

Game Playing

Louis Oliphant
(slides borrowed from Burr H. Settles)
CS-540-2, UW-Madison
www.cs.wisc.edu/~cs540-2
Fall 2005

1

Announcements

- Read:
 - Chapter 6 – Adversarial Search
 - Chapter 17.6-17.7 – Game Theory
- Homework 1 due on Thursday
 - Written portion to me by beginning of class
 - Programmed portion handed in electronically by beginning of class

2

AI for Game Playing

- Game playing is (was?) thought to be a good problem for AI research
- Game playing is non-trivial
 - Players need “human-like” intelligence
 - Games can be very complex (e.g. chess, go)
 - Requires decision making within limited time
- Games usually are:
 - Well-defined and repeatable
 - Limited and accessible
- Can directly compare humans and computers

3


AI for Game Playing

	Deterministic	Chance
Accessible: perfect info	Tic-tac-toe, checkers, chess, mancala	backgammon, monopoly
Inaccessible: imperfect info	???	bridge, poker, scrabble

4

Greedy Search for Games


- Expand the search tree to the terminal states
- Evaluate utility of each terminal board state
- Make the initial move that results in the board configuration with the maximum value



9

Greedy Search for Games

- But this still ignores what the opponent is likely to do...
 - Computer chooses C because its utility is 9
 - Opponent chooses J and wins!



10


The MiniMax Principle

- Assuming the worst (i.e. the opponent plays optimally):
 - Given there are two plays till the terminal states
 - Low utility numbers favor opponent
 - Smart opponent chooses minimizing moves
 - High utility numbers favor computer
 - Computer should choose maximizing moves

11

The MiniMax Principle

- The computer assumes after it moves the opponent will choose the minimizing move
 - Therefore, it chooses the best move considering both its move and the opponent's best move



12


Propagating MiniMax Values

- Explore the tree to the terminal states
- Evaluate utility of the resulting board configurations
- The computer makes a move to put the board in the best configuration for it, assuming the opponent makes its best moves on its turn:
 - Start at the leaves
 - Assign value to the parent node as follows
 - Use minimum when children are opponent's moves
 - Use maximum when children are computer's moves

13

Deeper Game Trees

- MiniMax can be generalized to more than 2 moves
- Propagate utility values upwards in the tree



14

General MiniMax Algorithm

```

for each move by the computer {
  perform DFS to terminal states
  evaluate each terminal state
  propagate MiniMax values upward
  - if opponent propagate min value of children
  - if computer propagate max value of children
  choose move with maximum MiniMax value
}
    
```

Note:

- MiniMax values gradually propagate upwards as DFS proceeds (i.e. MiniMax values propagate up in "left-to-right" fashion)
- MiniMax values for sub-tree propagate upwards "as we go", so only $O(bd)$ nodes need to be kept in memory at any time

15

Complexity of MiniMax

- Space complexity
 - depth-first search (no closed list necessary), so $O(bd)$
- Time complexity
 - given branching factor b , $O(b^d)$
- Time complexity is a *major problem* since computer typically only has a finite amount of time to make a move!!

16

Complexity of MiniMax

- Direct MiniMax algorithm is impractical
 - Instead do depth-limited search to depth limit l
 - But evaluation defined only for terminal states
 - We need to know the value of non-terminal states
- Static board evaluator (SBE) functions use heuristics to estimate utility for non-terminal states

17

Static Board Evaluators (SBE)

- A static board evaluation function is used to estimate how good the current board configuration is for the computer
 - Reflects computer's chances of winning from that state
 - Must be easy to calculate from board configuration

- For Example, Chess:

$SBE = \alpha \times materialBalance + \beta \times centerControl + \gamma \times \dots$
 $material\ balance = Value\ of\ white\ pieces - Value\ of\ black\ pieces$
 $pawn = 1, rook = 5, queen = 9, etc...$

18

Static Board Evaluators (SBE)

- Typically, one subtracts how good it is for the opponent from how good it is for the computer
- If the board evaluation has utility x for a player, then it is usually considered $-x$ for opponent
- Must agree with the utility function that is calculated at terminal nodes

19

MiniMax Algorithm with SBE

```
function minimax (STATE, DEPTH, LIMIT) {  
  // base cases  
  if STATE is terminal then  
    return utility(STATE)  
  if DEPTH = LIMIT then  
    return sbe(STATE)  
  // continue search  
  else {  
    CHILDREN = empty list  
    foreach CHILD of STATE {  
      add to CHILDREN:  
        minimax(CHILD, DEPTH+1, LIMIT)  
      if computer's turn then  
        return max(CHILDREN)  
      else  
        return min(CHILDREN)  
    }  
  }  
}
```

20

MiniMax with SBE

- The same as general MiniMax, except
 - Only goes to depth l
 - Estimates using SBE function
- How would this algorithm perform at chess?
 - If could look ahead ~4 pairs of moves (*i.e.* 8 ply) would be consistently beaten by average players
 - If could look ahead ~8 pairs (16 ply) as done in typical PC, is as good as human master

21

Summary So Far

- MiniMax can't search to the end of the game
 - Otherwise, choosing a move is trivial
- SBE isn't perfect at estimating utility
 - If it was, just choose best move without searching
- Since neither is feasible for interesting games, combine MiniMax with SBE
 - MiniMax to depth l
 - Use SBE to score board configuration

22

Alpha-Beta Pruning

- Some of the branches of the game tree won't be taken if playing against an intelligent opponent
- We can "prune" those branches from the tree
- Keep track while doing DFS of game tree of:
 - Maximizing level: **alpha**
 - Highest value seen so far
 - Lower bound on node's utility or score
 - Minimizing level: **beta**
 - Lowest value seen so far
 - Higher bound on node's utility or score

23


Alpha-Beta Pruning

- When **maximizing** (computer's turn):
 - If $alpha \geq parent's\ beta$, stop expanding
 - Opponent shouldn't allow the computer to make this move
- When **minimizing** (opponent's turn):
 - If $beta \leq parent's\ alpha$, stop expanding
 - Computer shouldn't take this route

24


Alpha-Beta Example

$\text{minimax}(A, 0, 4)$




Alpha-Beta Example

$\text{minimax}(B, 1, 4)$




Alpha-Beta Example

$\text{minimax}(F, 2, 4)$




Alpha-Beta Example

$\text{minimax}(N, 3, 4)$




Alpha-Beta Example

$\text{minimax}(F, 2, 4)$ returned to
 $F: \alpha = 4$, maximizing




Alpha-Beta Example

$\text{minimax}(O, 3, 4)$




Alpha-Beta Example

$\text{minimax}(W, 4, 4)$



Alpha-Beta Example


$\text{minimax}(O, 3, 4)$ returned to
 $O: \beta = -3$, minimizing



Alpha-Beta Example

$\text{minimax}(O, 3, 4)$ returned to


$O: \beta \leq F: \alpha$, stop expanding O (alpha cutoff)



Alpha-Beta Example

Why?


Smart opponent will choose W or less, thus O's upper bound is -3
Computer shouldn't take O: -3 because N: 4 is better



Alpha-Beta Example

$\text{minimax}(F, 2, 4)$ returned to


F: α not updated, maximizing



Alpha-Beta Example


$\text{minimax}(B, 1, 4)$ returned to

B: $\beta = 4$, minimizing



Alpha-Beta Example


$\text{minimax}(G, 2, 4)$



37

Alpha-Beta Example


$\text{minimax}(B, 1, 4)$ returned to
 $B: \beta = -5$, minimizing



38

Alpha-Beta Example


$\text{minimax}(A, 0, 4)$ returned to
 $A: \alpha = -5$, maximizing



39

Alpha-Beta Example


$\text{minimax}(C, 1, 4)$



40

Alpha-Beta Example


minimax(H, 2, 4)



DFS Stack
H
C
A

Alpha-Beta Example


minimax(C, 1, 4) returned to
C: $\beta = 3$, minimizing



DFS Stack
C
A

Alpha-Beta Example


minimax(I, 2, 4)



DFS Stack
I
C
A

Alpha-Beta Example


minimax(C, 1, 4) returned to
C: β not updated, minimizing



DFS Stack
C
A


Alpha-Beta Example

minimax (J, 2, 4)




Alpha-Beta Example

minimax (P, 3, 4)




Alpha-Beta Example

minimax (J, 2, 4) returned to
J: $\alpha = 9$, maximizing



Alpha-Beta Example


minimax (J, 2, 4) returned to
J: $\alpha \geq C:\beta$, so stop expanding J (beta cutoff)



Alpha-Beta Example

Why?


Computer should choose P or higher, so J's lower bound is 9
A smart opponent shouldn't take J:9 because H:3 is lower



Alpha-Beta Example

$\text{minimax}(C, 1, 4)$ returned to


C: β not updated, minimizing



Alpha-Beta Example


$\text{minimax}(A, 0, 4)$ returned to

A: $\alpha = 3$, maximizing




Alpha-Beta Example

$\text{minimax}(D, 1, 4)$



Alpha-Beta Example


$\text{minimax}(A, 0, 4)$ returned to
 A: α not updated, maximizing



53

Alpha-Beta Example


Finishing off the search...
 $E:\beta \leq A:\alpha \dots$ so stop expanding E (alpha cutoff)



54

Alpha-Beta Example


Why?
 A smart opponent will choose L or less, thus E's upper bound is 2.
 The computer shouldn't choose E: 2 since C: 3 is better



55

Alpha-Beta Example


Result: Computer chooses move C



56

Now It's Your Turn!

- Let's try working out a MiniMax search with alpha-beta pruning on this game tree (going from left to right):
 - If $alpha \geq parent's\ beta$, stop expanding
 - If $beta \leq parent's\ alpha$, stop expanding



57

Effectiveness of Alpha-Beta

- Effectiveness depends on the order in which successors are examined (more effective if best are examined first)
 - Best Case:
 - Each player's best move is evaluated first (left-most)
 - Worst Case:
 - Ordered so that no pruning takes place
 - No improvement over exhaustive search
- In general, performance is closer to the best case than the worst case

58

Effectiveness of Alpha-Beta

- In practice often get $O(b^{d/2})$ rather than $O(b^d)$
 - Same as having a branching factor of \sqrt{b} since $(\sqrt{b})^d = b^{d/2}$
- Example: chess
 - Branching factor goes from ~ 35 to ~ 6
 - Allows for a much deeper search given the same amount of time
 - Allows computer chess to be competitive with humans

59

The Horizon Effect

- Sometimes disaster is just beyond the depth limit
 - Computer captures queen, but a few moves later the opponent checkmates and wins
- The computer has a limited horizon, it cannot see that this significant event could happen
- How do you avoid catastrophic losses due to "short-sightedness"?
 - Quiescence search
 - Secondary search

60

The Horizon Effect

- Quiescence Search
 - When evaluation frequently changing, allow looking deeper than the limit
 - Looking for a point when game quiets down
- Secondary Search
 1. Find best move looking to depth d
 2. Look k steps beyond to verify it still looks good
 3. If it doesn't, repeat step 2 for next best move

61


Stochastic Game Environments

- Some games involve chance, for example:
 - Roll of a die
 - Spin of a game wheel
 - Deal of cards from shuffled deck
- Extend the game tree representation:
 - Computer moves
 - Opponent moves
 - *Chance nodes*

62

Stochastic Game Environments


The game tree representation is extended:



63

Stochastic Game Environments


- Weight score by the probabilities that move occurs
- Use expected value for move: sum of possible random outcomes



64

Stochastic Game Environments

- Choose move with highest expected value



65

Stochastic Game Environments

- Stochastic elements increase the branching factor
 - 21 possible number rolls with 2 dice
 - The value of look-ahead diminishes: as depth increases, probability of reaching a particular node decreases
- Alpha-beta pruning is less effective
- ▢ See *AI: A Modern Approach* for more details

66

Limiting Search Time

* *In real games there is usually some time limit T on making a move*

- How do we take this into account?
 - Can't stop alpha-beta midway and expect to use results with any confidence
 - So, we could set a conservative depth-limit that guarantees we will find a move in time $< T$
 - But then, the search may finish early and the opportunity to search deeper is wasted

67

Limiting Search Time

- In practice, we use an iterative-deepening (IDS) approach
 - Run MiniMax with alpha-beta pruning at increasing depth limits
 - When the clock runs out, use the solution found for the last complete alpha-beta search (*i.e.* the deepest search that was completed)
- As with all heuristics, there is also a speed vs. accuracy tradeoff for board evaluation functions

68

Using Book Moves

- For well-studied games, maybe we know the move we should make without having to searching for it
- Build a database of opening moves, end-games, and common board configurations
- If the current game state is in the lookup table, use database:
 - To determine the next move
 - To evaluate the board
- Otherwise do alpha-beta search

69

Evaluation Functions

- * *The board evaluation function estimates how good the current board state is for the computer*
- Heuristic function of the features of the board
 - i.e. $\text{function}(f_1, f_2, f_3, \dots, f_n)$
- The features are numeric characteristics
 - f_1 = # of white pieces
 - f_2 = # of black pieces
 - $f_3 = f_1 / f_2$
 - f_4 = estimate of "threat" to white king, etc...

70

Linear Evaluation Functions

- A linear evaluation function of the features is a weighted sum of f_1, f_2, f_3, \dots
 $(w_1 \times f_1) + (w_2 \times f_2) + (w_3 \times f_3) + \dots + (w_n \times f_n)$
 - where f_1, f_2, \dots, f_n are features
 - and w_1, w_2, \dots, w_n are their weights
- * *More important features get more weight*

71

Linear Evaluation Functions

- The quality of play depends directly on the quality of the evaluation function
- To build an evaluation function we have to:
 - Construct good features using expert knowledge of the game
 - Choose good weights... or *learn* them

72

Learning Weights

- **Q:** How can we learn the weights for a linear evaluation function?
- **A:** Play lots of games against an opponent!
 - For every move (or game)
 $error = true\ outcome - evaluation\ function$
 - If error is positive (underestimating)
adjust weights to *increase* the evaluation function
 - If error is zero do nothing
 - If error is negative (overestimating)
adjust weights to *decrease* the evaluation function

73

Learning Checkers

- ▣ A. L. Samuel, “Some Studies in Machine Learning using the Game of Checkers,” *IBM Journal of Research and Development*, 11(6):601-617, 1959
- Learned linear weights by playing copies of itself thousands of times
- Used only an IBM 704 with 10,000 words of RAM, magnetic tape, and a clock speed of 1 kHz
- Successful enough to be competitive in human tournaments

74

Learning Backgammon

- ▣ G. Tesauro and T. J. Sejnowski, “A Parallel Network that Learns to Play Backgammon,” *Artificial Intelligence*, 39(3), 357-390, 1989
- Also learned by playing copies of itself
- Used a non-linear evaluation function: a neural network (we’ll discuss these models in the machine learning section of the course)
- Rates in the top three players in the world

75

IBM’s Deep Blue

- Current world chess champion
- Parallel processor, 8 dedicated VLSI “chess chips”
- Can search 200 million configurations/second
- Uses MiniMax, alpha-beta pruning, very sophisticated heuristics
- It can search up to 14 ply (*i.e.* 7 pairs of moves)
- Can avoid horizon by searching as deep as 40 ply
- Uses book moves


76

IBM's Deep Blue

- Kasparov vs. Deep Blue, May 1997
 - 6-game full-regulation chess match sponsored by ACM
 - Kasparov lost the match 2.5 to 3.5
- This was a historic achievement for computer chess because it became the *best chess player on the planet!!*
- Note: Deep Blue still searches “brute force,” and still plays with little in common with the intuition and strategy humans use

77

Chess Rating Scale



78


AI for Other Games

- Checkers
 - Current world champion is Chinook
 - Blondie24 won a 2001 online checkers tournament
 - Learned to play checkers with genetic algorithms
 - Used a neural network: wasn't even programmed with rules!
- Go
 - Branching factor is ~360 on average, very large!
 - Pretty much still play at novice levels these days
 - \$2 million prize for any system that can beat a world expert

79

AI in Modern Computer Games

- Modern computer games (*i.e.* “Doom,” “Civilization,” etc.) usually still use rudimentary AI
 - Finite state machines, simple reflex agents
 - *e.g.* the “scientist” AI schema for Half-life:



80

AI in Modern Computer Games

- Path-finding for FPS-type tournament arena games is often done using A* search with straight-line distance as a heuristic
 - Often makes the agent's moves "look like it's drunk"
- Remember: reflex agents aren't very adaptable, and behave very deterministically (not very human-like)
- ▢ S. Rabin, editor, *AI Game Programming Wisdom*, Charles River Media, 2002

81

AI in Modern Computer Games

- Genetic algorithms and genetic programming have been used and shown some success in "evolving" realistically-acting agents for games
 - Certainly appropriate for "Sim"-type games
- ▢ B. Geisler, "An Empirical Study of Machine Learning Algorithms Applied to Modeling Player Behavior in a 'First Person Shooter' Video Game," M.S. Thesis, UW-Madison, 2002
 - Used machine learning to learn typical player actions
 - Created a computer agent player based on learned behavior

82

Summary

- Classic game playing is best modeled as a search problem
- Search trees for games represent alternate computer/opponent moves
- Evaluation functions estimate the quality of a given board configuration for each player
 - good for opponent
 - + good for computer
 - 0 neutral

83

Summary

- MiniMax is a procedure that chooses moves by assuming that the opponent always choose their best move
- Alpha-beta pruning is a procedure that can eliminate large parts of the search tree enabling the search to go deeper
- For many well-known games, computer algorithms using heuristic search can match or out-perform human world experts

84

Summary

- Initially thought to be good area for AI research
 - But brute force has proven to be better than a lot of knowledge engineering
 - More high-speed hardware issues than AI
 - AI relatively simple, enabled scaled-up hardware
 - Still a good test-bed for machine learning
- * *Perhaps machines don't have to think like us?*