


Final Lecture: Fun, mainly



Today's Plan

Minesweeper

Conway's Game of Life

The Busy-Beaver function

Eliza

The Turing Test: Can a machine be "intelligent" ?

The Chinese Room: Maybe not.

A Story about a Barometer


Minesweeper Consistency Problem

Minesweeper.

- Start: Blank grid of squares, some conceal mines.
- Goal: Find location of all mines without detonating any.
- Repeatedly choose a square.
 - if mine underneath, it detonates and you lose
 - otherwise, computer tells you # neighboring mines



MINESWEEPER. Given a state of what purports to be a N-by-N Minesweeper game, is it logically consistent?



yes




no

Minesweeper Consistency Problem


SAT reduces to MINESWEEPER.

- Build circuit by laying out appropriate minesweeper configurations.
- Minesweeper game is consistent if and only if circuit is satisfiable.



A Minesweeper Wire

Exactly one of x and x' is a bomb.



A Minesweeper NOT Gate

Minesweeper Consistency Problem

SAT reduces to MINESWEEPER.

- Build circuit by laying out appropriate minesweeper configurations.
- Minesweeper game is consistent if and only if circuit is satisfiable.



A Minesweeper OR Gate




A Minesweeper AND Gate

Reference: these are due to Robert Krone '08

Conway's Game of Life



John Conway



hacker's emblem

Conway's Game of Life

In an infinite square grid, critters live and die.

Time proceeds in discrete steps ("generations").

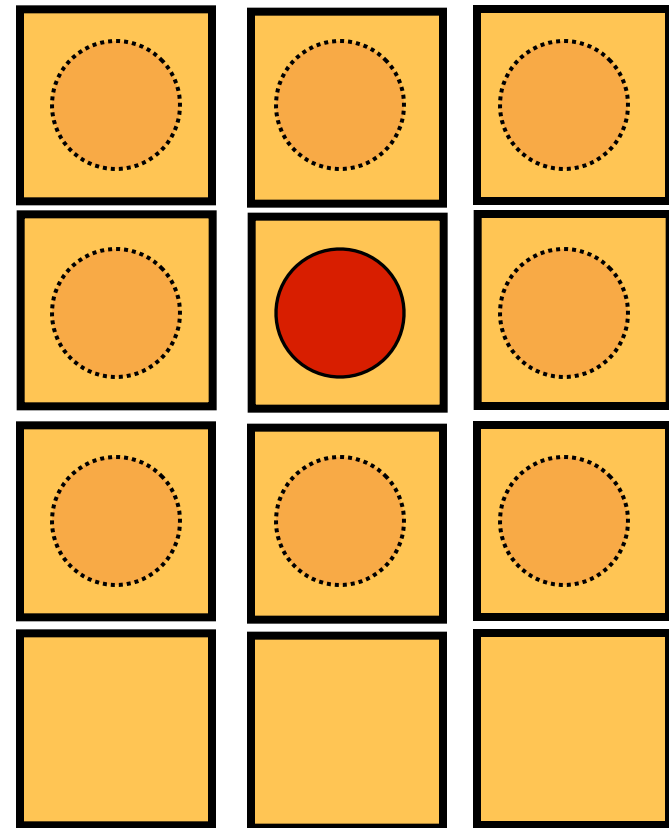
Here's one live critter: 

Its 8 neighbors look like: 

(Here's Conway: )


A critter's life and death depend on its 8 neighbors:

- too few? (0 or 1)
 - die of loneliness
- too many? (4-8)
 - die of overcrowding
- just right? (2 or 3)
 - survive to next generation
- exactly 3 parents?
 - critter born in empty square




Conway's Game of Life


Glider. Propagates a signal.




time t




time t+1



time t+2




time t+3



time t+5

What could this be?




Busy Beaver function



n	$BB(n)$
1	1
2	4
3	6
4	13
5	

$BB(n)$ is the largest number of 1's printed by any n -state Turing Machine (that halts). (Tape alphabet is blanks and 1's.)

Artificial Intelligence (a Tiny Glimpse)



de Kempelen del. *Cher à Mochel occid. Basilea.* P. 6.
Der Schachspieler im Spiele begriffen. | Le Joueur d'Échecs tel qu'on le voit pendant.

Origins

Idea of programming computers for "intelligent" behavior.


- First suggested by--who else?--Alan Turing, 1950.

Term "artificial intelligence" coined by John McCarthy in 1955.

Dartmouth summer conference, 1956.

- Gathering of field's luminaries
- Very optimistic!

"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."




J. McCarthy, *51

Optimistic predictions very common in 50's and 60's.

- Actual progress much slower than predicted.
- Now (2013): some striking successes; still lots to do.

Chess

Challenge. [Claude Shannon] Develop a computer program to play chess.



Number of possible games. $\approx 10^{23}$.

N-by-N version. EXPTIME-complete.

Kasparov vs. Deep Blue

Deep Blue. [IBM]

- Supercomputer, augmented by VLSI chess chips.
- 200 million board positions per second.



Machine beats man. [February 1996] First computer program to win a chess game against reigning world champion.

Q. Does Deep Blue appear intelligent?

← modern chess programs run on a laptop are now even better than Deep Blue

Checkers (Draughts)




Number of possible games. $\approx 10^{31}$.
N-by-N version. EXPTIME-complete.

Checkers (Draughts)

Chinook. [Jon Schaeffer, 1989] Computer program for checkers.

Man vs. machine. Chinook awarded man-machine world championship in 1994 after 6 draws with Marion Tinsley (who withdrew).



Checkers is solved! [Science, September 2007]

- Proof that black will never lose with optimal play.
- Proof that white will never lose with optimal play.

Backgammon

Backgammon. Two-player game of skill and luck.



TD gammon. [Gerry Tesauro 1980s]

- Program was given no expert backgammon knowledge.
- Learned strategy by playing itself 300,000 times.
- Among best backgammon players in the world

Poker

Poker. Multi-player game of skill and luck.



Texas hold 'em poker bots. Not currently competitive with pros.

Eliza

Weizenbaum, 1966.

- First "chatterbot": program that mimics human conversation.
- Just 200 lines of code!
- Works by generating scripted responses to keywords.

Program imitates Rogerian psychiatrist.

- "Active listening"
- Psychiatrist doesn't disagree; does not delve into the unconscious.

Is Eliza "intelligent"?

- You be the judge: here's Eliza on the Web.

Eliza Pattern Matching Rules

. . .

key: perhaps

decomp: *

reasmb: You don't seem quite certain.

reasmb: Why the uncertain tone ?

reasmb: Can't you be more positive ?

reasmb: You aren't sure ?

reasmb: Don't you know ?

key: francais

decomp: *

reasmb: goto xforeign

reasmb: I told you before, I don't understand French.

key: sorry

decomp: *

reasmb: Please don't apologise.

reasmb: Apologies are not necessary.

reasmb: I've told you that apologies are not required.

key: apologise

decomp: *

reasmb: goto sorry

. . .


Turing Test

Judge types questions to both.

- Computer claims to be human.
- (So does human, btw!)



If judge cannot tell which entity is human, we say computer has passed the test and is "intelligent."



Turing Test

Loebner Prize

- \$100,000 to first program to pass modified Turing Test.
- Annual competition held since 1995.
- Small awards given to best programs.



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 your output paper in response to various inputs.



Chinese Room Experiment (Searle 1980)

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 down, and pass it out.

什麼帶來
快樂
爲天下式

If you see this shape,
"什麼"
followed by this shape,
"帶來"
followed by this shape,
"快樂"

then produce this shape,
"爲天"
followed by this shape,
"下式".



http://www.mind.ilstu.edu/curriculum/searle_chinese_room/searle_chinese_room.php

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?



Chinese Room Experiment

"Weak" AI hypothesis.

- Machines can be programmed to **exhibit** intelligent behavior.
- Surely true: witness Deep Blue, TD-Gammon, others.
- Programs use methods very different from humans.
 - Performance (of task) vs. Simulation (of human methods).

"Strong" AI hypothesis.

- Machines can be programmed to **possess** intelligence.
- Must they use brain-like methods (e.g., neural nets) ?
 - "Connectionism"

Searle used Chinese Room as absolute refutation of the possibility of strong AI. But many disagree!

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

"Reverse" Turing Test

Standard Turing Test: judge is human.

Reverse Turing Test: judge is computer!

Why?

- Google allows each user 10 Gbytes of Web storage.
 - You write a "bot" to sign up 1 million users.
 - Congratulations. You now have 10 Petabytes of storage !
- PayPal once offered \$5 for each user who opens a new account.
 - You write a bot to sign up 1 billion users.
 - Congratulations. You now have \$5,000,000,000 !
- Both need to distinguish real humans from bots (programs).



CAPTCHA.

- Completely Automated Public Turing test to tell Computers and Humans Apart

Captchas

OCR. Given degraded text, find original text.

CAPTCHA. [completely automated public Turing test to tell computers and humans apart]



http://online.wsj.com/public/resources/images/OB-AB313_captch_20060524170113.gif

reCAPTCHAs

The screenshot shows the reCAPTCHA website interface. On the left is a navigation menu with the following items: HOME, WHAT IS reCAPTCHA (with sub-items: DIGITIZATION ACCURACY, WHAT IS A CAPTCHA, SECURITY), GET reCAPTCHA, MY ACCOUNT, EMAIL PROTECTION, and RESOURCES. The main content area features the heading "Digitizing Books One Word at a Time" and a challenge box. The challenge box contains the word "dressed" and a handwritten word "maoti". Below this is a text input field with the prompt "Type the two words:" and a "Submit" button. To the right of the input field are icons for refreshing, audio, and help. The reCAPTCHA logo and tagline "stop spam. read books." are also visible. Below the challenge box, a message states: "The words above come from scanned books. By typing them, you help to digitize old texts." At the bottom of the page, a paragraph explains: "reCAPTCHA is a free CAPTCHA service that helps to digitize books, newspapers and old time radio shows. Check out [our paper](#) in Science about it (or read more below)."

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 artificial intelligence 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, *54)

"Either artificial intelligence is possible...or we're not." (Herb Simon)

A Final Thought

“[Princeton’s Neuroscience program] has defined as its mission the understanding of how the brain works, how the wiring diagram is assembled during development, and how that wiring diagram results in complex processes like learning and memory and complex thinking...”

– Shirley Tilghman



Q. If we can understand the brain, can we simulate it?

