

Sizes of Games

Minimax(pos)

b = branching factor = avg moves possible
 If pos is terminal, return value determined by rules \rightarrow check who won
 +1: P1
 0: draw
 -1: P2

d = depth = avg length

Else if pos is P1's turn then return $\max_{pos \rightarrow pos'} MM(pos')$

pos $\approx b^d$

Else return $\min_{pos \rightarrow pos'} MM(pos')$

Tic-Tac-Toe	very rough est of # pos $\leq 3^9 \approx 20000$
Mancala	$\leq \binom{49}{36} \approx 262$ billion
2-player Yahtzee	$\leq 4 \cdot 10^{18}$
Checkers	$\leq 10^{20}$
Chess	$\leq 10^{43}$
Go	$\leq 10^{172}$

What to do with games of high complexity?

heuristics - estimate of position value

Ex: checkers % of remaining pieces that are black (scaled to (-1,1)) $\Delta 8$
 04 heuristic = $1/3$

chess assign value to each type of piece
 pawn = 1
 knight, bishop = 3
 rook = 5
 queen = 9
 % of value of remaining pieces belonging to white (scaled)

Minimax(pos, h, depth) \rightarrow depth bound; higher bounds yield results closer to unbounded MM

If pos is terminal, return value(pos)

If depth == 0, return h(pos)

Else if pos is P1's turn then return $\max_{pos \rightarrow pos'} MM(pos', h, depth)$

...

Else if pos is P1's turn then return $\max_{pos \rightarrow pos'} MM(pos', h, depth-1)$
 Else return $\min_{pos \rightarrow pos'} MM(pos', h, depth-1)$

Negamax(pos, h, depth, sign)

If pos is terminal, return $value(pos) \cdot sign$

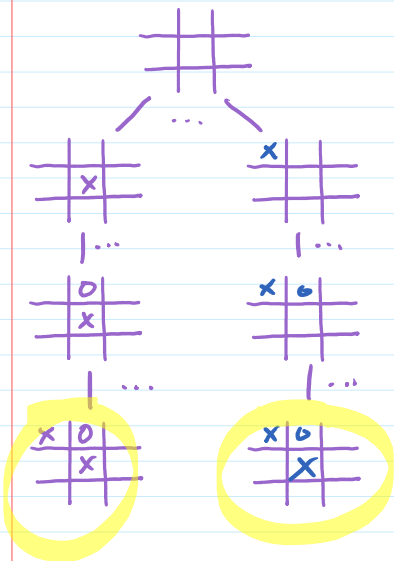
If depth == 0, return $h(pos) \cdot sign$

Else return $\max_{pos \rightarrow pos'} -MM(pos', h, depth-1, -sign)$

Iterative Deepening - to get value from MM in set time w/ depth high

depth $\leftarrow 2$
 while not out of time
 do MM(pos, h, depth)
 depth \leftarrow depth + 1
 return value, best move returned by last call to MM that finished

Transposition Table (memo)



multiple paths to same pos

If pos is terminal, return $value(pos)$ to un sound

If depth == 0, return $h(pos)$

if pos in TT return TT[pos]


Else if pos is P1's turn then return $\max_{TT(pos)=pos \rightarrow pos'} MM(pos', h, depth-1, TT)$
 return TT[pos]

Else return $\min_{pos \rightarrow pos'} MM(pos', h, depth-1, TT)$
 return TT[pos]

(memory intensive)

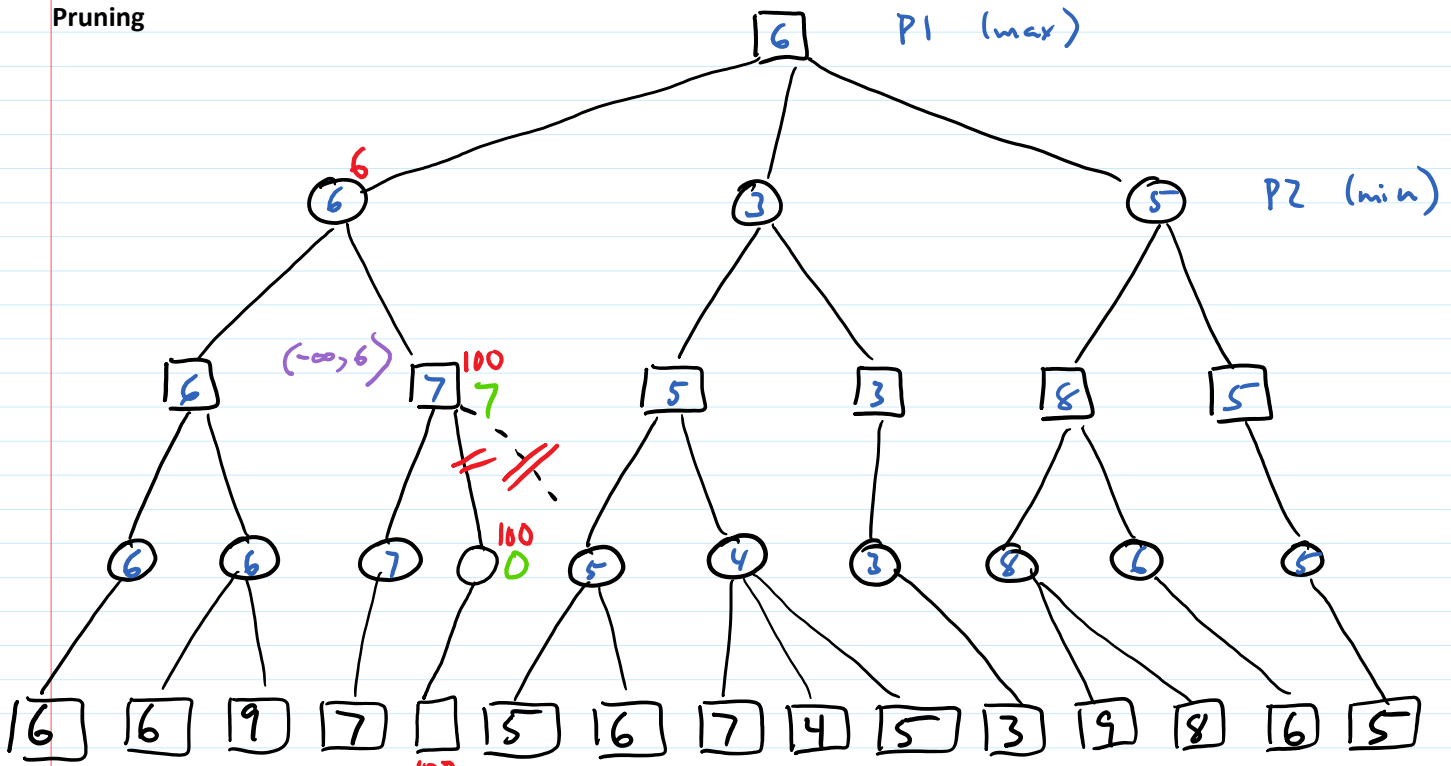
↳ replacement policy?

Minimax Example



Modified example from http://en.wikipedia.org/wiki/Alpha%E2%80%93beta_pruning

Pruning



↑
value here does not affect value of root
so don't search this branch!

Modified example from http://en.wikipedia.org/wiki/Alpha%E2%80%93beta_pruning

Alpha-Beta Pruning

bounds - if MM value in (α, β) , want exact value
 if $\leq \alpha$ or $\geq \beta$, don't need exact value

Alpha-Beta $(p, \alpha, \beta, \text{depth}, h)$ returns

value from MM using same depth, h \leftarrow $MM(p)$

$\left. \begin{array}{l} \text{value}(p) \\ h(p) \\ \text{if } p \text{ terminal} \\ \text{if } \text{depth} = 0 \\ \text{if } \alpha \leq MM(p) \leq \beta \\ \text{if } MM(p) \geq \beta \\ \text{if } MM(p) \leq \alpha \end{array} \right\} \text{postconditions}$

lower bound on $MM(p)$ some a s.t. $\beta \leq a \leq MM(p)$ if $MM(p) \geq \beta$
 upper bound $MM(p)$ some b s.t. $MM(p) \leq b \leq \alpha$ if $MM(p) \leq \alpha$

if $\text{depth} = 0$ then return heuristic(p)

if p is terminal then return value(p)

if p is a max position

$a \leftarrow -\infty$

for each position p' reachable in one move from p while $\alpha < \beta$

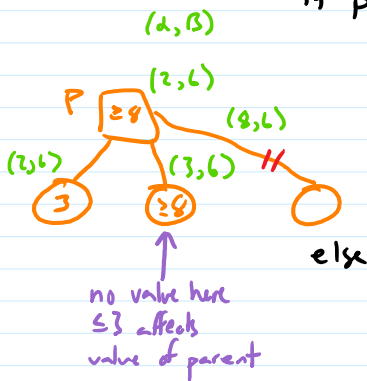
$a \leftarrow \max(a, AB(p', \alpha, \beta, \text{depth}-1, h))$

$\alpha \leftarrow \max(\alpha, a)$

return a fail-soft (fail-hard returns α after pruning)

$(\alpha, \beta) = (2, 6)$

$1^{\text{st}} p' = 3, 2^{\text{nd}} p' = 8$



$b \leftarrow \infty$

for each position p' reachable in one move from p while $\alpha < \beta$

$b \leftarrow \min(b, AB(p', \alpha, \beta, \text{depth}-1, h))$

$\beta \leftarrow \min(\beta, b)$

return b


start with call $AB(\text{curr}, -\infty, \infty, \text{depth}, h)$ on current pos curr

postconditions guarantee returned value is $= MM(\text{curr}, \text{depth}, h)$

if could guess value, could start w/narrower window $AB(\text{curr}, 2, 10, \text{depth}, h)$ returns ∞

$AB(\text{curr}, 0, 10, \text{depth}, h)$ and widen window if guess was wrong


⋮




child 3

child 4

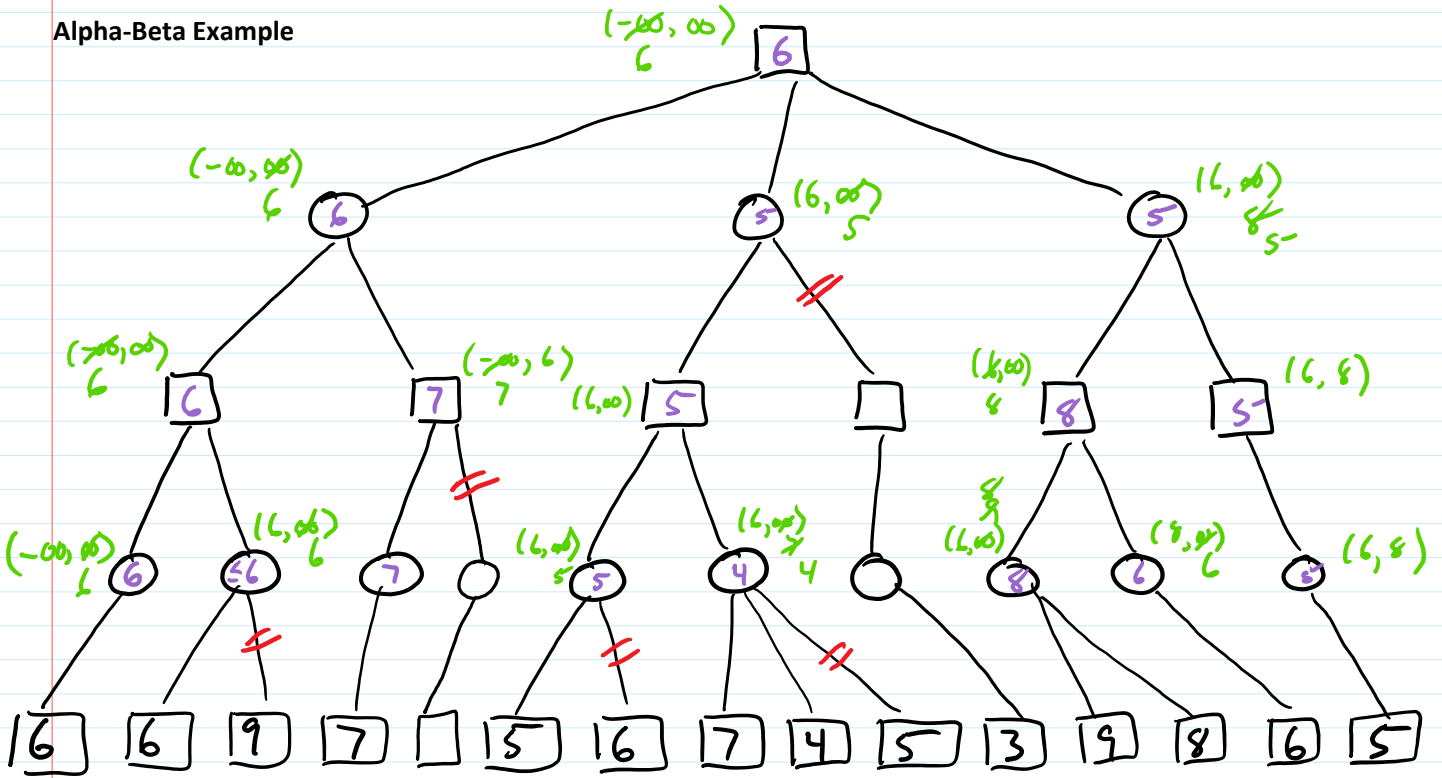
(if only 2 children)



child 5

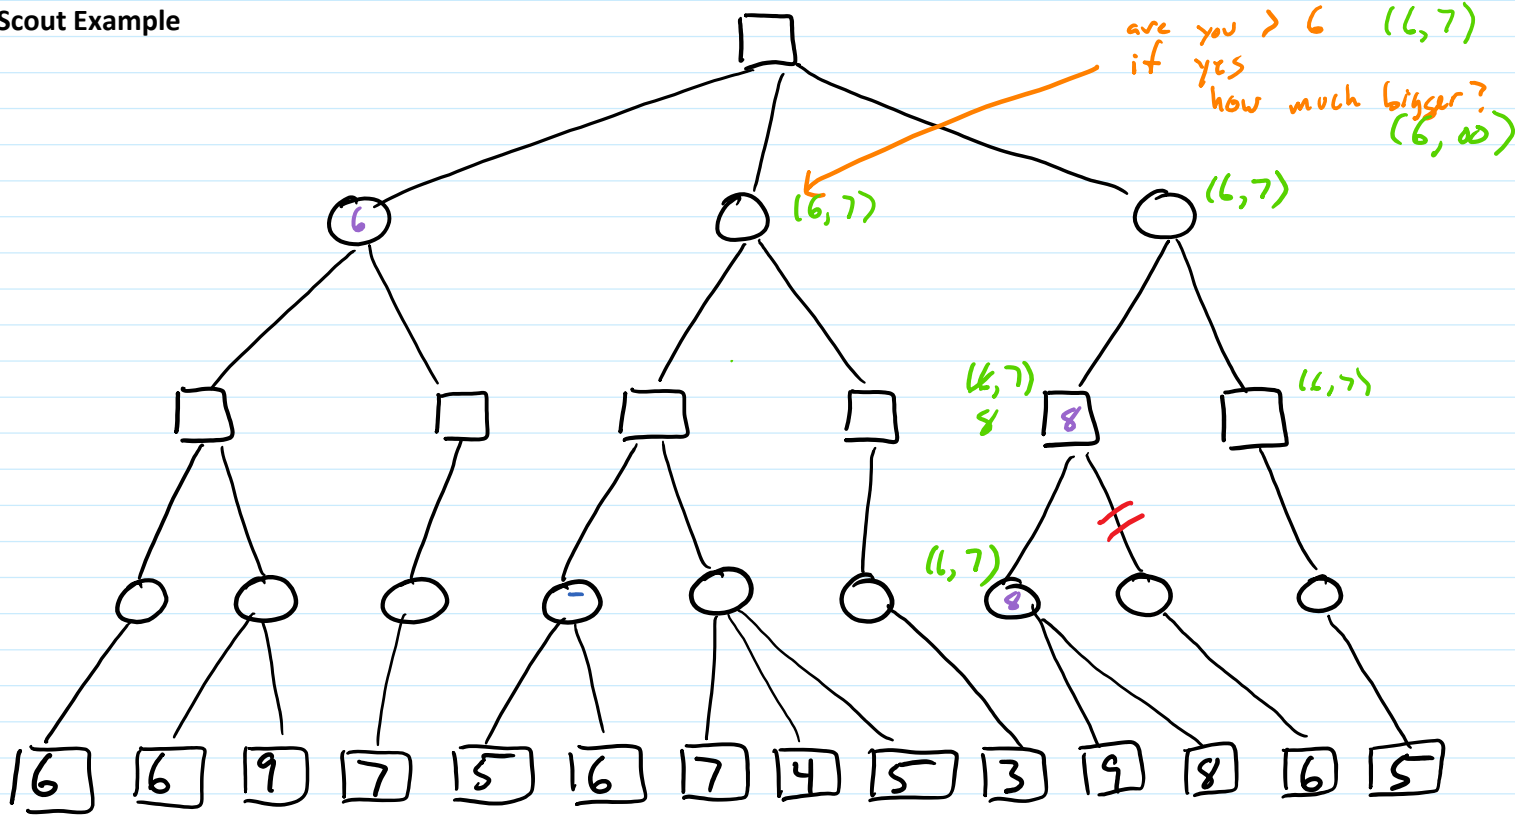


Alpha-Beta Example



Modified example from http://en.wikipedia.org/wiki/Alpha-Beta_pruning

Scout Example



Scout ($p, \alpha, \beta, \text{depth}, h$)

if p is terminal then return $\text{value}(p)$

if $\text{depth} = 0$ then return $\text{heuristic}(p)$

if p is a max position

$\text{best} \leftarrow -\infty$

for each reachable position p'_n and while $\alpha < \beta$

if p' is first pos
 $\text{score} \leftarrow \text{Scout}(p', \alpha, \beta, \text{depth}-1, h)$

else

$\text{score} \leftarrow \text{AB}(p', \alpha, \alpha+1, \text{depth}-1, h)$

if $\alpha < \text{score} < \beta$

$\text{score} \leftarrow \text{Scout}(p', \text{score}, \beta, \text{depth}-1, h)$

$\text{best} \leftarrow \max(\text{best}, \text{score})$

$\alpha \leftarrow \max(\text{best}, \alpha)$

return best

(or fail-hard version)

else

⋮ min position; symmetric
⋮

already passing
null window, so
no reason for Scout

null window (assuming integers)

already know $\text{value}(pos) \geq \text{score}$