

EECS 391: Introduction to AI

Soumya Ray

Website: http://enr.case.edu/ray_soumya/eecs391_sp13/

Email: sray@case.edu

Office: Olin 516

Office hours: M 9-10:30

Announcements

- PA2 out later today
- HW2 due, quiz next week
 - All of search, emphasizing optimization and games
- Read chapter 7 of book

Today

- Case study: Deep Blue (not on quiz)
- Propositional Logic (chapter 7)

Case Study: Deep Blue

- ○



How Does Tournament Chess Work?

- Chess time = 90m for 40 moves
 - And then 30m + 30s/move for rest
- Average time 2¼ minutes
- Threefold repetition
 - Board in same state 3 times (not necc. in a row)
- 50 move with no capture or pawn moves



In the beginning...

- How well does minimax do in chess?
 - Chess: branching factor $b=35$
 - Suppose we could process 1,000,000 positions/sec
 - Chess time $\approx 2\frac{1}{2}$ minutes
 - Can process 150m positions/move
 - What is our search depth?
 - Minimax time complexity $O(b^m)$
 - $35^4 = 150\text{m}$
 - $35^5 = 500\text{m}$ (5 plies = $2\frac{1}{2}$ moves)
 - $35^6 = 1,838\text{m}$
 - One more ply – Importance?
 - Believed: one ply = 50-70 Elo points

Elo


- Way to rank chess players
 - Named after professor Arpad Elo, physicist
- Score based on opponents
- Popular
 - Replaced Harkness, USCF 1960, FIDE 1970
 - Used in other games (including video games)
- Not apples to apples
 - Have to set a few parameters
 - USCF, FIDE, etc. have unique implementations
- FIDE the most important

Elo	Title	Quantity
2800+		4
2700		50
2600		200
2500	Grandmaster	600
2400	Master/Grandmaster	1,400
2300	Master	3,000
2200	Candidate Master	6,000

Can't really compare computer Elo to human

Alpha-Beta to the rescue

- Lets you ignore subtrees that are irrelevant
 - $\downarrow b$, means can search deeper in same amount of time
- Random ordering, α - β $O(b^{3m/4})$
- Best case: $O(b^{m/2})$
 - b effectively becomes \sqrt{b}
 - Chess $b = 35$ becomes $b = 6!$
 - $6^5 \approx 8k$, $6^6 \approx 47k$, $6^{10} \approx 60m$
 - 150m positions = 10 ply (5 moves)
 - Increase of 5 ply = 250-350 Elo
- Goal: Find accurate α and β bounds early
 - Requires evaluating boundary (best) nodes first



Elo	Title	Quantity
2800+		4
2700		50
2600		200
2500	Grandmaster	600
2400	Master/Grandmaster	1,400
2300	Master	3,000
2200	Candidate Master	6,000

Node ordering heuristics

1. Captures,
2. Threats,
3. Forward Moves,
4. Backwards Moves

Null Move Heuristic

- Observe:
 - Generally every move made by opponent will improve their situation
 - Corollary: If I am able to skip my turn (give opponent free move) and I'm still really good, I have a good position
 - If skip a move and still above β , Min won't let you get here

Singular Extension

- General idea – spend more time looking deeply at moves that are good
 - Not necessarily non-quiet
- More specifically – look at “forced” moves
 - Includes “qualitatively forced”
- Forced 1 – you have no other choice
- Forced 2 – one move significantly better than all the rest
 - v =value, Δv =how much better, parameter we choose
- In general, SE turns out to be really good
 - Helped Deep Blue go a lot deeper

One More Cutoff Issue

- Average time $2\frac{1}{4}$ minutes
- But some moves deserve more attention than others
- Try to save a little each time
 - Put it in the time bank
- Write function to detect danger

State Evaluation Function

- Weighted linear function
 - $v = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s) = \sum w_i f_i(s)$
 - s = state, $f(s)$ might be # pawns, # bishops, etc.
 - Simplest version normally just material
 - Works because values/winnability in chess roughly linear to pieces
- More features: pieces in center, pawn structure, king protection, etc.
 - Deep Blue has 8,000
 - Many respond to weird, rare situations
 - Written by chess experts (grandmasters)

How Else Can We Go Fast?

- Pondering
 - Search during *opponent's* turn
- Representation
 - BitBoard
- Opening Book
- Endgame Database

Opening Book

- Start of game has a REALLY deep and wide search tree
- But, it's the same every single time you play the game
- Every opponent normally plays the same moves
 - Well, one of a couple 100 patterns
- Why not just memorize the beginning?
- Works for roughly 10 moves
 - After that, end up in unexplored territory

Endgame Database

- Says “Here’s every single set of moves that can result”
 - “These are the best”
- Hard-coded moves, not strategy
 - Experts know strategy for playing KRK, a lot less memory required
- Nalimov tables the most famous; on Web for free
- Heavily compressed to save room. Even so...
 - 5 piece – 7GB
 - 6 piece – 1.2 TB
- Issue: DB competes with search memory, transposition table
 - Less transpose = more searching / less search depth

Endgame Database (contd.)

- Pro: Can tell when you're going to win
- Con: Can't actually tell when you'll lose or draw
 - Assumes opponent has perfect play
 - Humans screw up a lot, especially on long sequences
 - Computers often resign too early
- DB not say how easy/hard a pattern is
 - Some patterns very tricky, easy for human to screw up
 - Computer won't pick those over easier ones (doesn't model human opponent)
- Nalimov tables and others don't know the 50-move rule

Legend of Deep Blue

- CMU had 2 chess enthusiasts
 - Professor Hans Berliner (chess/AI expert)
 - Student Feng-hsiung Hsu, (hardware expert)
- 1989 IBM hires Hsu's Deep Thought team
- 1996, Kasparov beats Deep Blue 4-2
 - Deep Blue does 100m positions/sec
 - Deep Blue says “I’ll be back” (metaphorically)
- 1997, Deep Blue upgraded for rematch
 - Deep Blue does 200m positions/sec, #259 on TOP500
 - 30 P2SCs, 480 custom VLSI
 - P2SC single chip PowerPC CPU, 135MHz, 11.38 GFLOPS
 - Comparison: Single Intel i7-3770: 109 GFLOPS

Legend of Deep Blue (contd.)

- 200m positions/sec means Deep Blue looks at 30 billion positions per move (assuming $2\frac{1}{2}$ m per move)
- *Mostly hardware* but some software
 - Specifically Singular Extension and null move
- Depth limit = 6 ply (maybe, actual system unknown)
 - SE = 40
- Branching factor = 3
 - Down from 35 using *lots* of pruning
 - Null move heuristic, alpha-beta
- Position evaluation function – 8,000 features
- Opening book DB – 4,000 positions
- History DB – 700,000 grandmaster games
 - Selected by GMs Illescas, Fedorowicz and de Firmian
- Endgame DB – All 5 piece endings, many 6

Legend of Deep Blue (contd.)

- Deep Blue's play style in the rematch unnerved Kasparov
 - Round 1, Kasparov did not understand why Deep Blue resigned
 - Checking the last moves, he found there was a forced mate---in 20 moves!!!
 - “The conclusion was a little bit scary... Deep Blue had actually worked it all out, down to the very end and simply chosen the least obnoxious losing line. ... Garry [was] thankful that he had been on the right side of these awesome calculations.”— Frederic Friedel, attached to Kasparov's team, chessbase.com
 - Round 2: Kasparov loses
 - “Deep Blue was playing deep strategic chess, "just like Karpov", the experts on the stage agreed. They could hardly believe it was the same machine that lost yesterday's game. Patrick Wolff said: "Show it to any expert and he'll guess it was played by a very strong grandmaster." ...Back at the hotel Garry was deep in his analysis of what had happened. There were many mystifying moments in the game. He was especially puzzled by 36.Qb6 instead of 36.axb6, and 37.Be4, instead of 37.Qb6. No amount of coaxing would persuade Fritz4 or Hiarc6.0 to waver from Qb6 in both cases. Deep Blue must contain some highly sophisticated positional instructions that have never before been seen in a chess program. Or it calculated the variations to depths that defy imagination.”

Legend of Deep Blue (contd)

- Rounds 3 and 4, Kasparov played “anti-computer” chess (later)
 - Draws
- Round 5, Kasparov was in a winning position, but Deep Blue forced the match to a draw
 - “But just when everybody at last seemed to agree that [Kasparov] was actually winning Deep Blue pulled out a most amazing [defense], forcing Garry's king into a perpetual. The [maneuver] left the entire auditorium gasping, but, Garry told us later that he had seen it coming...But he was baffled by another move of the computer.”
 - In the conference after the game, Kasparov admits he “is afraid of the machine”

Legend of Deep Blue (contd)

- Round 6

- “Disaster. I could sense it already early that morning. In the two weeks we had all been together in New York I had never seen Garry so tense. He hardly spoke, and on the way to the site the mood was dark. When I saw the game start I really had a sinking feeling. In 19 moves it was all over, with the world champion falling into a well-known openings trap...Just over an hour later the world champion had resigned against the computer.”

Legacy of Deep Blue

- Houdini (3311 Elo CCRL)
 - alpha-beta search, bitboard, late move reduction, Nalimov end-game tablebases
 - Released 2011
- Rybka (3261 Elo)
- Deep Junior (3121 Elo)
- Deep Fritz (3095 Elo)
- Hydra (3000+, guesstimated)
 - Supercomputer with fancy hardware
- Deep Blue?
 - Didn't play enough games. Maybe 2870?
- Kasparov 2851

Elo	Title	Quantity
2800+		4
2700		50
2600		200
2500	Grandmaster	600
2400	Master/Grandmaster	1,400
2300	Master	3,000
2200	Candidate Master	6,000

Legacy of Deep Blue (contd.)

- 2002, Kramnik vs. **Deep Fritz**, tie
 - Five years after Deep Blue, computers are four times as fast
- 2003, Kasparov vs. **Deep Junior**, tie
- 2005, Adams (#7) vs. **Hydra**, 5½-½ Hydra
 - 64 CPU, custom FPGA chips, 150m positions/sec
 - D=18 with more pruning
- 2006, Kramnik vs. **Deep Fritz**, 4-2 Fritz
 - Dual Core 2 Duo, 1.3 GHz, 8m positions/sec
 - Two orders of magnitude less than Hydra
 - D=18
 - “The science is done” – Professor Monty Newborn
- 2009, **Pocket Fritz** wins Copa Mercosur in the National Master category
 - HTC Touch HD, 20,000 positions/second
 - It’s a freaking phone!

So Chess is pretty much over

...or is it?



BATTLESTAR
GALACTICA

THE RESISTANCE


Anti-Computer Chess

- Move quick
 - Reduce pondering
- Use unusual (non-book) opening (Kasparov did this in games 3 and 4)
 - Force them to use their time early
 - Reduces # possible plies later
- Don't swap pieces
 - Keep the branching factor high
 - Beware end games; 5 piece DBs are common
- Control the center from somewhere else
 - Computer likes center, rates it high
 - Fianchetto (Sicilian Dragon, Pirc Defense, Benko, Grob, etc.)
 - Fianchetto = stay out of center, put bishop on long diagonal

Anti-Computer Chess

- Poisoned pawn
 - It takes pawn, puts it ahead in material but hurts position
- Force it to trade
 - Most won't draw @ 50 if ahead
 - Forced trade takes it out of end game DB
- Fortress
 - 8 pawn chess – make a fence around the center

2008 Rybka vs. Nakamura



Move 111, White up 4

<http://www.chessgames.com/perl/chessgame?gid=1497429>

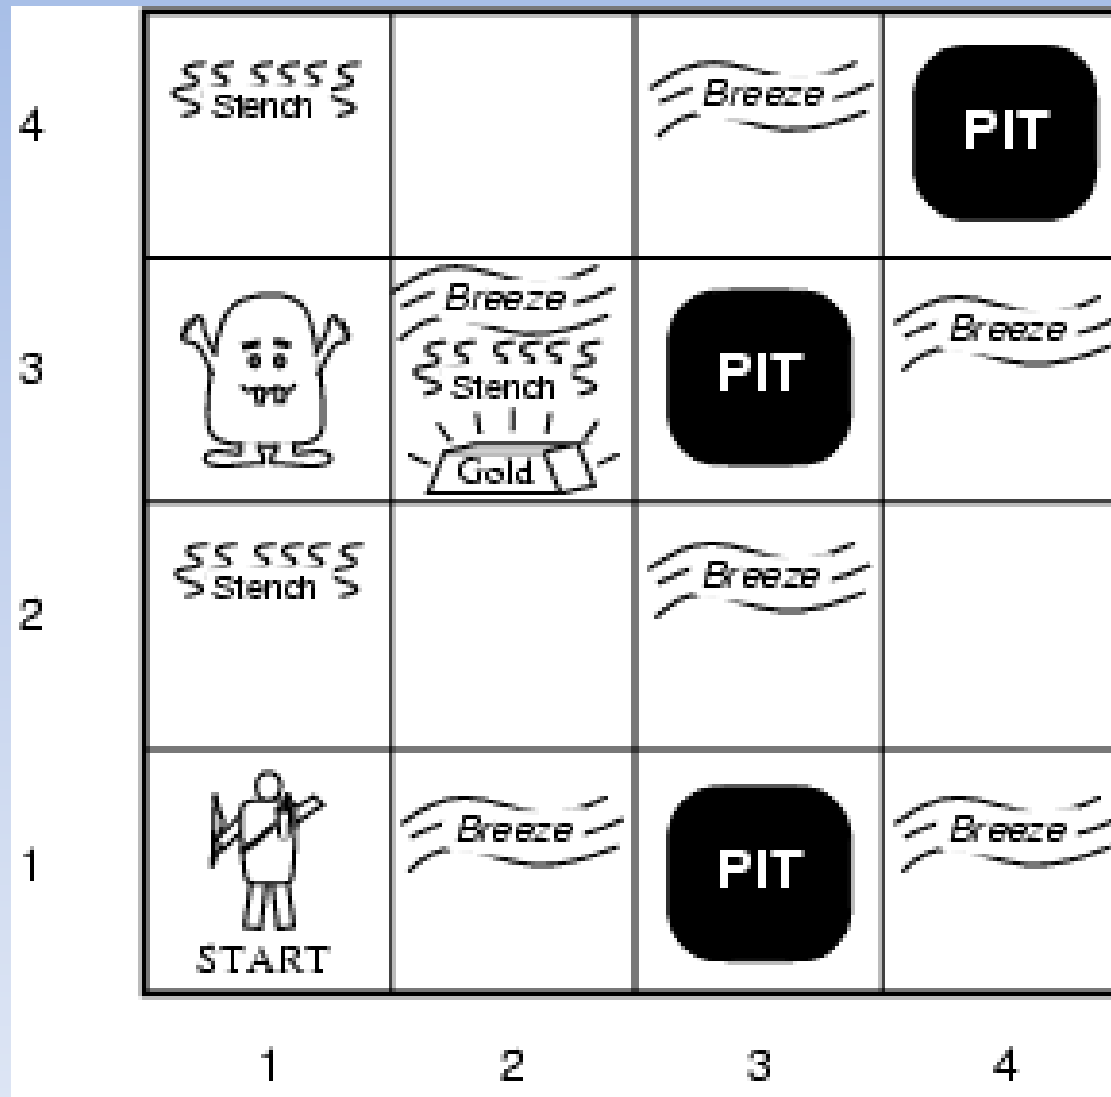
Summary

- We learned about:
 - Deep Blue was awesome
 - But humans aren't done yet!
 - And neither are the computers...

Representing the World

- Search is a very general problem solving technique, but it does not take advantage of *structure* in the world
- To represent structure, we need a *language* and some system of *reasoning* with it
 - Propositional and first order logics are such languages

Wumpus World



Formalizing the reasoning process

- To build an automated system that can carry out this sort of reasoning, we'll equip them with *knowledge bases of facts* and *inference algorithms* to derive new facts
- **Propositional logic** is a way to formalize the reasoning process described in the previous slide

Syntax and Semantics

- **Syntax**: what strings of symbols are *allowable sentences* in the logic
- **Semantics**: Given a well-formed sentence, *what does it “mean”*?

Possible Worlds and Models

- A “possible world” gives specific meanings to symbols in sentences
- A “model” abstracts possible worlds to assign values to symbols
- The semantics then defines the truth value of a sentence with respect to the model
















Example

- Consider the sentence $x+y=5$
- A possible world: “ x ” represents the number of books on a table and “ y ” represents the number of pens on the table
- A model: $x=2$ and $y=4$
 - The semantics then assigns the truth value “false” to the sentence above
- We say a model M is “a model of a sentence α ” if α is true in M

Entailment

- Given a bunch of sentences, we will need to identify what else follows from this
 - This is *logical inference (deduction)*
- Inference in logic involves establishing an *entailment* relationship: $\alpha \models \beta$
 - α entails β iff every model of α is also a model of β

Wumpus World

4	 Stench		 Breeze	 PIT
3	 Wumpus	 Breeze  Stench  Gold	 PIT	 Breeze
2	 Stench		 Breeze	
1	 START	 Breeze	 PIT	 Breeze
	1	2	3	4

Suppose our knowledge base has facts “Breeze in 2,1” and “Nothing in 1,1”

This KB entails “Nothing in 2,1” but does not entail “Nothing in 2,2”