

CS 188: Artificial Intelligence

Spring 2010

Lecture 7: Minimax and Alpha-Beta Search

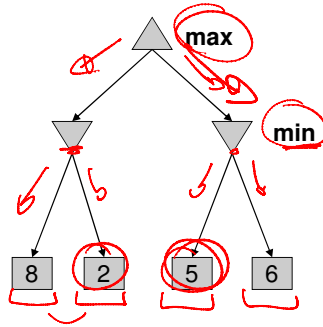
2/9/2010

Pieter Abbeel – UC Berkeley
Many slides adapted from Dan Klein

Announcements


- Section format
- Written 2: due Thursday

Simple two-player game example



3


Tic-tac-toe Game Tree



4

Deterministic Two-Player

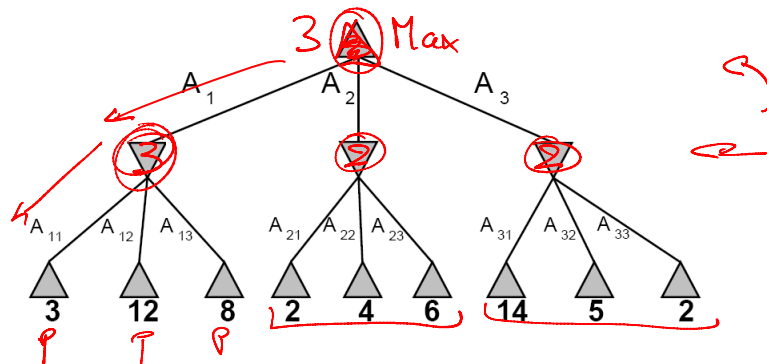
- E.g. tic-tac-toe, chess, checkers
- **Zero-sum games**
 - One player maximizes result
 - The other minimizes result
- **Minimax search**
 - A state-space search tree
 - Players alternate
 - Each layer, or ply, consists of a round of moves*
 - Choose move to position with highest **minimax value** = best achievable utility against best play



* Slightly different from the book definition

7

Minimax Example



8

Minimax Search

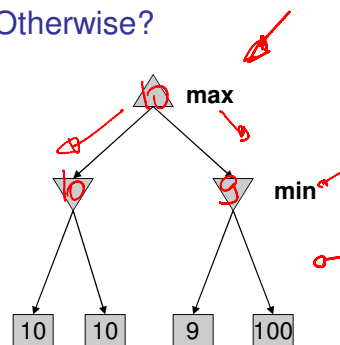
```
function MAX-VALUE(state) returns a utility value
  if TERMINAL-TEST(state) then return UTILITY(state)
  v ← -∞
  for a, s in SUCCESSORS(state) do v ← MAX(v, MIN-VALUE(s))
  return v
```

```
function MIN-VALUE(state) returns a utility value
  if TERMINAL-TEST(state) then return UTILITY(state)
  v ← ∞
  for a, s in SUCCESSORS(state) do v ← MIN(v, MAX-VALUE(s))
  return v
```

9

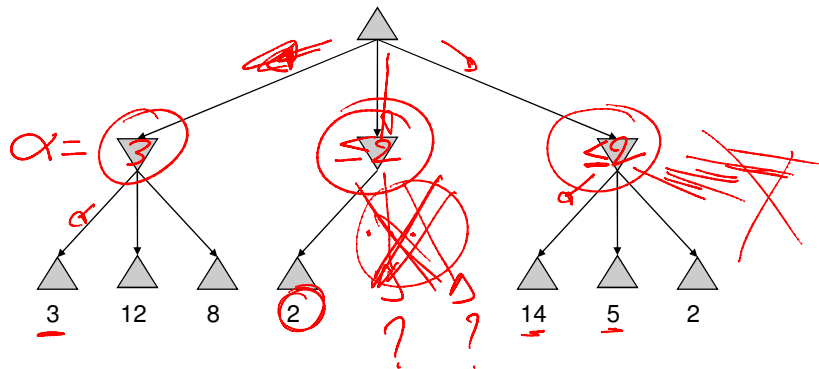
Minimax Properties

- Optimal against a perfect player. Otherwise?
- Time complexity?
 - $O(b^m)$
- Space complexity?
 - $O(bm)$
- For chess, $b \approx 35, m \approx 100$
 - Exact solution is completely infeasible
 - But, do we need to explore the whole tree?



11

Pruning



14

Alpha-Beta Pruning

General configuration


- We're computing the MIN-VALUE at n
- We're looping over n 's children
- n 's value estimate is dropping
- a is the best value that MAX can get at any choice point along the current path
- If n becomes worse than a , MAX will avoid it, so can stop considering n 's other children
- Define b similarly for MIN

MAX

MIN

MAX

MIN



16

Alpha-Beta Pseudocode

```


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  $v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s))$ 
  return v
    
```

minimax

```

function MAX-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
   $\beta$ , the value of the best alternative for MIN along the path to state
  if TERMINAL-TEST(state) then return UTILITY(state)
   $v \leftarrow -\infty$ 
  for a, s in SUCCESSORS(state) do
     $v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s, \alpha, \beta))$ 
    if  $v \geq \beta$  then return v
     $\alpha \leftarrow \text{MAX}(\alpha, v)$ 
  return v
    
```

$\alpha - \beta$

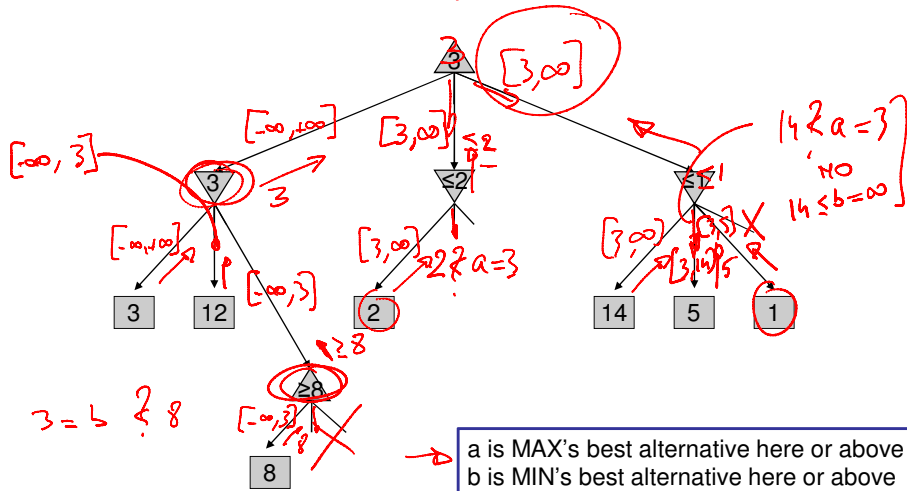


Alpha-Beta Pruning Properties

- This pruning has no effect on final result at the root
- Values of intermediate nodes might be wrong!
- Good child ordering improves effectiveness of pruning
- With “perfect ordering”:
 - Time complexity drops to $O(b^{m/2})$
 - Doubles solvable depth!
 - Full search of, e.g. chess, is still hopeless...
- This is a simple example of **metareasoning** (computing about what to compute)

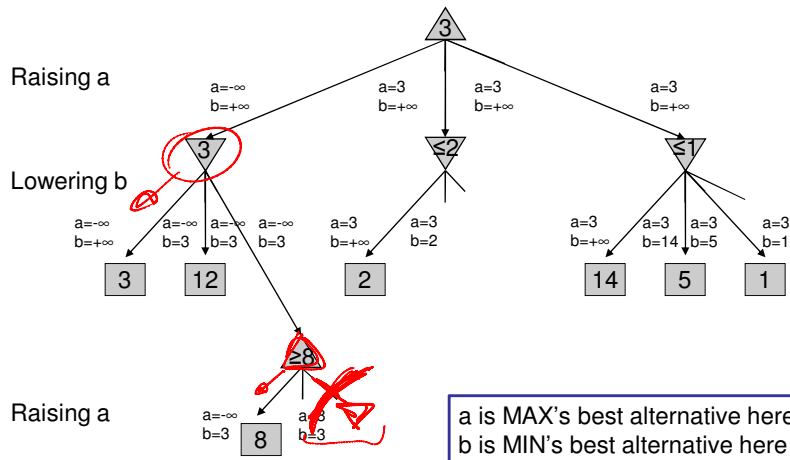
Alpha-Beta Pruning Example

Starting a/b $[-\infty, \infty]$



Alpha-Beta Pruning Example

Starting a/b $a=-\infty, b=+\infty$




Resource Limits

- Cannot search to leaves
- Depth-limited search
 - Instead, search a limited depth of tree
 - Replace terminal utilities with an eval function for non-terminal positions

- ☞ Guarantee of optimal play is gone
- ☞ More plies makes a BIG difference


- Example:
 - Suppose we have 100 seconds, can explore 10K nodes / sec
 - So can check 1M nodes per move
 - α - β reaches about depth 8 – decent chess program



23

Evaluation Functions

- Function which scores non-terminals



- Ideal function: returns the utility of the position
- In practice: typically weighted linear sum of features:

$$\hookrightarrow Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$

- e.g. $f_1(s) = (\text{num white queens} - \text{num black queens})$, etc.

24


utility = $x \rightarrow$ for max
utility (MIN) = $-x$

28

Non-Zero-Sum Games

- Similar to minimax:

- Utilities are now tuples
- Each player maximizes their own entry at each node
- Propagate (or back up) nodes from children



29