#### **Adversarial Search**

Dr. Melanie Martin CS 4480 October 1, 2010 Based on slides from http://aima.eecs.berkeley.edu/2nd-ed/slides-ppt/

# Outline

- Optimal decisions
- α-β pruning
- Imperfect, real-time decisions

#### **Games vs. search problems**

- "Unpredictable" opponent → specifying a move for every possible opponent reply
- Time limits → unlikely to find goal, must approximate

## **Game as Search Problem**

- Initial State
  - Board position and player to move
- Successor Function
  - Returns list of (state, move) pairs
    - Legal moves and resulting states
- Terminal (Goal) Test
  - When game is over
- Utility (Objective) Function
  - Assigns numeric outcome to terminal states
    - E.g. +1, -1, 0 for win, lose, draw

# Game tree (2-player, deterministic, turns)



## **Optimal Strategy**

- Leads to outcomes as good as any other strategy when playing an infallible opponent
- Tree where max takes a turn and min takes a turn is ONE MOVE DEEP made up two half-moves - each half move is called a ply

## Minimax

- Perfect play for deterministic games
- Idea: choose move to position with highest minimax value = best achievable payoff against best play
- E.g., 2-ply game:



## Minimax

- Minimax(node) is utility for max of being in corresponding state
- Max prefers a state with maximum value
- Min prefers a state with minimum value



## **Minimax algorithm**

function MINIMAX-DECISION(state) returns an action

 $v \leftarrow MAX-VALUE(state)$ return the *action* in SUCCESSORS(*state*) with value v

function MAX-VALUE(state) returns a utility value if TERMINAL-TEST(state) then return UTILITY(state)  $v \leftarrow -\infty$ for a, s in SUCCESSORS(state) do  $v \leftarrow MAX(v, MIN-VALUE(s))$ return v

function MIN-VALUE(state) returns a utility value if TERMINAL-TEST(state) then return UTILITY(state)  $v \leftarrow \infty$ for a, s in SUCCESSORS(state) do  $v \leftarrow MIN(v, MAX-VALUE(s))$ return v

#### Minimax

 Recursion proceeds to leaves and based on utility function assigns minimax values at the level above and so on



## **Properties of minimax**

- <u>Complete?</u> Yes (if tree is finite)
- Optimal? Yes (against an optimal opponent)
- <u>Time complexity?</u> O(b<sup>m</sup>)
- <u>Space complexity?</u> O(bm) (depth-first exploration)
- For chess, b ≈ 35, m ≈100 for "reasonable" games
   → exact solution completely infeasible

## Problem

- Number of game states exponential in number of moves
- Can cut in half with pruning
  - Still exponential
  - Get rid of braches that can't influence final decision











## **Properties of α-β**

- Pruning does not affect final result
- Good move ordering improves effectiveness of pruning
- With "perfect ordering," time complexity = O(b<sup>m/2</sup>)
   → doubles depth of search
- A simple example of the value of reasoning about which computations are relevant (a form of metareasoning)

#### Why is it called α-β?

- α is the value of the best (i.e., highest-value) choice found so far at any choice point along the path for max
- If v is worse than α, max will avoid it
   → prune that branch
- Define β similarly for min



## The α-β algorithm

function ALPHA-BETA-SEARCH(state) returns an action inputs: state, current state in game

 $v \leftarrow \text{MAX-VALUE}(state, -\infty, +\infty)$ return the *action* in SUCCESSORS(*state*) with value v

function MAX-VALUE(state,  $\alpha$ ,  $\beta$ ) returns a utility value

inputs: *state*, current state in game

lpha, the value of the best alternative for  $_{
m MAX}$  along the path to state

 $\beta,$  the value of the best alternative for  $_{\rm MIN}$  along the path to state

if TERMINAL-TEST(state) then return UTILITY(state)

 $v \leftarrow -\infty$ 

for a, s in SUCCESSORS(state) do

```
v \leftarrow Max(v, MIN-VALUE(s, \alpha, \beta))
```

```
 {\rm if} \ v \geq \beta \ {\rm then} \ {\rm return} \ v \\
```

```
\alpha \leftarrow Max(\alpha, v)
```

return v

## The α-β algorithm

```
function MIN-VALUE(state, \alpha, \beta) returns a utility value
```

inputs: *state*, current state in game

 $\alpha$ , the value of the best alternative for MAX along the path to *state* 

eta, the value of the best alternative for  $_{
m MIN}$  along the path to state

if TERMINAL-TEST(*state*) then return UTILITY(*state*)

 $v \leftarrow +\infty$ 

```
for a, s in SUCCESSORS(state) do
```

```
v \leftarrow \operatorname{MIN}(v, \operatorname{MAX-VALUE}(s, \alpha, \beta))
```

```
if v \leq \alpha then return v
```

```
\beta \leftarrow \operatorname{Min}(\beta, v)
```

return v

#### **Resource limits**

Suppose we have 100 secs, explore  $10^4$  nodes/sec  $\rightarrow 10^6$  nodes per move

Standard approach:

- cutoff test:
  - e.g., depth limit
- evaluation function
  - = estimated desirability of position
  - Usually a heuristic function

#### **Evaluation functions**

- For chess, typically linear weighted sum of features
   Eval(s) = w<sub>1</sub> f<sub>1</sub>(s) + w<sub>2</sub> f<sub>2</sub>(s) + ... + w<sub>n</sub> f<sub>n</sub>(s)
- e.g.,  $w_1 = 9$  with
  - $f_1(s) = (number of white queens) (number of black queens), etc.$

## **Cutting off search**

MinimaxCutoff is identical to MinimaxValue except

- 1. Terminal? is replaced by Cutoff?
- 2. Utility is replaced by Eval

Does it work in practice?  $b^m = 10^6, b=35 \rightarrow m=4$ 

4-ply lookahead is a hopeless chess player!

- 4-ply  $\approx$  human novice
- 8-ply  $\approx$  typical PC, human master
- 12-ply ≈ Deep Blue, Kasparov

# Deterministic games in practice

- Checkers: Chinook ended 40-year-reign of human world champion Marion Tinsley in 1994. Used a precomputed endgame database defining perfect play for all positions involving 8 or fewer pieces on the board, a total of 444 billion positions.
- Chess: Deep Blue defeated human world champion Garry Kasparov in a six-game match in 1997. Deep Blue searches 200 million positions per second, uses very sophisticated evaluation, and undisclosed methods for extending some lines of search up to 40 ply.
- Othello: human champions refuse to compete against computers, who are too good.
- Go: human champions refuse to compete against computers, who are too bad. In go, b > 300, so most programs use pattern knowledge bases to suggest plausible moves.

## Summary

- Games are fun to work on!
- They illustrate several important points
   about AI
- perfection is unattainable → must approximate
- good idea to think about what to think about