Lecture 23: Artificial Intelligence

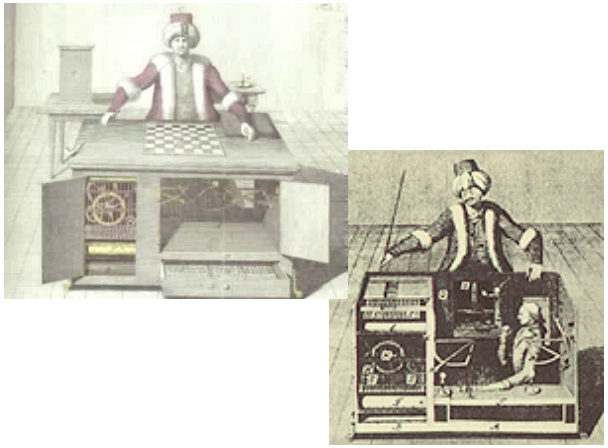

Overview

A whirlwind tour of Artificial Intelligence.

We spend just one lecture, but there are:

- Many courses on AI, but only one at Princeton :- (
- Professorships of AI.
- Entire university departments of AI.

Today's level of aspiration.

- A quick survey of several important topics.

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."

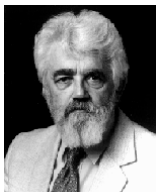

J. McCarthy, *51

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

- Actual progress much slower than predicted.
- Now (2004): some striking successes; still lots to do.

8-puzzle

Slide tiles until they're in numerical order.

Of course there is: 8-puzzle on the Web.

What strategy to use to solve puzzle ?

- No obvious algorithm.

From this position, →

- Three successor positions.
- From each of these,
  - Two, three, or four successors.
  - And so on.

7	3	5
1	8	4
2		6

Sounds like a tree?

- Yup.

8-puzzle: Breadth-First Search

Search tree for Goal node, one level at a time.

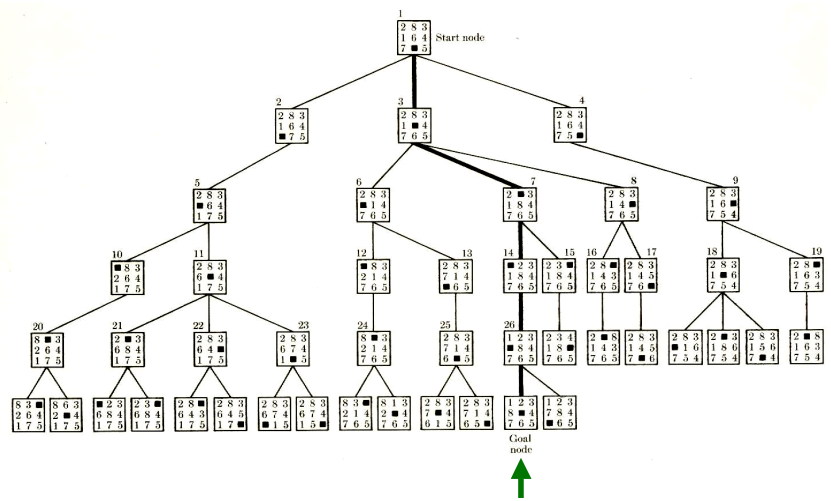

8-puzzle: Depth-First Search

Search for Goal down to some depth (the search "horizon").

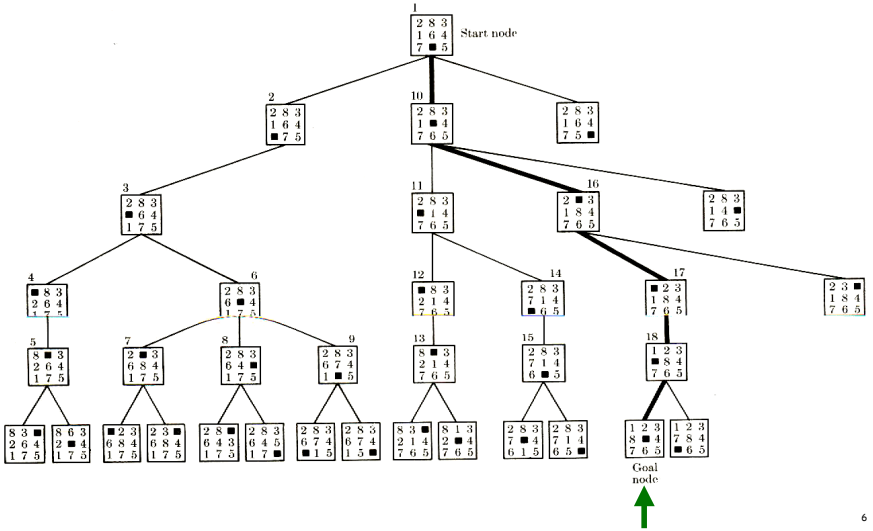

8-puzzle: Heuristic Search

Breadth-First and Depth-First are "blind" searches.

- Exhaustive methods for finding path to goal.
- Often infeasible because too many nodes expanded.
- Success (eventually) guaranteed.

"Heuristic" search.

- Uses "evaluation function" to rank successor nodes; pick best.
- No guarantee of success.
- Example uses distance from start plus number of tiles out of place
- Many other functions possible.
- Note: only 6 nodes expanded

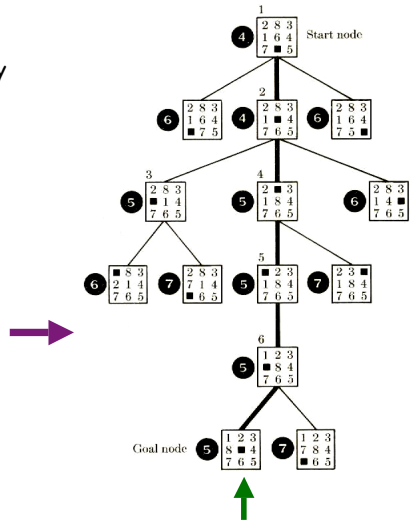

Game Trees

What if you have an opponent?

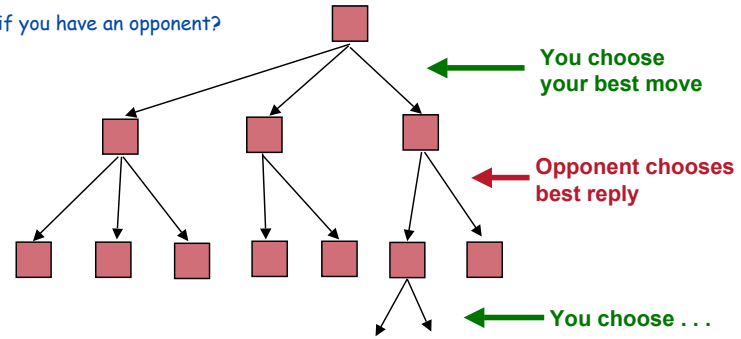

You should choose move that minimizes opponent's best chance.

- And your opponent should do likewise.
- "Minimax" methods use heuristics that assume you make choice that's best for you (max) and opponent makes choice that's worst for you (min).

Tic-tac-toe

Partial game tree.

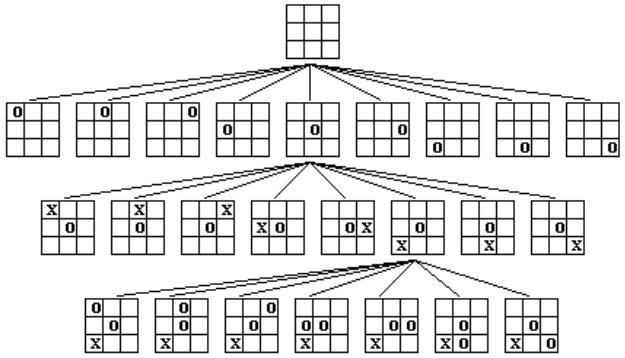

Tic-tac-toe

Reduce tree by recognizing symmetries.

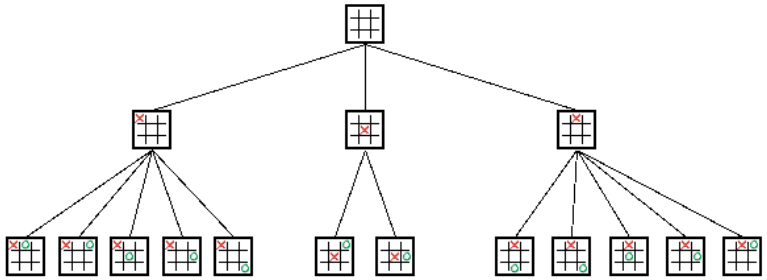

Sizes of game trees

8-puzzle.

- First tile can be in any one of 9 squares.
- Second tile in one of 8.
- Total number of nodes = 9! / 2 = 181,440

Tic-tac-toe (ignoring symmetries).

- First move: one of 9 spots.
- Second move: one of 8.
- Some games terminate before 9 moves made.
- Total number of nodes < 9! = 362,880

Both numbers small, so exhaustive search feasible.

But what about some bigger game, for instance . . .

Chess

A favorite target of AI researchers since 1950's.

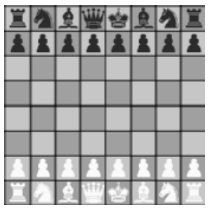

How big is game tree?

- 20 possible first moves.
- 35 possible next moves on average.
  - called "branching factor"
- Suppose games average 40 moves (each side).
- Complete game tree therefore has

$$35^{80} = 10^{120} \text{ nodes !}$$

So if each electron in the universe were a supercomputer, etc., etc.

Any computer (and any person) can search only the tiniest fraction of the complete game tree.

Chess

Lots of effort, and steady improvements, since 1950's.

Deep Blue chess machine developed at IBM

- Hsu, Campbell, Tan, Hoane, Brody, Benjamin, 1990's.
- 32-processor parallel computer.
- 512 special chess chips.
- Evaluates 200 million positions/sec.
- Huge amount of chess knowledge programmed in.
- Uses heuristic functions and minimax tree search.

In a 6-game match in 1997, Deep Blue defeated Garry Kasparov, the World Chess Champion.

Backgammon

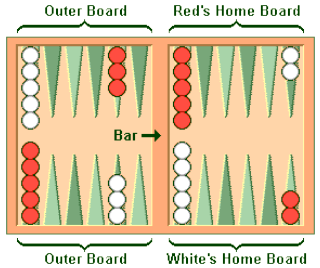

Moves use roll of 2 dice.

- 21 different outcomes.

About 20 possible moves, on average, per dice roll.

Game tree has branching factor of about 400. (That's big.)

Searches of even moderate depth computationally prohibitive.

- Really good heuristic function VERY important.

Something better must be done. . .

TD-Gammon

G. Tesauro, 1995.

Uses *neural network* and *reinforcement learning*.

- Each connection has a weight.
- Each node outputs weighted sum of inputs.
- Weights adjusted at each step.

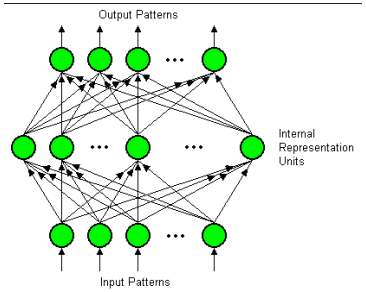

Input patterns.

- Describe board position.

Output patterns (just 4).

- Probability of each outcome.

Weight adjustment function.

- Tries to make successive predictions agree.
- Hence: Temporal Difference.

TD-Gammon

Initial state of neural net: random weights!

- Program begins "training" with practically no knowledge of game!
- (Well, OK, it knows what a legal move is, but that's almost all.)

Training ("learning" ?) process.

- Program evaluates predicted game outcomes for all legal next moves.
- Program picks move with best outcome for current player.
- Then:
  - Adjusts weights according to TD function.
  - Switches sides and repeats until game over (plays against itself).

Results

- Early random games took thousands of moves to finish!
- After 1.5 million games, program became *best player in the world!*

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.

Turing Test

Judge types questions to both.

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

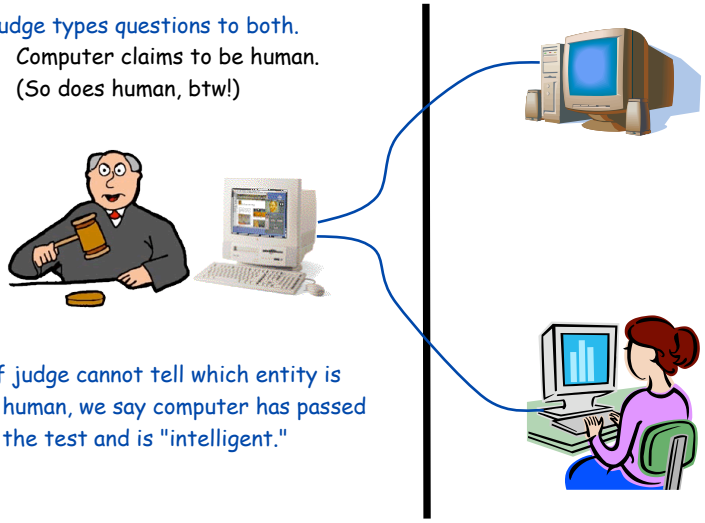

If judge cannot tell which entity is human, we say computer has passed the test and is "intelligent."

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 are 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.

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.

The Big Question.

- The folks outside think you understand Chinese. Do you?
- 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!

"Reverse" Turing Test

Standard Turing Test: judge is human.

Reverse Turing Test: judge is computer!

Why?

- Yahoo allows each user 15 Mbytes of Web storage.
  - You write a "bot" to to sign up 1 million users.
  - Congratulations. You now have 15 Terabytes 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.

- **C**ompletely **A**utomated **P**ublic **T**uring test to tell **C**omputers and **H**umans **A**part

Visual CAPTCHAs

Samples from CMU:

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)

AI in Everyday Life

Many examples of AI methods at work in the real world.

Microsoft Office's helpful talking paperclip.

- R.I.P. ?

Google (and other "intelligent" search engines).

Google

Speech recognition.

- Speak slowly and clearly to the telephone robot.

Optical character recognition (OCR).

- Makes U.S. Postal Service happy.

Control of spacecraft !

- AI system given control of Deep Space 1 for 2 days in May 1999.

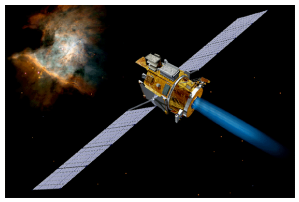

Omitted Topics

Knowledge representation.

Reasoning.

Expert systems.

Natural language understanding.

Speech recognition.

Computer vision.

And dozens more. . .

(But hey, we only had the one lecture.)

Summary

Games computers play.

- Some easy: 8-puzzle, tic-tac-toe.
- Some formerly hard, now easy: checkers, chess, backgammon.
- Some still hard: Go.

Methods.

- Exhaustive search (small problems only).
- Heuristic search.
- Neural networks.
- Special-purpose machines with built-in knowledge.

Turing Test

- Convincing conversation still a challenge!
- Chinese Room experiment refutes (?) possibility of strong AI.

Interested? Try COS 402, Introduction to Artificial Intelligence.

Announcements

Exam still not cancelled

- Wed April 27, 7:30 PM, right here
- Closed book, but
- You can bring one *cheatsheet*
  - *both sides* of one (8.5 by 11) sheet, handwritten by you
- P.S. No calculators, laptops, Palm Pilots, talking watches, etc.

Helpful review session

- Tonight, 7:30 PM, COS 105
- Not a canned presentation
- Driven by your questions, so bring some!