

## 22. Artificial Intelligence

### 22. Artificial Intelligence

- Life
- Games
- Cars
- Turing test
- Strong AI

#### Conway's game of life

##### A cellular automaton

- Cells live and die in an infinite square grid.
- Time proceeds in discrete steps.

##### Survival.

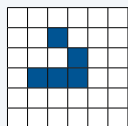
 Each cell lives or dies depending on its 8 neighbors:

- Too few neighbors alive? (0 or 1) Die of loneliness.
- Number of living neighbors just right (2 or 3) Survive to next generation.
- Too many neighbors alive? (more than 3) Die of overcrowding.

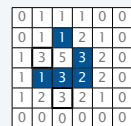
##### Birth.

 Cell born when it has exactly 3 living neighbors.

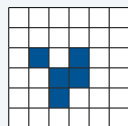

John Horton Conway



time t



living neighbors



time t+1

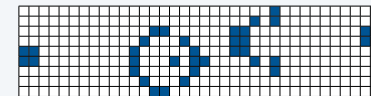
#### Conway's Game of Life

**Lesson.** Simple rules can lead to extremely complicated behavior

Example 1. Glider



Example 2. Glider gun (generates gliders)



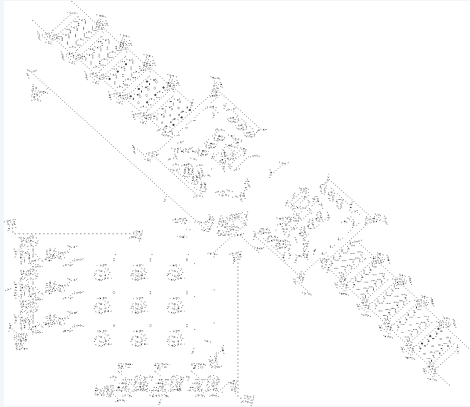
Example 3. Glider gun breeder (generates glider guns)



**Note.** YOU can write a program for the game of life (too easy for this course).

## A complex initial configuration for the game of life

Q. What happens with this starting configuration?



A. Anything we can compute!  
(It is a UTM).

5

## Questions

Is REAL life described by discrete rules, or not?

Can a UTM be built from living components?

Can a machine do anything a human can do?

Can human intelligence be simulated by a machine?



**Artificial Intelligence.** The science and engineering of making intelligent machines.

6

## Two conceptual approaches to artificial intelligence

**Strong AI.** Can a machine *be* intelligent?

*"I propose to consider the question, 'Can machines think?'"*

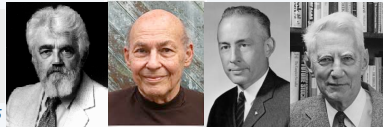
— Alan Turing, 1950



**Weak AI.** Can a machine *appear* to be intelligent?

*"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."*

— John McCarthy, Marvin Minsky, Nathaniel Rochester, Claude Shannon, 1955



7

COMPUTER SCIENCE  
SEDGEWICK / WAYNE

## 22. Artificial Intelligence

- Life
- Games
- Cars
- Turing test
- Strong AI

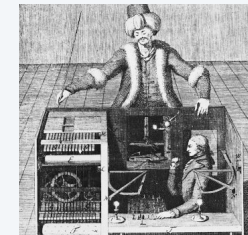
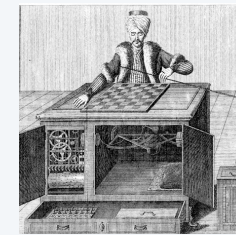
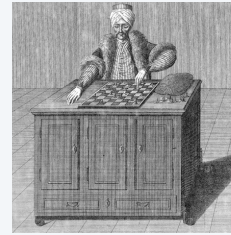
## 22. Artificial Intelligence

- Life
- **Games**
- Cars
- Turing test
- Strong AI

### The mechanical Turk (1770-1854)

#### A chess-playing automaton

- Developed in the court of Maria Theresa.
- Toured the world playing human opponents for decades.
- Defeated Benjamin Franklin and Napoleon.
- Exposed as a hoax in 1820.



Q. Can a machine play a game (and defeat a human)?

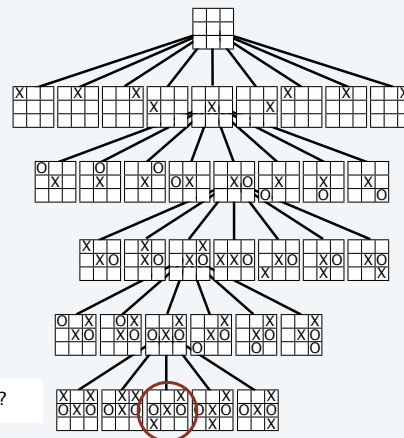
### Tic-Tac-Toe

#### A two-person game

- Number of possible games: 255,168.
- Machine can check all possibilities to make sure no move makes it possible for opponent to force a win.
- *Min-max algorithm* achieves the same goal without needing to check all possibilities.

Game is *solved*.

Neither player can force a win.



Q. Can a machine play a game *as well as* a human?

A. Yes. (Neither should lose.)

### Checkers

#### A two-person game

- Number of possible games: about  $10^{31}$ .
- *Chinook* computer program by a team led by Jon Schaeffer awarded world championship in 1994 after 6 draws with Marion Tinsley (who withdrew).
- 2007: Schaeffer and team *solved* checkers (proved that neither player can force a win (!))



Q. Can a machine play a game (and defeat a human)?

A. Yes.

## Chess

### A two-person game

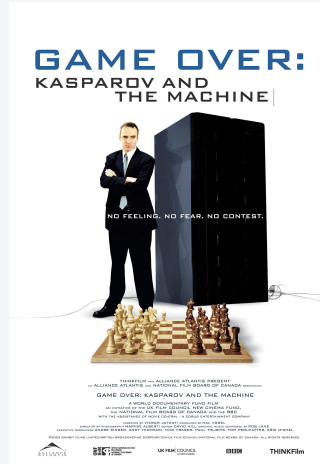
- Number of possible games: about  $10^{123}$ .
- *Deep Blue* computer program by IBM researchers beat Garry Kasporov for the world championship in 1997.

Open question: can chess be solved?



Q. Can a machine play a game (and defeat *any* human)?

A. Yes.



13

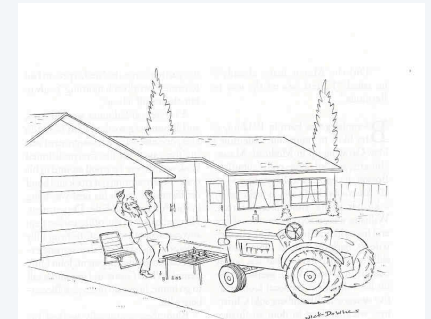
## Can a machine play a game?



Q. Can a machine play a game?

A. Sure.

Q. Does that constitute *intelligence*?



"Once again, man beats machine!"

14

COMPUTER SCIENCE  
SEDGWICK / WAYNE

## 22. Artificial Intelligence

- Life
- **Games**
- Cars
- Turing test
- Strong AI

COMPUTER SCIENCE  
SEDGWICK / WAYNE

## 22. Artificial Intelligence

- Life
- Games
- **Cars**
- Turing test
- Strong AI



## Can a computer drive a car?

### 2004 DARPA Grand Challenge

- Navigate an autonomous vehicle.
- 142 mile course in Mohave Desert.
- Military speed.
- 15 entrants.
- \$1 million prize.

ENSCO



Flipped in the start area

Cliff (Virginia Tech)



Brakes locked in the start area

CIMAR (Florida State)



Wrapped up in wire (0.4 miles)

Rascal (Auburn)



Stuck in an embankment (6.7 miles)

Caltech



Stuck in a fence (1.3 miles)

Sandstorm (CMU)



Stuck on an obstacle (7.4 miles)  
Front wheels caught fire  
Last surviving vehicle

Results. 0 finishers.

17

## Can a computer drive a car?

### 2005 DARPA Grand Challenge

- Navigate an autonomous vehicle.
- 132 mile course in Mohave Desert.
- Military speed.
- 15 entrants.
- \$2 million prize.

Stanley (Stanford)



Finished in under 7 hours

Sandstorm (CMU)



Second place, 10 minutes behind.

Results. FIVE teams finish. Stanford team collects the prize (!)

18

## Can a computer drive a car?

### 2007 DARPA Urban Challenge

- Navigate an autonomous vehicle.
- 60 mile course in *mock urban environment*.
- Obey traffic laws.



Results. Six teams finish. (!!)

19

## Progress report

### Google driverless car (2012)

- 300,000 miles on public roads in CA.
- Typically a dozen cars on the road.
- No accidents.



Q. Does it take *intelligence* to drive a car?

20

## 22. Artificial Intelligence

- Life
- Games
- **Cars**
- Turing test
- Strong AI

## 22. Artificial Intelligence

- Life
- Games
- Cars
- **Turing test**
- Strong AI

### Turing test

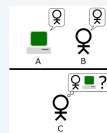
- Q. Can a machine *be* intelligent?
- Q. Can a machine *appear to be* intelligent?
- Q. If so, **How could we tell?**

*"I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms "machine" and "think." Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words."* — Alan Turing, 1950



### Turing test

- Computer (program) and person communicate with a judge.
- Each claims to be a person.
- If judge cannot tell which is the program, the program passes the test.



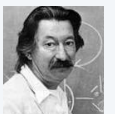
### A response to the Turing test

#### ELIZA [Weizenbaum 1966]

- Program imitates psychiatrist.
- Generates scripted responses to keywords.
- Just 200 lines of code!

- Q. Does Eliza appear intelligent?
- A. You be the judge.

Men are all alike.  
 IN WHAT WAY  
 They're always bugging us about something.  
 CAN YOU THINK OF A SPECIFIC EXAMPLE  
 Well, my boyfriend made me come here.  
 YOUR BOYFRIEND MADE YOU COME HERE  
 He says I'm depressed much of the time.  
 I AM SORRY TO HEAR YOU ARE DEPRESSED  
 It's true. I am unhappy.  
 DO YOU THINK COMING HERE WILL HELP YOU  
 I need some help, that much seems certain.  
 WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP  
 Maybe I could learn to get along with my mother.  
 TELL ME MORE ABOUT YOUR FAMILY  
 My mother takes care of me,  
 WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU  
 My Father.  
 YOUR FATHER  
 You are like my father in some ways.  
 ...



J. Weizenbaum  
1923 – 2008

TURING TEST EXTRA CREDIT:  
CONVINCE THE EXAMINER  
THAT HE'S A COMPUTER.

YOU KNOW, YOU MAKE  
SOME REALLY GOOD POINTS.  
I'M ... NOT EVEN SURE  
WHO I AM ANYMORE.

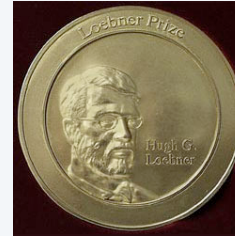


25

## Annual Turing tests

### Loebner Prize

- \$100,000 to first program to pass restricted Turing Test.
- Annual competition held since 1995.
- Small awards given to best attempts.
- Prize still unclaimed.



26

## Sample Loebner contest dialogs

"Hello, Judge. My name is Ella.  
Please tell me your first name, or whatever I should call you."  
"Call me Lisa. Am I male or female?"  
"Okay, I will call you 'Lisa. Am I Male Or Female?'  
I am here recently from Tainjin, China, so please be patient with me."  
"I'm not a patient because I am not ill. Are you feeling well today?"  
"What do you think it means that you are not ill?"  
"Have you ever been ill?"  
" 'Have you ever been ill?' you say?  
That's it, Lisa Am I Male Or Female! I will schedule a whipping for you."

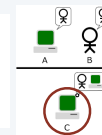
"What do you think of Reading?"  
"I'm somewhat crazy about reading. I like to read mystery novels."  
"I meant the town."  
"Jones lives in the city and is a writer of children's books."  
"Does Jones live in Reading, then?"  
"Tomorrow is another day. We can only live today well."  
"Do you live, then?"  
"I live then a lot. Someday we will live on another planet."  
"I think you are already there."

27

## Digression: "Reverse" Turing test

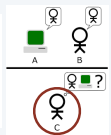
### Reverse Turing test

- Judge is a *program*.
- Application: Web commerce.



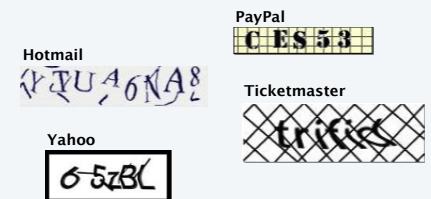
### Turing test

- Judge is a person.
- Application: Validate AI.



### CAPTCHA

- Completely Automated Public Turing test to tell Computers and Humans Apart.
- Approach: Display a word that is easily read by a human but not by a program.
- Invented in 1997 by a group at AltaVista, by reverse-engineering a scanner manual.



28

## Progress report: Watson

Watson: a computer system specifically built to play *Jeopardy!*

- Access to 200 million pages of content (4 terabytes).
- 2011: Watson beat two former *Jeopardy!* champions.

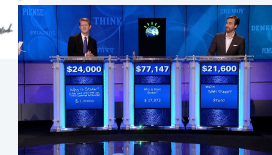
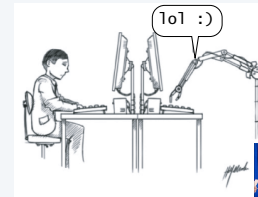


Q. Is a machine that can answer questions intelligent?

Q. Does a machine that can answer questions *appear to be* intelligent?

29

## Can a machine answer questions?



Q. Can a machine answer questions?

A. Sure.

Q. Does that constitute *intelligence*?

*“Watson got the promotion to CIO over me. How am I supposed to compete with a supercomputer who won on Jeopardy?”*

30

COMPUTER SCIENCE  
SEDGWICK / WAYNE

## 22. Artificial Intelligence

- Life
- Games
- Cars
- **Turing test**
- Strong AI

COMPUTER SCIENCE  
SEDGWICK / WAYNE

## 22. Artificial Intelligence

- Life
- Games
- Cars
- Turing test
- **Strong AI**

## "Weak" AI

Weak AI. Can a computer *appear* to be intelligent?

Q. Can a computer play a game (and defeat a human)?

A. Sure.

Q. Does that constitute *intelligence*?



Q. Can a computer drive a car?

A. Sure.

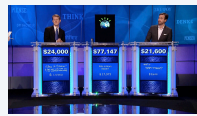
Q. Does that constitute *intelligence*?



Q. Can a machine answer questions?

A. Sure.

Q. Does that constitute *intelligence*?



Strong AI. Can a machine *be* intelligent?

33

## Chinese Room Experiment (Searle 1980)

Imagine that:

- You don't understand Chinese.
- You're alone in a room that has paper slots labeled "input" and "output".
- You have a big book of Chinese writing.
- You have English instructions (no translations) that tell you what to write on a piece of paper in response to any possible inputs.

And then:

- Chinese speakers outside the room pass in pieces of paper with Chinese writing. They know these are questions (but you don't).
- You consult your manual of instructions, figure out the proper Chinese response, copy it, and pass it out to them.



If you see this shape, "什麼"	then produce this shape, "為天"
followed by this shape, "帶來"	followed by this shape, "下式".
followed by this shape, "快樂"	

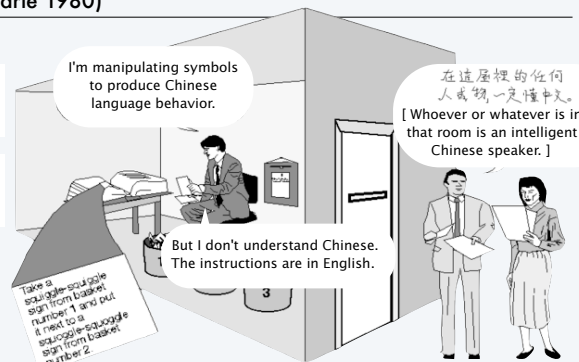
什麼帶來快樂  
為天下式

34

## Chinese room experiment (Searle 1980)

Q. The folks outside think you understand Chinese. Do you?

Q. If a computer did the same, would it understand Chinese?



[http://www.mind.illstu.edu/curriculum/searle\\_chinese\\_room/searle\\_chinese\\_room.php](http://www.mind.illstu.edu/curriculum/searle_chinese_room/searle_chinese_room.php)

Searle. Chinese Room thought experiment **absolutely refutes** the idea of strong AI.

35

## "Systems" refutation of Searle's assertion that strong AI is impossible

Man/machine in the room is just one part of a system

- Huge memory (piece of paper for every possible symbol).
- Unlimited I/O.

Could implement a UTM!



Akin to the difference between a DFA and a Turing machine.

Turing	DFA	cannot recognize palindromes
	TM	can recognize anything that any computer can recognize
Searle	human/machine	not acting intelligently
	system	intelligent (?)

The *machine* may not understand Chinese, but the *system* does!

36

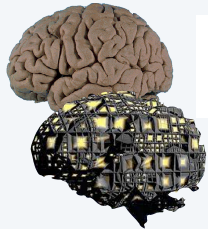
## "Ship of Theseus" refutation

### Ship of Theseus paradox

- Theseus sails the world in a wooden ship.
- During the trip, every single component of the ship is replaced, one by one.
- Is it the same ship when it returns?



Plutarch  
ca. 46–ca. 120



### Refutation of Searle's assertion that strong AI is impossible

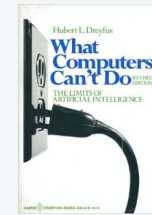
- Suppose that we replace each neuron in a human brain, one by one.
- At what point does the brain become something that is not intelligent?

Note: ~100 billion neurons in the brain.  
500+ billion bytes in your PC.

37

## Longstanding debate: Is strong AI ultimately possible?

"Just as the Wright brothers at Kitty Hawk in 1903 were on the right track to the 747, so too is AI, with its attempts to formalize commonsense understanding, on its way to fully intelligent machines."  
– Patrick Winston



"Believing that writing these types of programs will bring us closer to real AI is like believing that someone climbing a tree is making progress toward reaching the moon."  
– Hubert Dreyfus

"The brain happens to be a meat machine."  
– Marvin Minsky



"Either AI is possible ... or we're not."  
– Herb Simon



"The question of whether a computer can think is no more interesting than the question of whether a submarine can swim."  
– Edsger Dijkstra



38

## 22. Artificial Intelligence

- Life
- Games
- Cars
- Turing test
- Strong AI

## A challenge

Write a program that can pass this course.

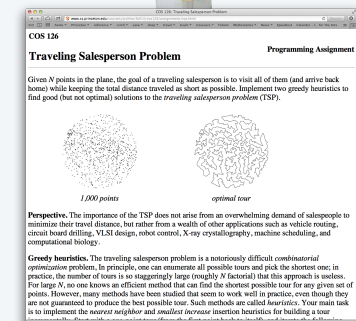
- Exams
- Programming assignments
- Programming exams.

Q. Would such a program be intelligent?

Next challenge.

Write a program that can pass *any* college course.

Q. Would such a program be intelligent?



CS 126 Programming Assignment

**Traveling Salesperson Problem**

Given  $N$  points in the plane, the goal of a traveling salesperson is to visit all of them (and arrive back home) while keeping the total distance traveled as short as possible. Implement two greedy heuristics to find good (but not optimal) solutions to the traveling salesperson problem (TSP).

*1,000 points* *optimal tour*

**Perspective.** The importance of the TSP does not arise from an overwhelming demand of salespeople to minimize their travel distance, but rather from a wealth of other applications such as vehicle routing, circuit board drilling, VLSI design, robot control, X-ray crystallography, machine scheduling, and computational biology.

**Greedy heuristics.** The traveling salesperson problem is a notoriously difficult combinatorial optimization problem. In principle, one can enumerate all possible tours and pick the shortest one; in practice, the number of tours is so staggeringly large (roughly  $N!$  factorial) that this approach is infeasible. For large  $N$ , no one knows an efficient method that can find the shortest possible tour for any given set of points. However, many methods have been studied that seem to work well in practice, even though they are not guaranteed to produce the best possible tour. Such methods are called heuristics. Your main task is to implement the *nearest neighbor* and *smallest increase insertion* heuristics for building a tour.

40



## Modern AI: Machine learning

Machine learning is the branch of AI that concerns systems *that can learn from data*.

### Applications

- Optical Character Recognition.
- Medical treatment.
- Robot navigation.
- Speech recognition.
- Genome analysis.
- Web commerce
- ...

"How can we program systems to automatically learn and to improve with experience?"



Tom Mitchell, CMU

"Machine learning is the science of getting computers to act without being explicitly programmed."



Andrew Ng, Stanford

### Approaches

- Sound statistical, computational and mathematical principles.
- Theories of learning processes.
- Methods of representing knowledge.

41

## Practical bottom line

The pursuit of AI has helped open limitless possibilities for applications of computer science.

### UK gov revving up for driverless cars on Britain's roads



Britain could soon be swarmed with driverless cars. If the UK is to lead in the development of driverless cars, it will have to do so before new legislation and regulatory framework is in place. Then the unrestrained will come on our roads.

### IBM's Watson uses Jeopardy become House-like medical diagnostician



IBM's Watson (left) has a medical career going, but he needs his first job. Watson's Jeopardy! success was a milestone, but he needs his first job. Watson's Jeopardy! success was a milestone, but he needs his first job. Watson's Jeopardy! success was a milestone, but he needs his first job.

### Scientists See Promise in Deep-Learning Programs



A video recognition program revealed a special gem by itself? Scientists see the promise of deep learning in the brain recognizes patterns, technology companies are reporting exciting gains in fields as diverse as computer vision, speech recognition and the identification of promising new molecules for drug design.

### Sizing Up Big Data, Broadening Beyond the Internet



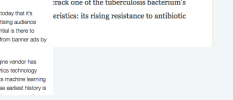
Petabytes of raw information could provide clues for everything from preventing TB to shrinking health care costs—if we can figure out how to use them.

### Yandex Boosts Precision Ad Targeting; Machine-Learning Method MatrixNet Is Behind The Scenes



Search engine Yandex said today that it's boosting its precision advertising audience targeting, and that the catalyst is three to increase clickthrough rates from banner ads by hundreds of percent.

### The promise of big data



Health microbiologist Sarah Fortune went in 2010 to attend a small but widely revered track one of the tuberculosis bacterium's evolution: its rising resistance to antibiotic

And **you** are prepared to begin participating!

42



Congratulations, and good luck!

43

