

# Artificial Intelligence

## CS 441/541

Instructor: Melanie Mitchell

1

- **Two-fold nature of AI:**

- **Philosophical:**

- “Can machines think, in principle?”
- “Is this particular machine thinking [conscious, creative, ...]”
- Will machine thought be different from human thought?

- **Practical:**

- Collection of techniques to automatically solve particular problems that require “intelligence” (whatever that is)

2

- What abilities would an intelligent program have?
- Are intelligent programs really intelligent? How can we tell?

3

*“I believe that in about fifty years’ time it will be possible to programme computers...to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning.*

— Alan Turing, *Computing machinery and intelligence*, 1950.

4

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

— Marvin Minsky to **Life** magazine, 1970

5

*We can expect computers to pass the Turing test, indicating intelligence indistinguishable from biological humans, by the end of the 2020s.*

— Ray Kurzweil, *The Singularity is Near*, 2005

“Long bets” prediction [website](#)

6

### Milestones in AI History (from AAAI website)

1950: Alan Turing publishes, "Computing Machinery and Intelligence."  
1956: John McCarthy coins the term, "Artificial Intelligence" at a Dartmouth computer conference.  
1956: Demonstration of the first running AI program at Carnegie Mellon University.  
1958: John McCarthy invents the Lisp language, an AI programming language, at Massachusetts Institute of Technology (MIT).  
1964: Danny Bobrow shows that computers can understand natural language enough to solve algebra word programs (MIT).

7

1965: Joseph Weizenbaum builds ELIZA, an interactive program that carries on a dialogue in English on any topic (MIT).  
1969: Shakey, a robot, combines locomotion, perception and problem solving (Stanford Research Institute).  
1979: The first computer-controlled autonomous vehicle, the Stanford Cart, is built.  
1983: Danny Hillis co-founds Thinking Machines, the first company to produce massively parallel computers.  
1985: The drawing program, Aaron, created by Harold Cohen, is demonstrated at AI conference.

8

1990s: Major advances in all areas of AI. Significant demonstrations in machine learning, intelligent tutoring, case-based reasoning, multi-agent planning, scheduling, uncertain reasoning, data mining, natural language understanding and translation, vision, virtual reality and games.  
1997: IBM computer Deep Blue beats world champion Garry Kasparov in chess match.  
Late 1990s: Web crawlers and other AI-based information-extraction programs become Web essentials.  
**2000s:**  
Interactive robot pets become commercially available.  
MIT displays Kismet, a robot with a face that expresses emotions.  
Carnegie Mellon robot Nomad explores remote regions of Antarctica and locates meteorites.  
Mars rovers Spirit and Opportunity autonomously explore Mars  
Stanford University robot car Stanley wins DARPA Grand Challenge.  
Robotic vacuum cleaners, floor scrubbers, and lawn mowers become mainstream

9

*"In from three to eight years, we'll have a machine with the general intelligence of an average human being."*

— Marvin Minsky to **Life** magazine, 1970

*"Easy things are hard."*

— Marvin Minsky, *Society of Mind*, 1988

10

### Famous AI Debates

"Strong" vs. "weak" AI

"Symbolic" vs. "sub-symbolic" AI

"Neats" vs. "scruffies"

11

### AI Demos



#### Tests of "true" intelligence

Turing Test (e.g., [Loebner Prize](#))

Captchas

[Words](#)

[Bongo](#)

[ESP-PIX](#)

[Breaking captchas](#)

- [Babelfish translation](#)
- [Sign-language translator glove](#)
- [Word-sense disambiguation](#)

#### Natural language processing

- [Eliza](#)
- [Jabberwock](#)
- [Start.NL question answering](#)
- [Dragon Naturally Speaking](#)

#### Arts

- [EMI](#)
- [Aaron](#)
- [Picbreeder](#)
- [Sims evolved art](#)
- [Sims virtual creatures](#)

12

## AI Demos

### Search and game-playing

- [Checkers](#)
- [Chess \(Fritz\)](#)

### Knowledge representation and reasoning

- [Open Mind](#)

### Vision

- [Face detection](#)
- [Content-based image retrieval](#) 1
- [Content-based image retrieval](#) 2

### Robotics

- [Lipson lab videos](#)
- [Kismet](#)
- [ASIMO](#)
- [Cockroach robots](#)
- [Roomba](#)
- [Stanley](#)

### Reasoning under uncertainty

- Bayesian reasoning about probabilities

13

## THE VALUE FUNCTION

(adapted from Kahneman & Tversky, 1979)

- **1. In addition to whatever you own you have been given \$1000. You are now asked to choose between:**
  - A: a 50 per cent chance of \$1000
  - B: \$500 for certain
- **2. In addition to whatever you own you have been given \$2000. You are now asked to choose between:**
  - C: a 50 per cent chance of losing \$1000
  - D: losing \$500 for certain

## THE VALUE FUNCTION

(adapted from Kahneman & Tversky, 1979)

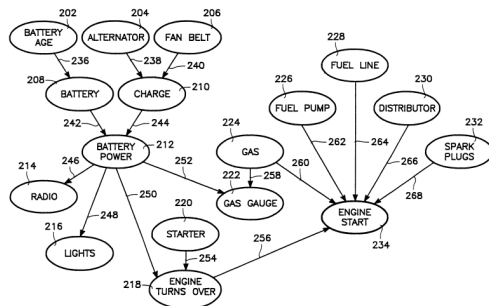
- **1. In addition to whatever you own you have been given \$1000. You are now asked to choose between:  $N=70$** 
  - A: a 50 per cent chance of \$1000  
[16 per cent]
  - B: \$500 for certain  
[84 per cent]
- **2. In addition to whatever you own you have been given \$2000. You are now asked to choose between:  $N=68$** 
  - C: a 50 per cent chance of losing \$1000  
[69 per cent]
  - D: losing \$500 for certain  
[31 per cent]

## The “Monte Hall problem”

Monte Hall shows you three doors. He tells you that behind one of them is a brand new car, and that there are goats behind the other two. He asks you to choose one of them. You choose door #1. He opens door #3 and reveals a goat. He then asks if you want to change your choice to door #1. Should you change?

16

## A Bayesian network for diagnosing automobile problems:



17

## Class overview

- **Class web page:**  
<http://www.cs.pdx.edu/~mm/ArtificialIntelligenceFall2007>
  - For schedule, slides, assignments, etc.
- **Office hours:**  
M, W 2-3pm, or by appointment
- **TA:** TBA

18

### • Topics we'll cover

- What would constitute an intelligent program?
- Problem-solving as search
- Heuristic search
- Game playing: adversarial search
- Knowledge representation, reasoning, and planning
- Natural language processing
- Vision
- Learning and information retrieval
- Reasoning under uncertainty, Bayesian networks
- Analogy

19

### • Assignments:

- Reading, reading questions
- Homework exercises (written)
- Paper presentations (undergrads)
  - Read an assigned paper
  - Present it in class (10-15 minutes)
- AI system presentations (required for grad students, optional for undergrads)
  - Research a particular AI system
  - Present it in class (20-30 minutes)
- Team AI programming project

20

### Paper assignments and presentations

- Use Powerpoint (or equivalent) slides
- Slides to me before class or link to web site, or put on flash drive, or bring your own computer
- Undergrads select preferred topic area for paper presentation by next Monday. I will schedule presentations and assign papers.
- Grads select AI system to present by next Monday. I will schedule presentations.
- Grading based on clarity and correctness of presentation

21

### Projects:

- **Goal:** Implement system that uses at least two different AI techniques, or compare performance of at least two different AI techniques.
- Work in teams of up to 4 people
- Final paper on project: 15-25 pages
- Next class: we'll do some matchmaking
- **By Oct 3:** Teams should be formed, each with rough idea of project area
- **Week of Oct. 3:** I'll meet with each team individually to discuss project plans

22

### Sample projects

- Implement a game playing program that uses adversarial search and a human-written evaluation function to determine next move; compare with using a neural network learn to perform the evaluation of a game configuration.
- Implement a simple face-detection program, combining two different methods.
- Implement a hand-written digit recognition program, combining two different methods.
- Implement a simple NLP technique for determining the topic of a text, and compare with latent semantic analysis.
- Implement a simulation of a few of the Braitenberg vehicles, and experiment with their behavior in different environments.
- Implement a chat-box program using fixed rules, and one using a Bayesian network to determine what to say next. Compare the behavior of the two.

23

### Grading:

Homework (reading questions and exercises): 50%

In-class presentation of paper or AI system: 20%

Final project, presentation, and paper: 30%

24

## Today's assignment

Complete by next class (Sept. 26)

- **Read:**

*Alan Turing: Computing Machinery and Intelligence*

(Linked from class web page)

- **Answer reading questions**

- Hand in computer formatted answers
- Check spelling!

- Why we're starting by reading Turing
- Turing history (Hilbert's questions, + Godel's theorem, halting problem)

25

[Vehicles simulation](#)

26

## Vehicles

Overview, why we are reading it

27

## Vehicles

- Overview, why we are reading it

28

### Vehicle 1: Alive

This and following slides adapted from  
"Notes on Braitenberg's Vehicles"


([http://instruct.westvalley.edu/lafave/Vehicles\\_online.html](http://instruct.westvalley.edu/lafave/Vehicles_online.html))



- Components: Sensor and motor.
- Principle: The more there is of the quality (e.g., heat) to which the sensor is tuned, the faster the motor goes.
- Description: alive, restless, doesn't like "heat" (will slow down in cold regions, speed up in hot regions)

29

### Vehicle 2a: Cowardly




Vehicle 2a:

- Components: 2 sensors, 2 motors, each sensor connected to the motor on the same side ("uncrossed")
- Principle: The more there is of the quality to which the sensor is tuned, the faster the motors go ("excitatory").
- Description: Will run away from source to which the sensor is tuned; occasionally "attacks" it

30

### Vehicle 2b: *Aggressive*



Vehicle 2a:

- Components: 2 sensors, 2 motors, each sensor connected to the motor on the opposite side ("crossed")
- Principle: The more there is of the quality to which the sensor is tuned, the faster the motors go ("excitatory").
- Description: dislikes source to which the sensor is tuned; "attacks" it

31

### Vehicle 3: *Loving*

- Components: 2 sensors, 2 motors, each sensor connected to the motor on the same side ("uncrossed")
- Principle: The more there is of the quality to which the sensor is tuned, the *slower* the motors go ("inhibitory").
- Description: loves the source, wants to be near it, comes to rest facing it

32

[Vehicles Simulation](#)

33