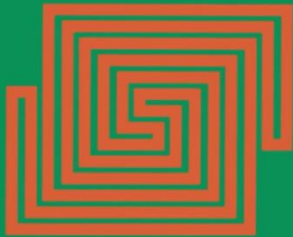
*Years*

*1969*

Expanded Edition



Perceptrons

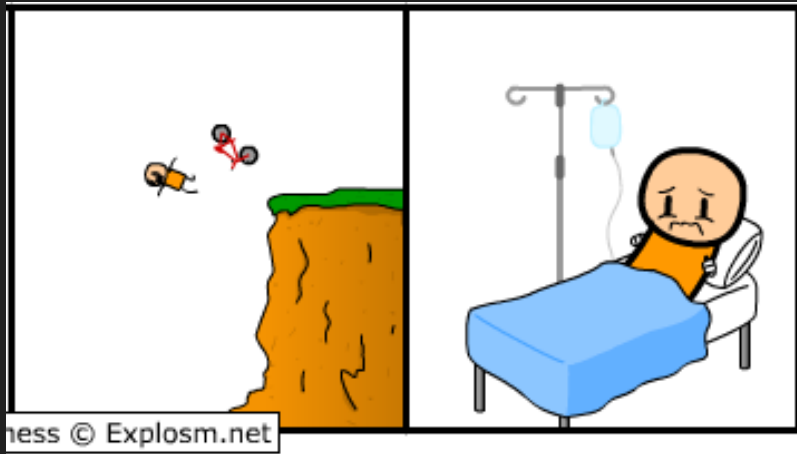


Marvin L. Minsky  
Seymour A. Papert

Marvin Minsky's and  
Seymour Papert's  
work show the limits  
of  
perceptron

# The 70<sub>s</sub> YEarS

## Artificial Intelligence Winter



Low computing power

State-space explosion

Knowledge/  
Rationality  
problem

# The 70<sub>s</sub> YEars

Early development of  
knowledge-based systems

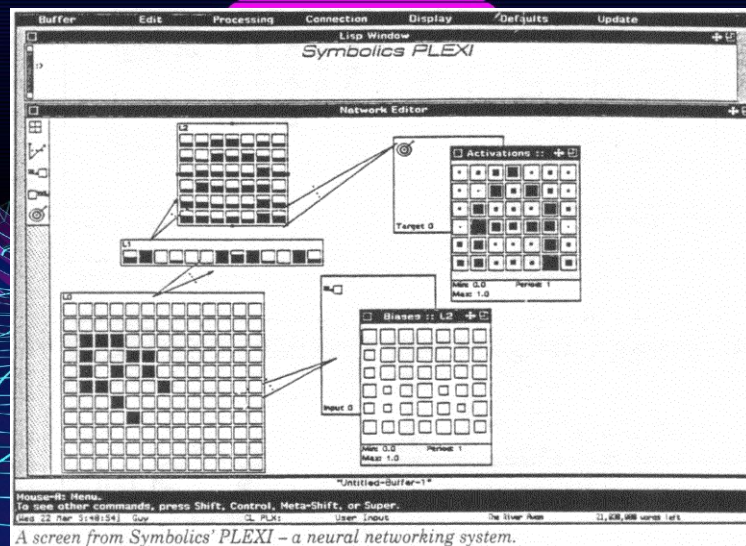
«One might say that to solve a hard problem, you  
have to almost know the answer»

# THE 80s

# YEARS

Backpropagation: overcome perceptron's limits

Hopfield networks: new model of neural networks



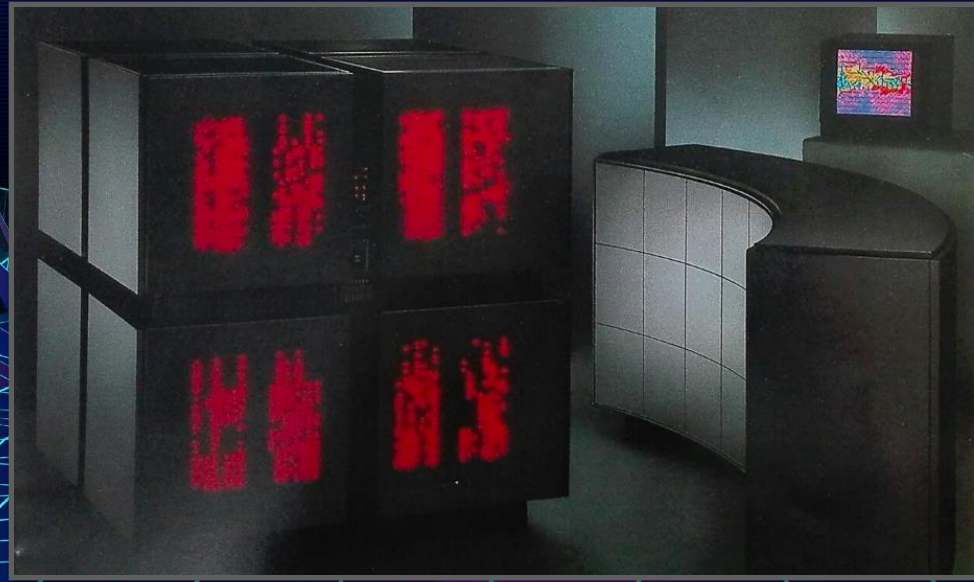
A screen from Symbolics' PLEXI - a neural networking system.

## Neural networks reborn



# THE 80s YEARS

Connection Machine: a different supercomputer



Precursor of modern GPU?

# THE 80s YEARS

Expert systems and knowledge engineering

Expert systems  
industry booms



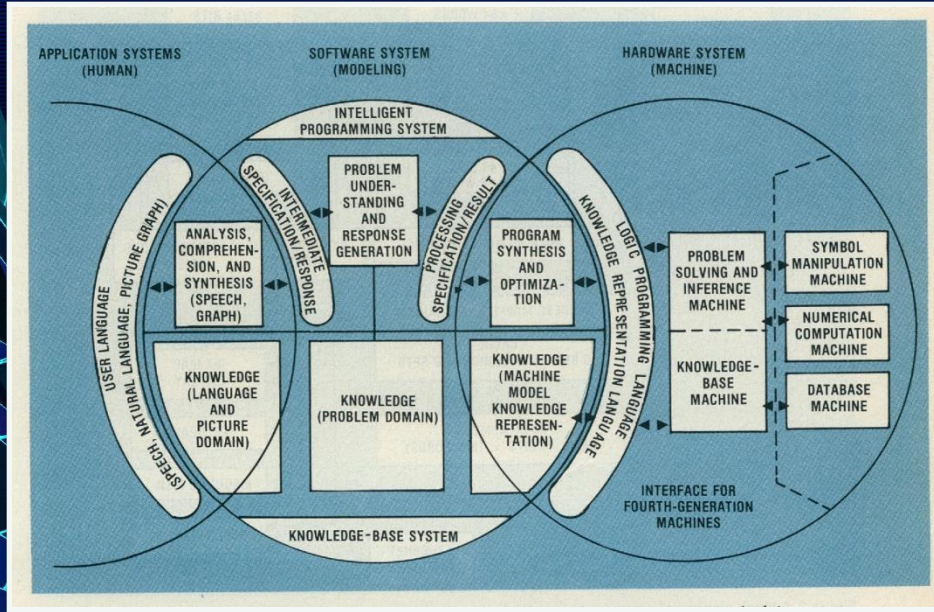
Workstation and systems fo AI



# THE 80s

# YEARS

The Japanese project of the fifth generation's computer



Developing of  
"Knowledge Information Processing systems"  
through logic programming

# The 90s years

## A second AI winter ☹️

Scalable Limits of the expert systems  
(what's the size of knowledge you "put inside"?)

Failure of the Japanese project of  
the fifth generation's computer

PC have similar performance as the  
dedicated workstation

# The 90s years

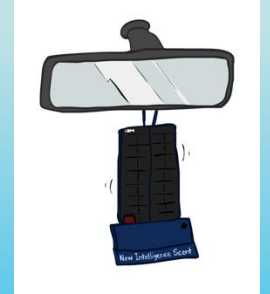
## "AI Spring"?

Let's use  
data

Resurgence of probability, focus on uncertainty

Machine learning rather than hand-coding

1996: Kasparov defeats Deep Blue at chess



"I could feel --- I could smell --- a new kind of intelligence across the table." ~Kasparov

# The 90s years



**1997:** Deep Blue defeats Kasparov at chess



“Deep Blue hasn't proven anything.”  
~Kasparov



# 2000—

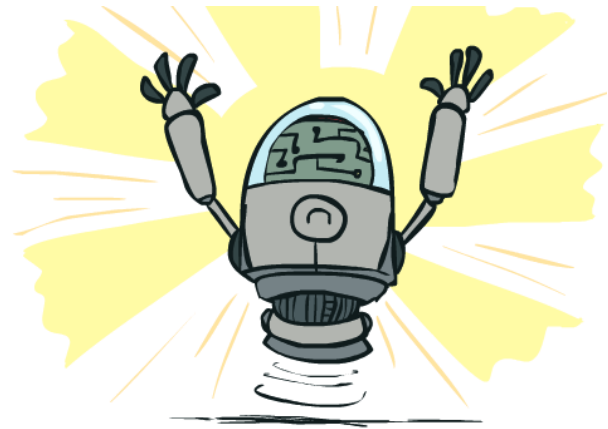
- Where are we now?
  - Big data, big computers, neural networks
  - Some re-unification of sub-fields
  - AI used in many industries
  - Chess engines running on ordinary laptops can defeat the world's best chess players
  - 2011: IBM's Watson defeats Ken Jennings and Brad Rutter at Jeopardy!
  - 2016: Google's AlphaGo beats Lee Sedol at Go



# What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of Jeopardy?
- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✓ Grab a particular cup and put it on a shelf?
- ✗ Unload any dishwasher in any home?
- ? Drive safely along the highway?
- ✗ Drive safely in Trieste?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Despar?
- ? Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✗ Unload a know dishwasher in collaboration with a person?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✗ Write an intentionally funny story?





# Course Topics

---

The so-called **Classical Artificial Intelligence**

**Intelligence from computation:** how we write algorithms or write the logic that will cause a programmer or a robot or an agent to do something that the human would or that we would consider intelligence



# Course Topics

---

- **Part 1: Problem-solving**
  - Search strategies and heuristic research (ch 3)
  - Search in complex Environments (ch 4)
  - Constraint Satisfaction Problems (ch 5)
  - Adversarial search and games (ch 6)
- **Part 2: Knowledge, Reasoning and planning**
  - Multiagent Decision Making (ch17)
  - Logical agent (ch 7)
  - First-order logic and inference (ch 8,9)
  - Knowledge representation (ch 10)
  - Automated planning (ch 11)



# Search problem

---

- Intelligent agents -> Chapter 2
- Search strategies -> Chapter 3
  - Search algorithms
  - Uninformed Search Strategies
  - Informed (Heuristic) Search Strategies
- Search in complex Environments -> Chapter 4
  - Local Search in discrete and continuous space
  - Search with non-deterministic action
  - Search with partially-observable environment

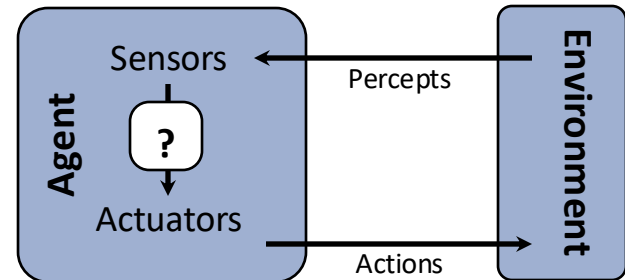
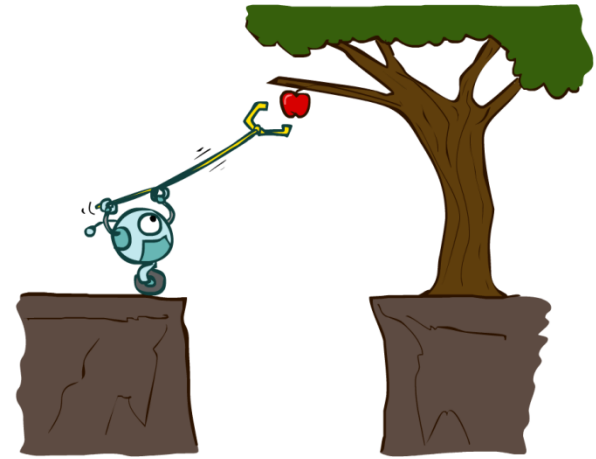
# Goal of the Course

---

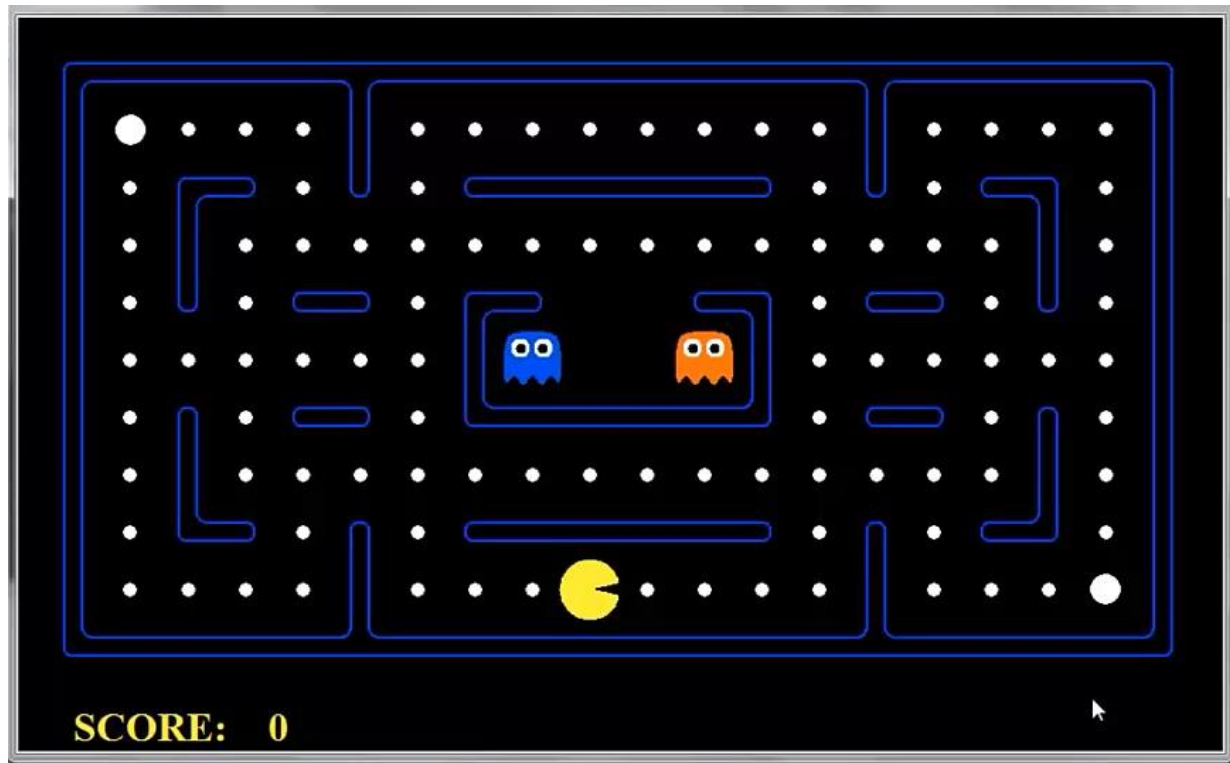
- It gives you an extra mathematical maturity
- It gives you a survey of fields that interact with AI

# AI as Designing Rational Agents

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that **maximize its (expected) utility**.
- Characteristics of the **sensors, actuators, and environment** dictate techniques for selecting rational actions
- **This course** is about:
  - Classical AI techniques for many problem types
  - Learning to choose and apply the technique appropriate for each problem



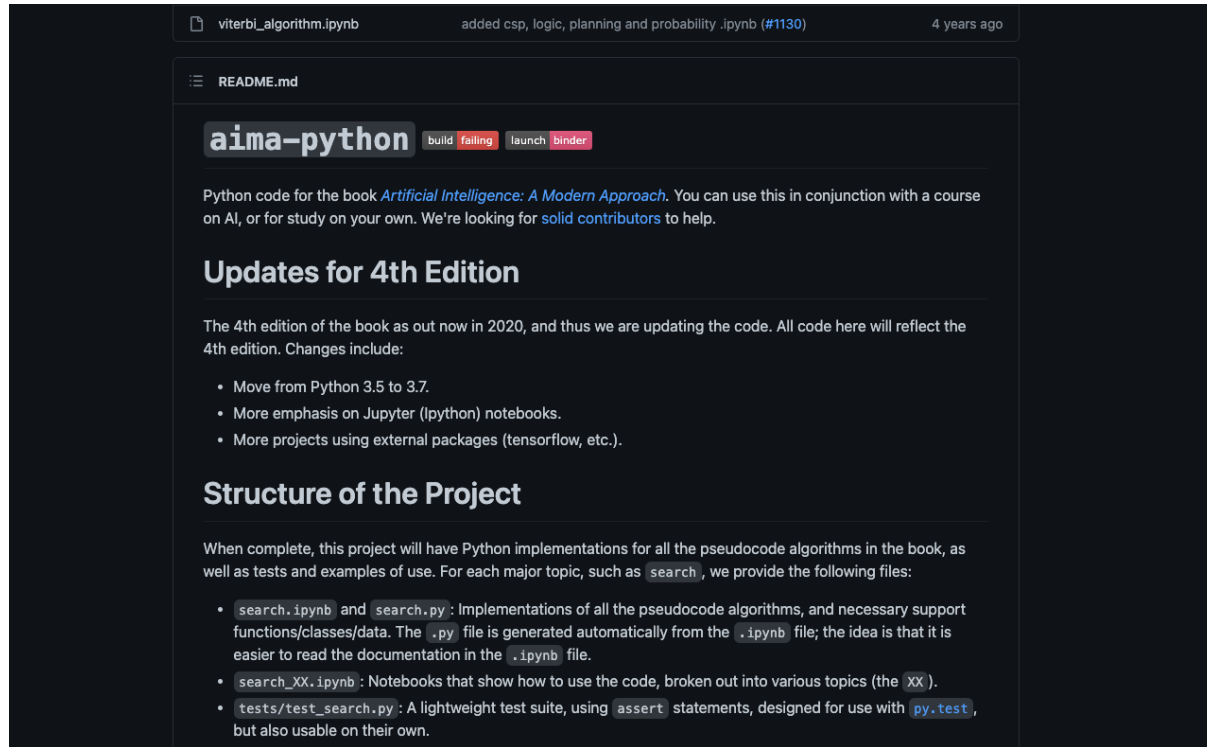
# Pac-Man as an Agent





# Aimacode

- <https://github.com/aimacode/aima-python>



The screenshot shows the GitHub repository page for `aimacode/aima-python`. At the top, it indicates the repository was updated 4 years ago. The main content area features the repository name `aima-python` with status indicators for `build` (falling) and `launch` (binder). Below this, there is a description of the project as Python code for the book *Artificial Intelligence: A Modern Approach*. A section titled "Updates for 4th Edition" lists changes such as moving to Python 3.7 and using external packages. Another section, "Structure of the Project", details the file organization, including `search.ipynb`, `search.py`, `search_XX.ipynb`, and `tests/test_search.py`.

viterbi\_algorithm.ipynb added csp, logic, planning and probability .ipynb (#1130) 4 years ago

README.md

## aima-python

build falling launch binder

Python code for the book *Artificial Intelligence: A Modern Approach*. You can use this in conjunction with a course on AI, or for study on your own. We're looking for [solid contributors](#) to help.

### Updates for 4th Edition

The 4th edition of the book as out now in 2020, and thus we are updating the code. All code here will reflect the 4th edition. Changes include:

- Move from Python 3.5 to 3.7.
- More emphasis on Jupyter (Ipython) notebooks.
- More projects using external packages (tensorflow, etc.).

### Structure of the Project

When complete, this project will have Python implementations for all the pseudocode algorithms in the book, as well as tests and examples of use. For each major topic, such as `search`, we provide the following files:

- `search.ipynb` and `search.py`: Implementations of all the pseudocode algorithms, and necessary support functions/classes/data. The `.py` file is generated automatically from the `.ipynb` file; the idea is that it is easier to read the documentation in the `.ipynb` file.
- `search_XX.ipynb`: Notebooks that show how to use the code, broken out into various topics (the `XX`).
- `tests/test_search.py`: A lightweight test suite, using `assert` statements, designed for use with `py.test`, but also usable on their own.