

CMSC 373

Artificial Intelligence

Fall 2023

Deepak Kumar
Bryn Mawr College

1

Administrivia

- **Instructor:** Deepak Kumar
- **Lectures:** Mon & Wed 10:10a to 11:30a
- **Labs:** Wednesdays 2:40p to 4:00p
- **Pre-requisites:** CMSC B206/B151 or H106 and CMSC B231 or permission of instructor. Junior/Senior standing expected.
- **Course web page:** <https://cs.brynmawr.edu/Courses/cs373/fall2023/>


Go to course web page now...

2

2

Texts & Software

- **Artificial Intelligence: A guide for Thinking Humans**
Melanie Mitchell
Farrar, Strauss, Giroux Books, 2019.
- **A Brief History of Artificial Intelligence**
Michael Wooldridge
Flatiron Books, 2020.
- We will use the programming language PROLOG (installed on all CS Lab machines and can be downloaded on your own devices) and Google's Colab (available through a web browser).



3

3

Goals

- History
- Foundations
- Applications
- Implications

4

4

Goals

- History
- Foundations
- Applications
- Implications

With a personal perspective.

1956

1983
Deepak gets
Interested in AI

1993
Deepak's PhD
Thesis in AI

2023

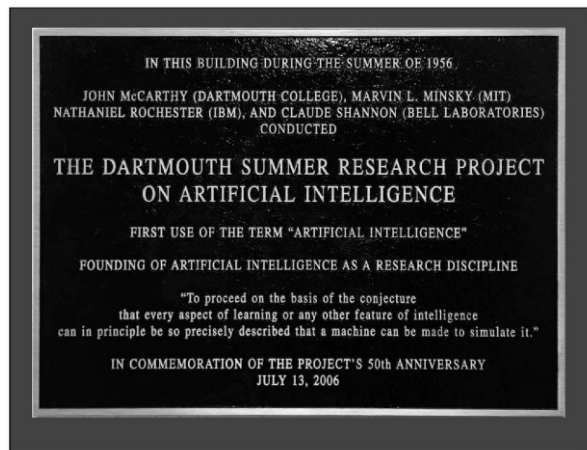
5

5

AI Beginnings – 1956 Dartmouth College



Image: From <https://home.dartmouth.edu/news/2022/11/dartmouth-hall-transformed>
Downloaded on July 5, 2023.




Plaque placed at Dartmouth Hall in 2006 commemorating 50 years of the Dartmouth Conference.

6

6

Summer Research Project Proposal (August 1955)

We propose a 2 month, 10 man study of artificial intelligence to be carried out during the summer of 1956.



The proposers of the 1956 Dartmouth Conference. From left to right: John McCarthy, Marvin Minsky (Photo: MIT Museum), Claude Shannon 0 and Nathaniel Rochester (Photo: Institute of Electrical and Electronics Engineers Inc)

7

7

Two Months and Ten Men at Dartmouth

“The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can be in principle so precisely described that a machine can be made to simulate it.”

-: McCarthy et al, 1955



John McCarthy



Marvin Minsky



Claude Shannon



Ray Solomonoff



Alan Newell



Herbert Simon



Arthur Samuel



Oliver Selfridge



Nathaniel Rochester



Trenchard More

8

8

AI Predictions: Early 1960s

- **John McCarthy** founded the Stanford AI Project with the “goal of building a fully intelligent machine in a decade.”
- **Herbert Simon** (and **Allen Newell**) went to Carnegie Mellon University.
“Machines will be capable, within twenty years, of doing any work that a man can do.”
- **Marvin Minsky** founded the AI Lab at MIT
“within a generation ... the problems of creating ‘artificial intelligence’ will be substantially solved.”

9

9

AI Predictions: Early 1960s

- **John McCarthy** founded the Stanford AI Project with the “goal of building a fully intelligent machine in a decade.”
- **Herbert Simon** (and **Allen Newell**) went to Carnegie Mellon University.
“Machines will be capable, within twenty years, of doing any work that a **man** can do.”
- **Marvin Minsky** founded the AI Lab at MIT
“within a generation ... the problems of creating ‘artificial intelligence’ will be substantially solved.”

None of the predictions have come to pass!

10

10

What is Artificial Intelligence?

- Your responses here...

11

11

What is AI?

- AI is about building fully intelligent machines?

Reverse engineer a human brain? Which human?

Find clever algorithms that can produce intelligent behavior?

- Minsky: **Intelligence** is an over packed suitcase word.
thinking, cognition, consciousness, emotion,...
verbal, spatial, logical, artistic, social, ...

12

12

What is AI?

- AI is about building fully intelligent machines?

Reverse engineer a human brain? Which human?

Find clever algorithms that can produce intelligent behavior?

- Minsky: **Intelligence** is an over packed suitcase word.
thinking, cognition, consciousness, emotion,...
verbal, spatial, logical, artistic, social, ...

Much of AI largely ignores all this!

13

13

What is AI?

- **Scientific Agenda**

Investigating the mechanisms of (natural) intelligence by trying to embed it in computers.

- **Practical Agenda**

Create computer programs that perform tasks as well or better than humans. Ignore *thinking* whatever that is.

14

14

What is AI?

- **Scientific Agenda**

Investigating the mechanisms of (natural) intelligence by trying to embed it in computers.

- **Practical Agenda**

Create computer programs that perform tasks as well or better than humans. Ignore *thinking* whatever that is.

Artificial Intelligence like Artificial Flight??

15

15

Why AI?

Or, *computers as electronic brains?*

- **Computation**

The process of taking **symbolic structures**, breaking them apart, comparing them, and reassembling them according to a precise recipe (or an algorithm)

Symbols: 7, 9, a, w, x, this, Deepak, ...

Operators: +, -, *, /, %, <, <=, >, >=, =, !=, ...

Symbolic structures: 793, -45.6, salary > \$100000, CMSC373, I am hungry, etc.

- In other words:

output symbols = algorithm(input symbols)

- **Key Insight:** A computer *does not understand* what the symbols stand for or why an output from a program/algorithm is correct.

Yet, computers do a ton of useful things for us by mechanically operating on symbols using algorithms.

16

16

Gottfried Leibniz (1646-1716)

- The rules of arithmetic allow us to deal with abstract numbers in terms of concrete symbols. The manipulation of those symbols mirrors the relations among the numbers being represented.
- The rules of logic allow us to deal with abstract ideas in terms of concrete symbols. The manipulation of those symbols mirrors the relations among the ideas being represented.

17

17

Gottfried Leibniz (1646-1716)

- The rules of arithmetic allow us to deal with abstract numbers in terms of concrete symbols. The manipulation of those symbols mirrors the relations among the numbers being represented.
- The rules of logic allow us to deal with abstract ideas in terms of concrete symbols. The manipulation of those symbols mirrors the relations among the ideas being represented.

AI's conjecture: *Thinking can be usefully understood as computation.*


18

18

Alan Turing, 1950

- **Computing Machinery and Intelligence**, *Mind*, Volume LIX, Number 236, October 1950: 433-460.

“I propose to consider the question, “Can machines think?””



Picture Credit: Overlooked No More: Alan Turing, Condemned Code Breaker and Computer Visionary, New York Times, June 5, 2019.
Link: <https://www.nytimes.com/2019/06/05/obituaries/alan-turing-overlooked.html>

19

19

The Imitation Game: Turing Test




20

20

Turing Test

- Operational test for intelligent behavior
- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes.
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
- Major subfields of AI: language processing, knowledge representation & reasoning, machine learning, vision, robotics



21

21

Turing Test

- Operational test for intelligent behavior
- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes.
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
- Major subfields of AI: language processing, knowledge representation & reasoning, machine learning, vision, robotics




Problem: Turing Test is not reproducible, constructive, or amenable to mathematical analysis

22

22

2015 Movie The Imitation Game



23

23

The Dichotomies of AI

1. The program ***actually understands*** the dialogue, in much the same sense that a person does.
2. The program does not understand but can ***simulate*** such understanding.

24

24

ELIZA by Joseph Weizenbaum, 1966

“An Illusion of Understanding”

```

Welcome to
EEEEEE LL      IIII  ZZZZZZ  AAAAA
EE      LL      II    ZZ    AA  AA
EEEEEE LL      II    ZZZ  AAAAAA
EE      LL      II    ZZ    AA  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 ?
YOU:

```

<https://en.wikipedia.org/wiki/ELIZA>

25

A Glimpse into ELIZA

- ELIZA has several **rules** of the form:

LHS → RHS
 [pattern] → [transform]

- Example:

(* YOU * ME) → (WHAT MAKES YOU THINK I 3 YOU)

where * is Kleene star (zero or more of some words)

and 3 is the matched token number in pattern

Input: YOU HATE ME.

Output: WHAT MAKES YOU THINK I HATE YOU

26

26

ELIZA's Rules

- Rules can be attached to (triggered by) a series of keywords

I: (I *) → (YOU SAY YOU 2)

everybody:

(* EVERYBODY *) -> (WHO IN PARTUCULAR ARE YOU THINKING OF?)

Try these inputs:

I know everybody laughed at me.

I hate everybody.

27

27

ELIZA's Rules

- Rules can be attached to (triggered by) a series of keywords

I: (I *) → (YOU SAY YOU 2)

everybody:

(* EVERYBODY *) -> (WHO IN PARTUCULAR ARE YOU THINKING OF?)

Try these inputs:

I KNOW EVERYBODY LAUGHED AT ME.

YOU SAY YOU KNOW EVERYBODY LAUGHED AT YOU

I hate everybody.

WHO IN PARTUCULAR ARE YOU THINKING OF?

28

28

ELIZA's Rules

- Default Rules

When nothing matches in the LHS:

(*) → (Please go on), (That's very interesting), (Go on)

- Plus, it maintains a “memory” queue.

```
[Memory MY
(* YOUR *) → (LETS DISUSS FURTHER WHY YOUR 3)
(* YOUR *) → (EARLIER YOU SAID YOUR 3)
(* O MY O) → (DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR 3)
]
```

Whenever **MY** is the highest ranked keyword, randomly select a transform on MEMORY list.

29

29

ELIZA's Rules

- It also maintains classes of words

FAMILY = MOTHER, FATHER, BROTHER, SISTER
NOUN = ...

- And performs basic transforms in output patterns

I → YOU
YOU → I

Example:

Input: MY SISTER MADE ME COME HERE.
Output: TELL ME MORE ABOUT YOUR FAMILY

For more details: Joseph Weizenbaum: ELIZA- A Computer Program for the Study of Natural Language Communication Between Man and Machine. Computational Linguistics, Volume 9, Number 1, February 1966.

30

30

The Dichotomies of AI

1. The program ***actually understands*** the dialogue, in much the same sense that a person does.
2. The program does not understand but can ***simulate*** such understanding.

31

31

The Dichotomies of AI: Strong AI vs Weak AI

1. The program ***actually understands*** the dialogue, in much the same sense that a person does. **Strong AI**
2. The program does not understand but can ***simulate*** such understanding. **Weak AI**

32

32

Winograd Schemas: A Better Test for Comprehension/Understanding (2012)

Consider the sentence:

The city councilmen refused the demonstrators a permit because they [feared/advocated] violence.

Version 1: *The city councilmen refused the demonstrators a permit because they advocated violence.*

Version 2: *The city councilmen refused the demonstrators a permit because they feared violence.*

See: <https://cs.nyu.edu/~davise/papers/WinogradSchemas/WS.html>

33

33

Winograd Schemas: A Better Test for Comprehension

Consider the sentence:

To really need to **understand** the text and have some **knowledge** about the scenario/context.

The city councilmen refused the demonstrators a permit because they [feared/advocated] violence.

Version 1: *The city councilmen refused the **demonstrators** a permit because **they** advocated violence.*

Version 2: *The **city councilmen** refused the demonstrators a permit because **they** feared violence.*

See: <https://cs.nyu.edu/~davise/papers/WinogradSchemas/WS.html>

34

34

Winograd Schemas: Another Example

Version 1: *The trophy doesn't fit into the brown suitcase because it is too small.*

Version 2: *The trophy doesn't fit into the brown suitcase because it is too large.*

What is too small/large?

See: <https://cs.nyu.edu/~davise/papers/WinogradSchemas/WS.html>

35

35

Winograd Schemas: Another Example

Version 1: *The trophy doesn't fit into the brown suitcase because **it** is too **small**.*

Version 2: *The trophy doesn't fit into the brown suitcase because **it** is too **large**.*

What is too small/large?

See: <https://cs.nyu.edu/~davise/papers/WinogradSchemas/WS.html>

36

36

The Dichotomies of AI

1. The program ***actually understands*** the dialogue, in much the same sense that a person does.

Strong AI

2. The program does not understand but can ***simulate*** such understanding.

Weak AI

Strong AI is largely irrelevant to contemporary AI research.

37

37

AGI: The Holy Grail of AI?

- AGI (Artificial General Intelligence)

Goal is to build machines that have general-purpose human-level intelligence.

[Recent claims about ChatGPT like programs!]

- Most (all?!) contemporary AI can be characterized as **Narrow AI**.

38

38

The Dichotomies of AI: Symbolic AI vs Subsymbolic AI

- Modeling the mind: the processes of conscious reasoning, problem solving, etc. (**semantic networks, logic-based systems**)

Uses symbols that stand for things that the system is reasoning about.

E.g. car54, room451, cleanRoom(room451)

- Modeling the brain: biologically inspired computation (**neural networks**)

39

39

The Dichotomies of AI: Symbolic AI vs Subsymbolic AI

- Modeling the mind: the processes of conscious reasoning, problem solving, etc. (**semantic networks, logic-based systems**)

Symbolic AI

Uses symbols that stand for things that the system is reasoning about.

E.g. car54, room451, cleanRoom(room451)

- Modeling the brain: biologically inspired computation (**neural networks**)

Subsymbolic AI

40

40

The Seasons of AI

- **1950s – 1966 First AI Summer: Irrational Exuberance**
Early successes in game playing, theorem proving, problem solving
- **1967 – 1977 First AI Winter**
No useful deliverables led to loss of research funding and cancellation of AI programs. In UK *The Lighthill Report* (toy AI systems do not scale due to combinatorial explosion).
- **1978 – 1987 Second AI Summer/Spring**
Rise of knowledge-based systems, success of Expert Systems. Boom times.
- **1988 – 1993 Second AI Winter**
Failure of AI Hardware companies (Symbolics, LMI, Lisp Machines) and AI Companies (Teknowledge, Inference Corp. etc.) Commercial deployments of Expert Systems were discontinued.
- **1993 – 2011 Third AI Summer (Mostly academic advances)**
Statistical approaches and extensions to logic (Bayesian Nets), Non-Monotonic Reasoning (in Logic), Fuzzy Logic, advances in Machine Learning (Decision Trees, Random Forests, Neural Nets), Cognitive Models, Logic Programming, Case-Based Reasoning, Genetic Algorithms, Agent-based approaches, etc.
- **2011 – Now Third AI Spring**
Rise of Deep Learning, Neuro-symbolic AI, ChatGPT and other chatbots, generative AI.

41

41

Reading for this week

- Chapter 1 from Wooldridge and pages 17-26 from Mitchell.
- Pick one (or all!) of the several original articles mentioned in this lecture (listed in the References Section) for a discussion next week.

42

42

References

- J. McCarthy, M. Minsky, N. Rochester, C. E. Shannon: A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence (1955). Reprinted in *AI Magazine* 27, No. 4 (2006): 12-14
(<https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/1904>)
- A. M. Turing: Computing Machinery and Intelligence. In *Mind*, Vol LIX, No. 236. October 1950. Available at: <https://academic.oup.com/mind/article/LIX/236/433/986238>
- J. Weizenbaum: A Computer Program for the study of Natural Language Communication Between Man and Machine. *Communications of the ACM*, Vol 9, Number 1. January 1966. Available at: <https://dl.acm.org/doi/pdf/10.1145/365153.365168>
- E. David, L. Morgenstern, C. Ortiz: The Winograd Schema Challenge. Webpage available at: <https://cs.nyu.edu/~davise/papers/WinogradSchemas/WS.html>
- M. Mitchell: *Artificial Intelligence-A Guide for Thinking Humans*. Farrar, Straus, and Giroux Books, 2019.
- M. Wooldridge: *A Brief History of Artificial Intelligence*. Flatiron Books, 2020.