

CSC384
Game Tree Search
Part 2

Bahar Aameri & Sonya Allin

Winter 2020

These slides are drawn from or inspired by a multitude of sources including :

Faheim Bacchus
Sheila McIlraith
Andrew Moore
Hojjat Ghaderi
Craig Boutillier
Jurgen Strum
Shaul Markovitch


- There are ways to **avoid** examining the **entire** tree to make correct Minimax decision.
- When using depth-first search of a game tree:
 - After generating value for only **some of s 's children** we can prove that we **never reach s** in a Minimax strategy.
 - So we need NOT generate or evaluate any further children of s . These other children can be **pruned**.

Cutting Max Nodes (Alpha Cuts)

At a **Max** node s :

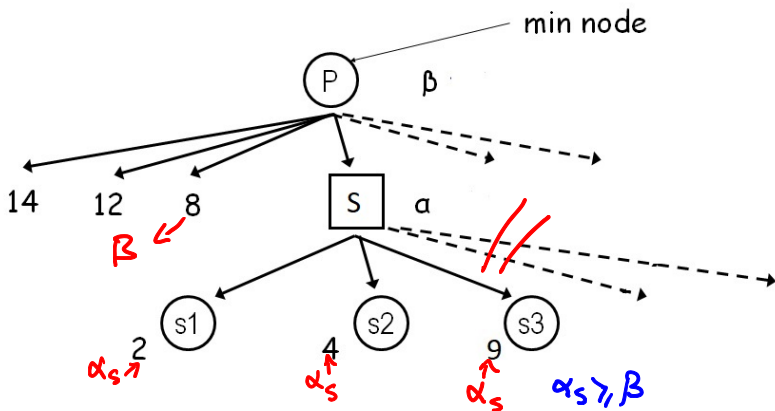
α_s : The **highest value of s 's children** examined so far (changes as children of s are examined).

β : The **best option for MIN** (i.e., lowest value) found so far (fixed as children of s are examined);



Cutting Max Nodes (Alpha Cuts)

If α_s becomes **greater than or equal** to β , we can stop expanding the children of s :
Min will **never choose** to move from s 's parent to s since it would choose one of s 's **lower valued siblings**.




Cutting Min Nodes (Beta Cuts)

At a **Min** node s :


α : The **best option for MAX** (i.e., highest value) found so far (fixed as children of s are examined).


β_s : The **lowest value of s 's children** examined so far (changes as children of s are examined);



Cutting Min Nodes (Beta Cuts)

- If α becomes **greater than or equal to** β_s , we can stop expanding the children of s :
Max will **never choose** to move from s 's parent to s since it would choose one of s 's higher value siblings.





Alpha-Beta Pruning

α : **Best** already explored option along the path to the root for **MAX**.

β : **Best** already explored option along the path to the root for **MIN**.


Alpha-Beta Pruning:


- Set initial values: $\alpha = -\infty$ and $\beta = \infty$
- While backing the utility values up the tree, identify α and β for each node.
- At every node s , if $\alpha \geq \beta$, **prune** (remaining) children of s .

α -cuts: Pruning of **MAX** nodes.

β -cuts: Pruning of **MIN** nodes.

Alpha-Beta Pruning – Example





Alpha-Beta Pruning Implementation

```
def AlphaBeta(s, Player, alpha, beta):
    // Return Utility of state s given that Player is MIN or MAX
    1. If s is TERMINAL
    2.     Return U(s)    # Return terminal states utility
    3. ChildList = s.Successors(Player)
    4. If Player == MAX
    5.     ut_val = -infinity
    6.     for c in ChildList
    7.         ut_val = max(ut_val, AlphaBeta(c, MIN, alpha, beta))
    8.         If alpha < ut_val
    9.             alpha = ut_val
    10.            If beta <= alpha: break
    11.     return ut_val
    12. Else # Player is MIN
    13.     ut_val = infinity
    14.     for c in ChildList
    15.         ut_val = min(ut_val, AlphaBeta(c, MAX, alpha, beta))
    16.         If beta > ut_val
    17.             beta = ut_val
    18.            If beta <= alpha: break
    19.     return ut_val
```

Ordering of Moves

- For **MIN** nodes the best pruning occurs if the **best move for MIN** (child yielding lowest value) is **explored first**.
- For **MAX** nodes the best pruning occurs if the **best move for MAX** (child yielding highest value) is **explored first**.
- We don't know which child has highest or lowest value without doing all of the work! But we can use **heuristics** to estimate the value, and then choose the child with highest (lowest) heuristic value.

Effectiveness of Alpha-Beta Pruning

- With no pruning, $\mathcal{O}(b^d)$ nodes are explored, which makes the run time of a search with pruning the same as plain Minimax.

If, however, the move **ordering** for the search is **optimal** (meaning the best moves are searched first), the number of nodes we need to search using alpha beta pruning is $\mathcal{O}(b^{d/2})$.

- In Deep Blue, they found that **alpha-beta pruning** meant the average branching factor at each node was about **6 instead of 35**.

- **Real** games are **too large** to enumerate tree.

Example:

- Chess branching factor is roughly 35.
 - Depth 10 tree: 2,700,000,000,000,000 nodes
 - Even Alpha-Beta pruning won't help here!
- We must **limit depth** of search tree:
 - Must stop the search at some **non-terminal nodes**.
 - We must make **heuristic estimates** about the values of the non-terminal positions where we terminate the search.
 - These heuristics are often called **evaluation functions**.
 - Evaluation functions are often **learned**.

Examples of Heuristics in Games:

- **Tic Tac Toe:**

$h(n) = [\text{\# of 3 lengths that are left open for player A}] - [\text{\# of 3 lengths that are left open for player B}]$.

- **Chess:** Alan Turing's function

$h(n) = A(n)/B(n)$ where $A(n)$ is the sum of the point value for player A's pieces and $B(n)$ is the sum for player B.

- Many evaluation functions can be specified as a **weighted sum of features**:

$h(n) = w_1 \times \text{feature}_1(n) + w_2 \times \text{feature}_2(n) + \dots w_i \times \text{feature}_i(n)$.

The weights can be **learnt**.

An Aside on Large Search Problems

- Inability to expand tree to terminal nodes is relevant even in **standard search**: Often we can't expect the search to reach a goal by expanding full frontier.
- **Real-time (or online) Search**: We **limit our look-ahead**, and make **moves before we actually know the true path** to the goal.
- In real-time search, we use the **heuristic function** not just to guide our search, but also to commit to **moves** we actually make.
In general, guarantees of **optimality are lost**, but we **reduce computational/memory expense** dramatically.

