



Why Are Games Good for AI?

- Games typically have concise rules
- Well-defined starting and end points
- Sensing and effecting are simplified
 - Not true for sports games
 - See robocup
- Games are fun!
- Downside: Getting taken seriously (not)
 - See robo search and rescue





Checkers: Tinsley vs. Chinook



| Name: | Marion Tinsley |
|----------------------------------|--------------------|
| Profession: | Taught mathematics |
| Hobby: | Checkers |
| Record: | Over 42 years |
| | loses only 3 games |
| | of checkers |
| World champion for over 40 years | |
| | |

Mr. Tinsley suffered his 4th and 5th losses against Chinook (1994)



| | ACM Chess Cha Garry Kaspar | · |
|---------------------|-------------------------------|--------------------------|
| Kasparov | | Deep Blue |
| 5'10" | Height | 6' 5″ |
| 176 lbs | Weight | 2,400 lbs |
| 34 years | Age | 4 years |
| 50 billion neurons | Computers | 32 RISC processors |
| | | + 256 VLSI chess engines |
| 2 pos/sec | Speed | 200,000,000 pos/sec |
| Extensive | Knowledge | Primitive |
| Electrical/chemical | Power Source | Electrical |
| Enormous | Ego | None |
| 1997: Deep Blue | wins by 3 wins, 1 lo | ss, and 2 draws |

Chess: Kasparov vs. Deep Junior

Deep Junior

2,000,000 pos/sec Available at ~\$1000

8 CPU, 8 GB RAM, Win 2000

(Note: Lesser hardware, but more clever software)



August 2, 2003: Match ends in a 3/3 tie!









Zero Sum Games

- Assign values to different outcomes
- Win = 1, Loss = -1
- With zero sum games every gain comes at the other player's expense
- Sum of both player's scores must be 0
- Are any games truly zero sum?































Why use backed-up values?

- If e is to be trusted in the first place, then the backed-up value is a better estimate of how favorable STATE(N) is than e(STATE(N))
- Why? Presumption that e() is more accurate closer to the end of the game











| Implementing alpha-beta | | |
|--|---|--|
| $\begin{array}{c} max_value(state, alpha, beta)\\ \mbox{if cutoff(state) then return eval(state)}\\ v = -\infty\\ \mbox{for each s in successors(state) do}\\ v = max(v, min_value(s, alpha, beta))\\ \mbox{if } v \ge beta then return v\\ \mbox{alpha} = max(alpha,v)\\ \mbox{end}\\ \mbox{return } v \end{array}$ | beta=value of best guaranteed option available to min | |
| alpha=value of best guaranteed option available to max Call: max_value(root, -∞, +∞) | $\begin{array}{l} \hline min_value(\text{state, alpha, beta}) \\ \text{if cutoff(state) then return eval(state)} \\ v = & \\ v = & \\ \text{for each s in successors(state) do} \\ v = min(v, max_value(s, alpha, beta)) \\ \text{if } v \leq alpha \text{ then return } v \\ beta = min(beta,v) \\ \text{end} \\ \text{return } v \end{array}$ | |



























































































How much do we gain?

- Assume a game tree of uniform branching factor b
- Minimax examines O(b^h) nodes = worst case for alpha-beta
- The gain for alpha-beta is maximum when:
 - The children of a MAX node are ordered in decreasing backed up values
 - The children of a MIN node are ordered in increasing backed up values
- Then alpha-beta examines O(b^{h/2}) nodes [Knuth and Moore, 1975]
- But this requires an oracle (if we knew how to order nodes perfectly, we would not need to search the game tree)
- If nodes are ordered at random, then the average number of nodes examined by alpha-beta is ~O(b^{3h/4})













