CS540 Introduction to Artificial Intelligence

Lecture 21

Young Wu

Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

August 9, 2022

Pirate Game Example

Quiz

- 5 pirates got 100 gold coins. Each pirate takes a turn to propose how to divide the coins, and all pirates who are still alive will vote whether to accept the proposal or reject the proposal, kill the pirate, and continue to the next round. Use strict majority rule for the vote, and use the assumption that if a pirate is indifferent, he or she will vote reject with probability 50 percent.
- How will the first pirate propose?
- $A : (0, 0, 0, 0, 100)$
- $B : (20, 20, 20, 20, 20)$
- $C : (94, 0, 1, 2, 3)$
- $D : (97, 0, 1, 0, 2)$
- $E : (98, 0, 1, 0, 1)$

Q1

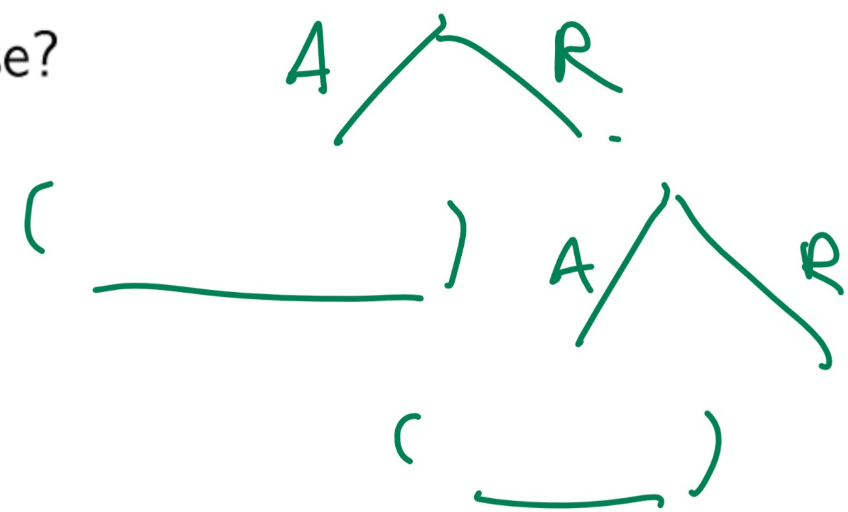

Pirate Game Example Diagram

Quiz

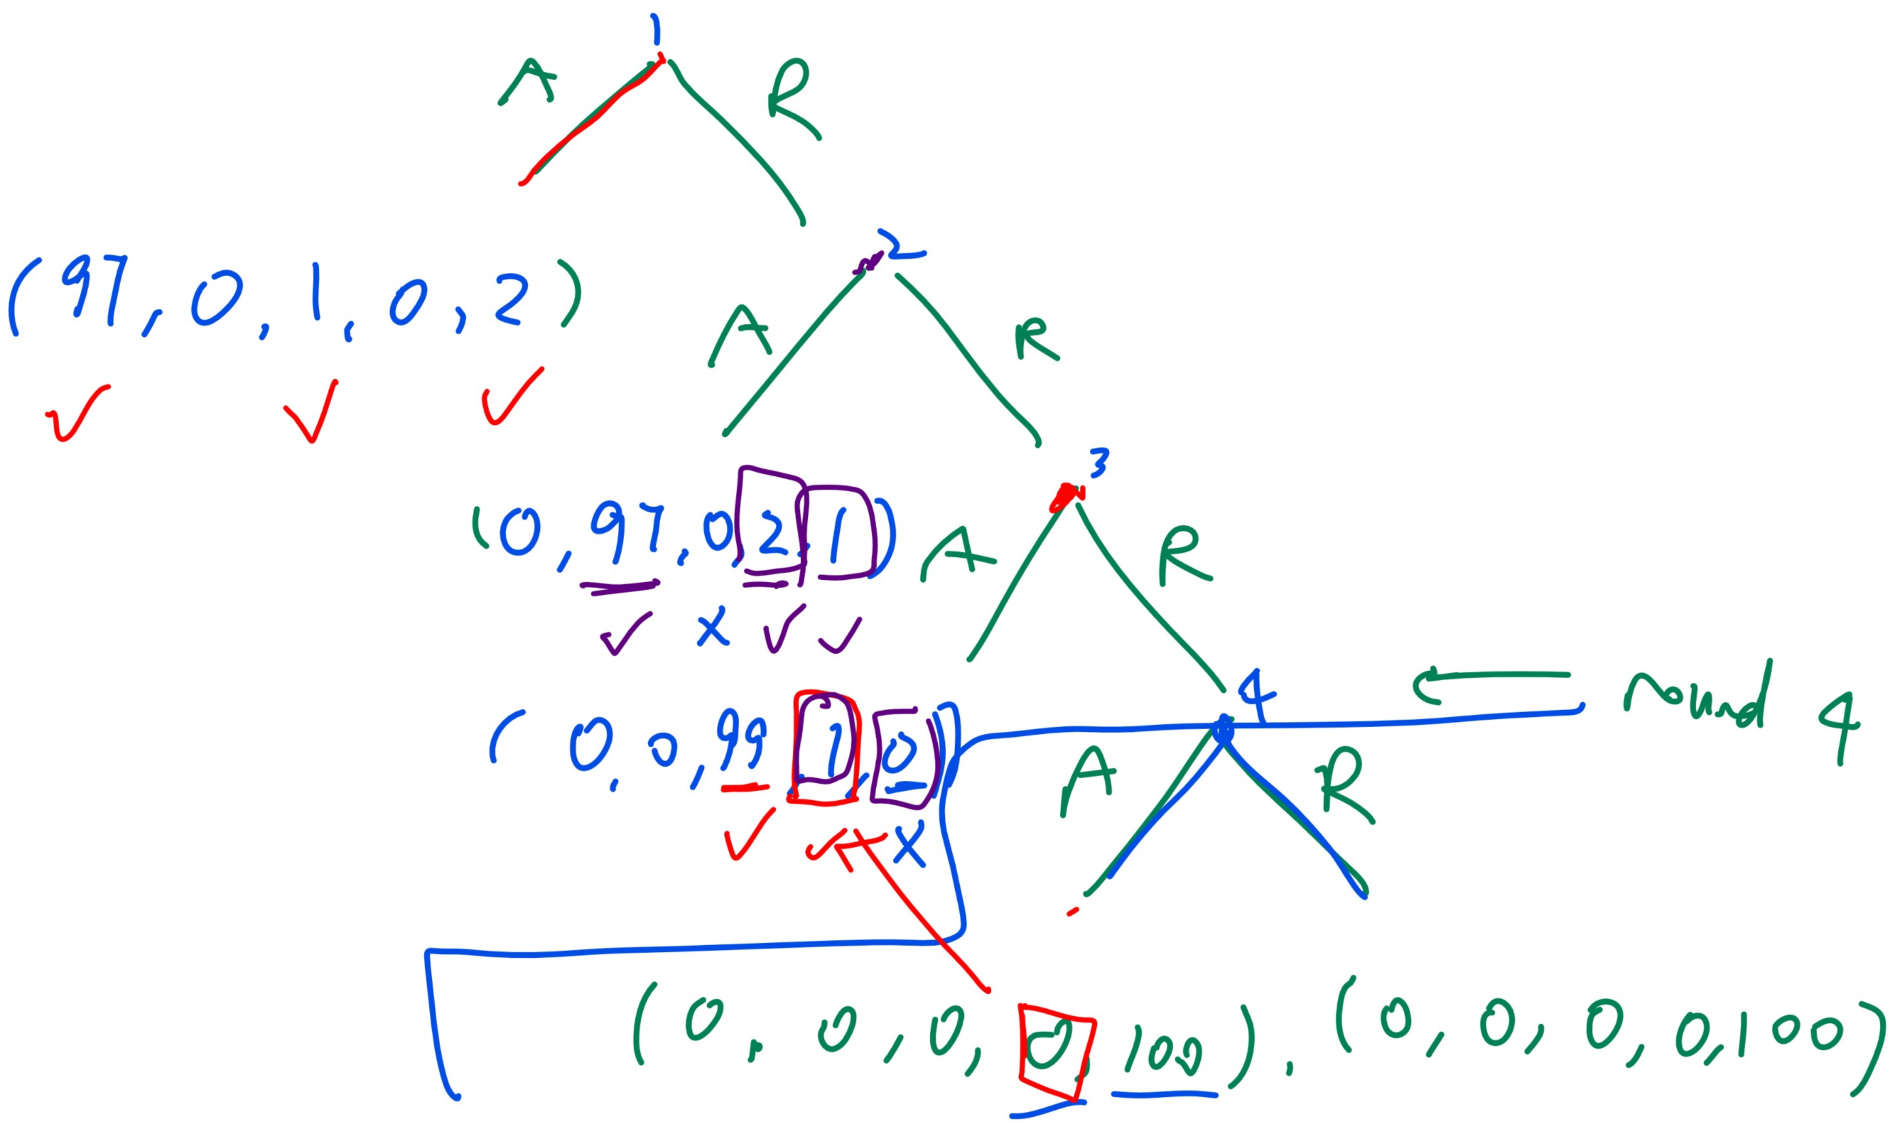

Summary

Discussion

- Search:

1. Uninformed.
2. Informed.
3. Local Search.
4. Adversarial Search: Sequential move games.
5. Adversarial Search: Simultaneous move games.

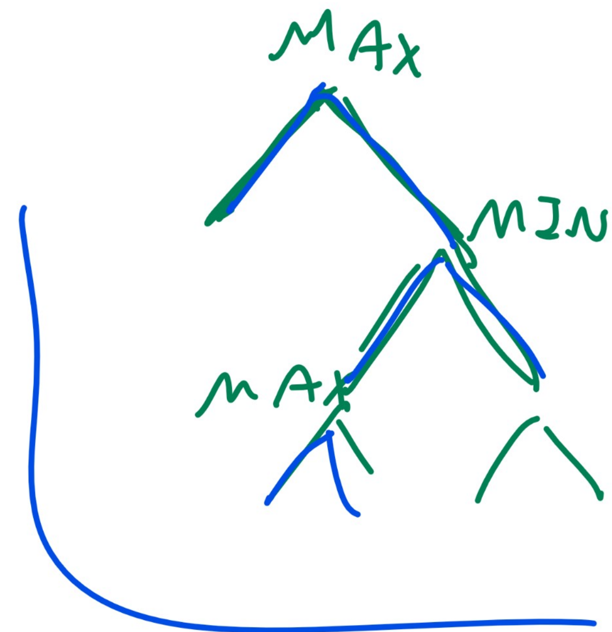

Tic Tac Toe Example

Motivation

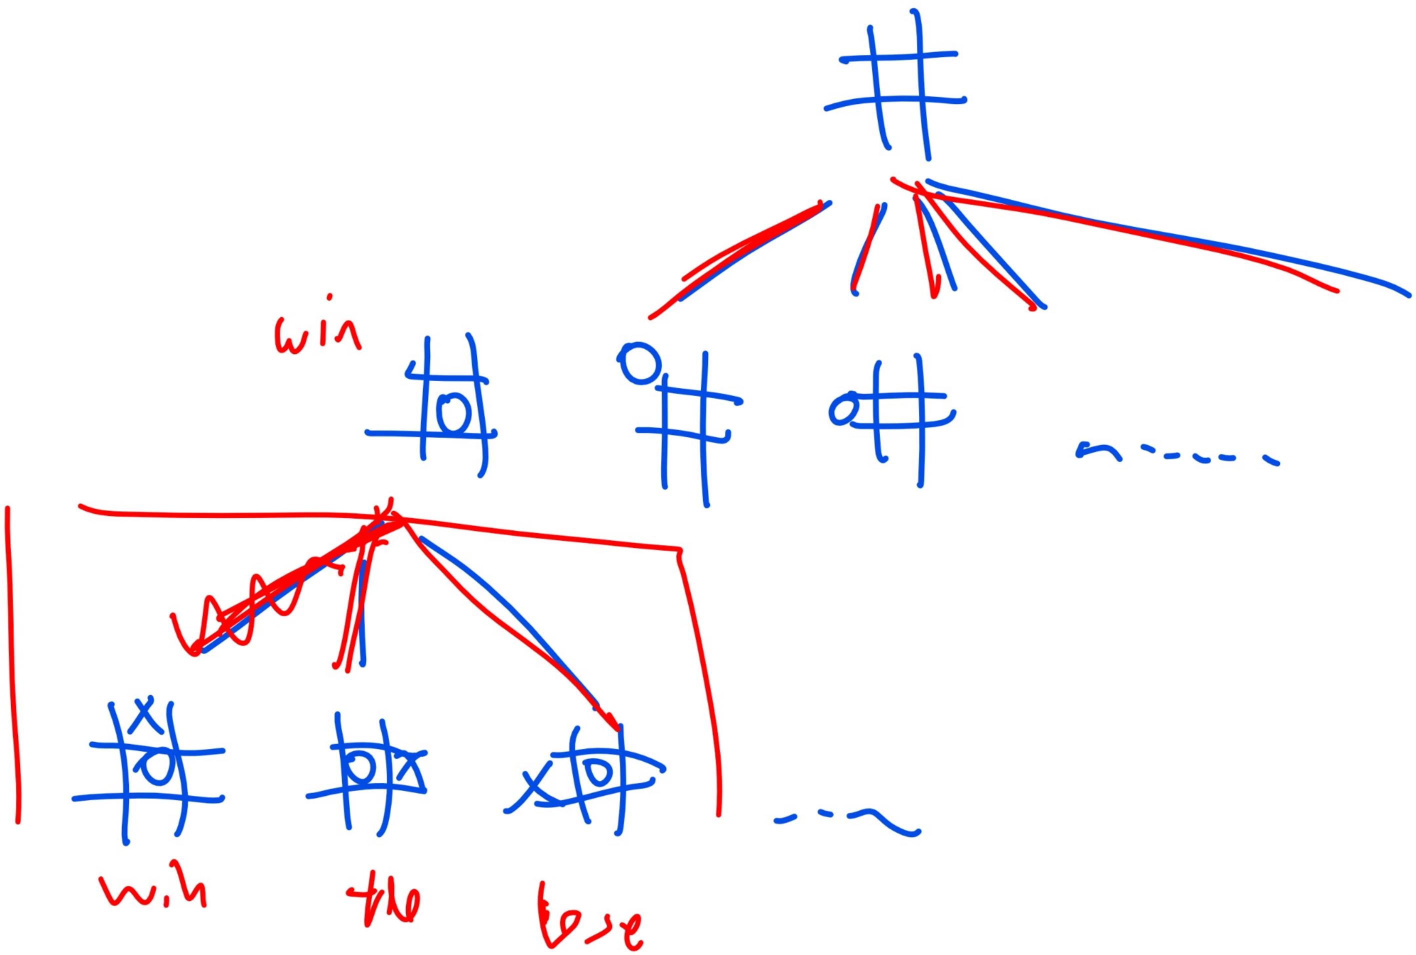

Nim Game Example

Motivation

- Ten objects. Pick 1 or 2 each time. Pick the last one to win.
- A : Pick 1.
- B : Pick 2.
- C, D, E : Don't choose.

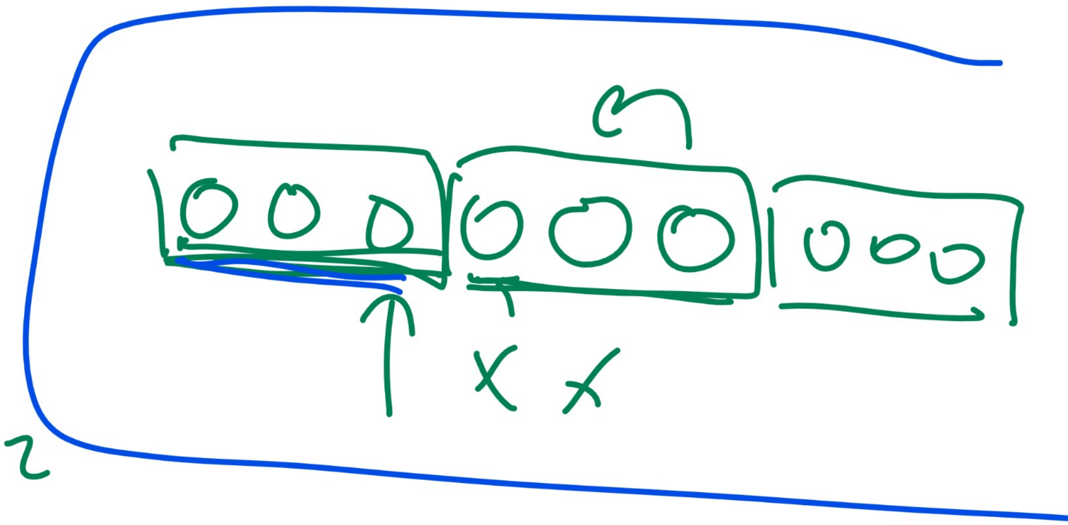

Minimax Algorithm

Description

- Use DFS on the game tree.

Minimax Example

Quiz

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. What is the value of the game?
- $A : 1$
- $B : 2$
- $C : 4$
- $D : 8$
- $E : 16$

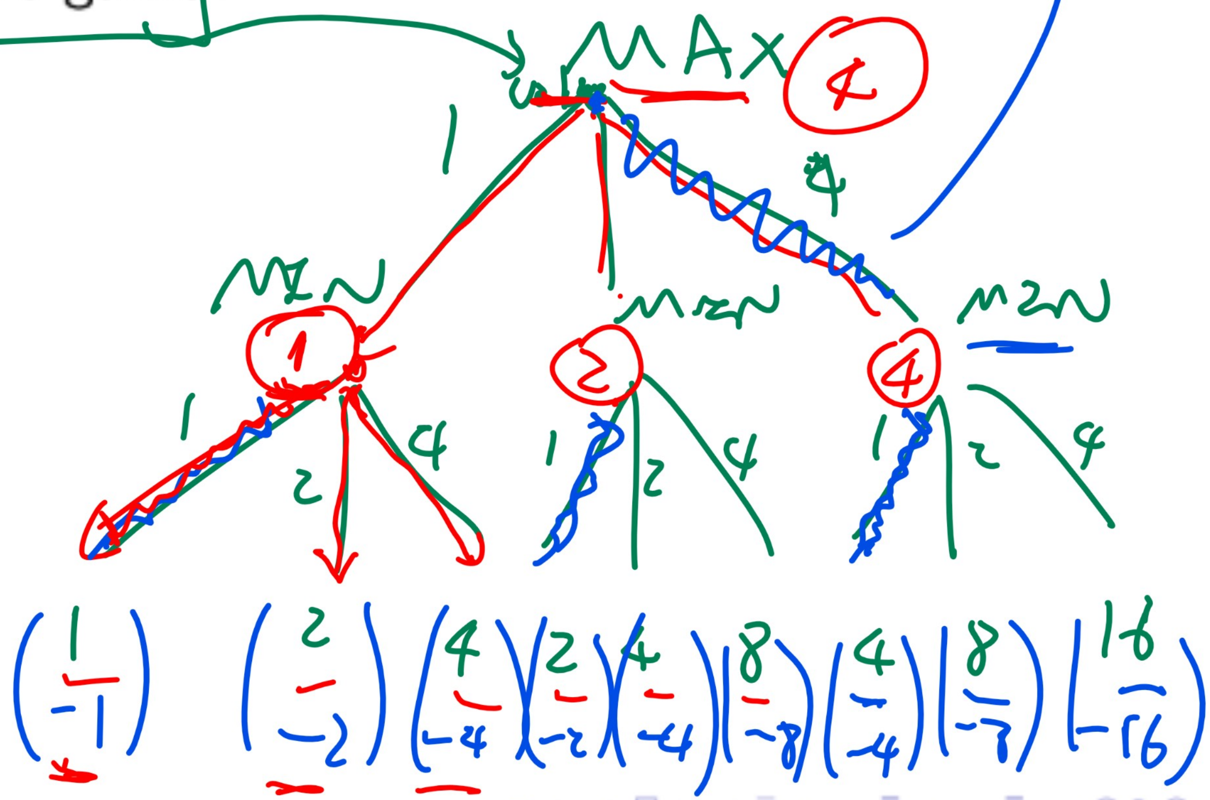

Minimax Performance

Discussion

- The time and space complexity is the same as DFS. Note that $D = d$ is the maximum depth of the terminal states.

$$T = 1 + b + b^2 + ... + b^d$$

$$S = (b - 1) \cdot d$$

Non-deterministic Game

Discussion

- For non-deterministic games in which chance can make a move (dice roll or coin flip), use expected reward or cost instead.
- The algorithm is also called expectiminimax.

Game Tree with Chance Example 1

Quiz

- Max can pick L or R. If Max picks L, Chance picks L with probability 0.3 and R with probability 0.7. If Chance picks L, Min picks L to get 3, R to get 2, and if Chance picks R, Min gets 7. If Max picks R, Min picks L to get -1 and R to get 2. What is the value of the game?

$0.3 \cdot 2 + 0.7 \cdot 7 = 5.5$

0.6 4.9

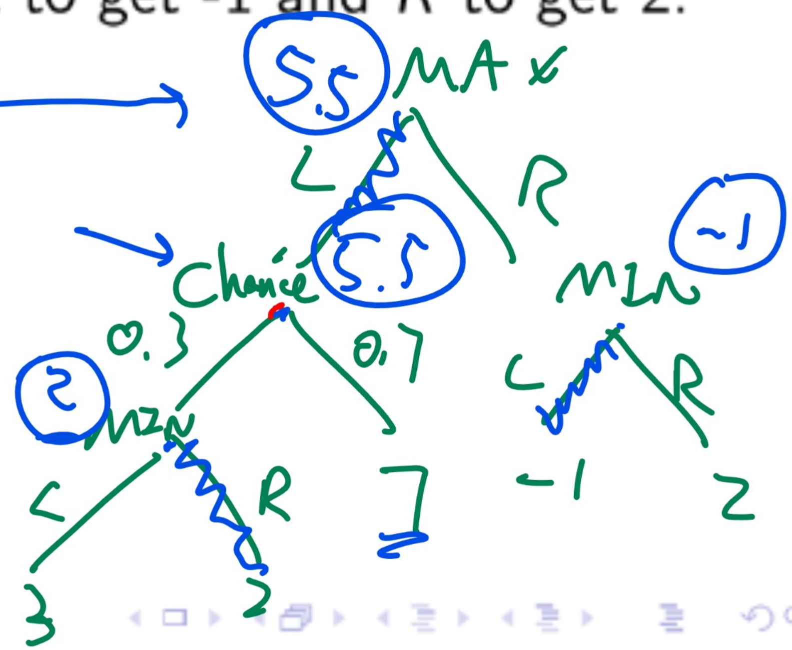

Game Tree with Chance Example 1 Diagram

Quiz

Pruning

Motivation

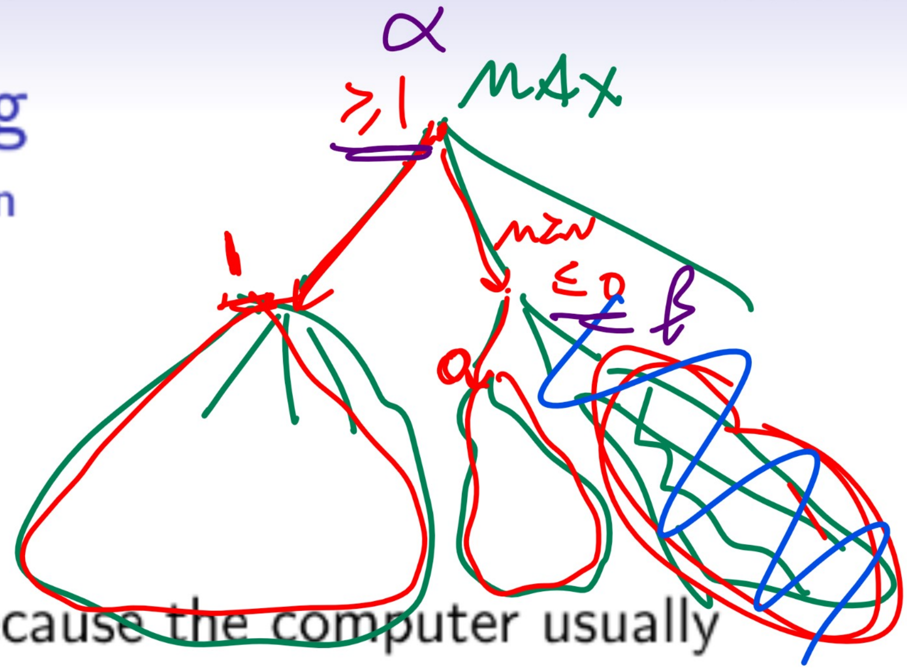

- Time complexity is a problem because the computer usually has a limited amount of time to "think" and make a move.
- It is possible to reduce the time complexity by removing the branches that will not lead the current player to win. It is called the Alpha-Beta pruning.

prune

max does not care about subgames → always get ≤ 0.

Alpha Beta Pruning

Description

- During DFS, keep track of both α and β for each vertex.
- Prune the subtree with $\alpha \geqslant \beta$.

Alpha Beta Example 1

Quiz

Alpha Beta Example 1 Continued

Quiz

Alpha Beta Example 2

Quiz

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. Alpha-Beta pruning is used. What is the number of branches (states) that can be pruned if the actions with larger labels are searched first?
- $A: 0$
- $B: 1$
- $C: 2$
- $D: 3$
- $E: 4$

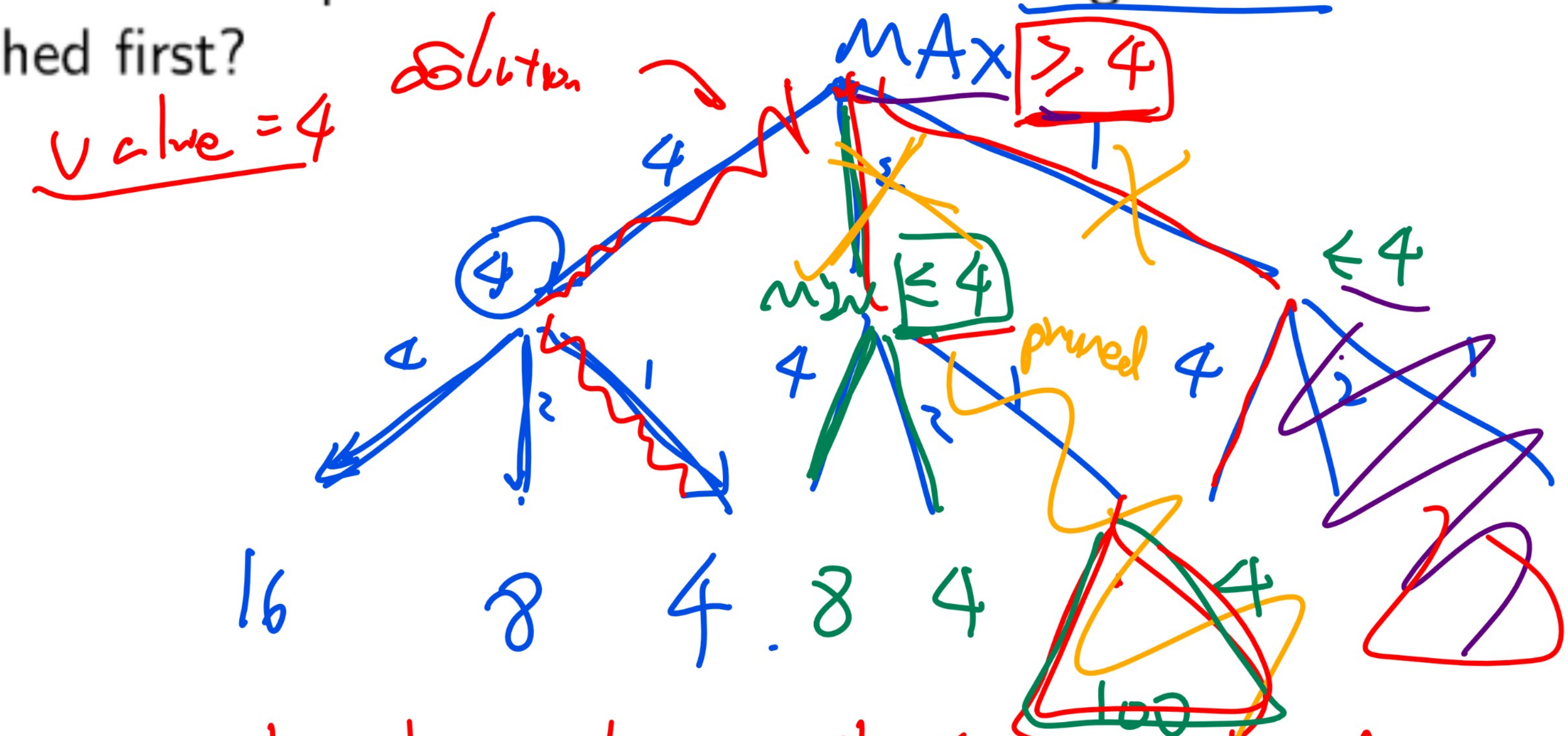

MAX does not care about how much worse than 4 they will get from action 2

Alpha Beta Example 3

Quiz

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. Alpha-Beta pruning is used. What is the number of branches (states) that can be pruned if the actions with smaller labels are searched first?
- A : 0
- B : 1
- C : 2
- D : 3
- E : 4

Alpha Beta Example 4

Quiz

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. Alpha-Beta pruning is used. What is the maximum number of branches (states) that can be pruned if the actions can be searched in any order?
- $A : 2$
- $B : 3$
- $C : 4$
- $D : 5$
- $E : 6$

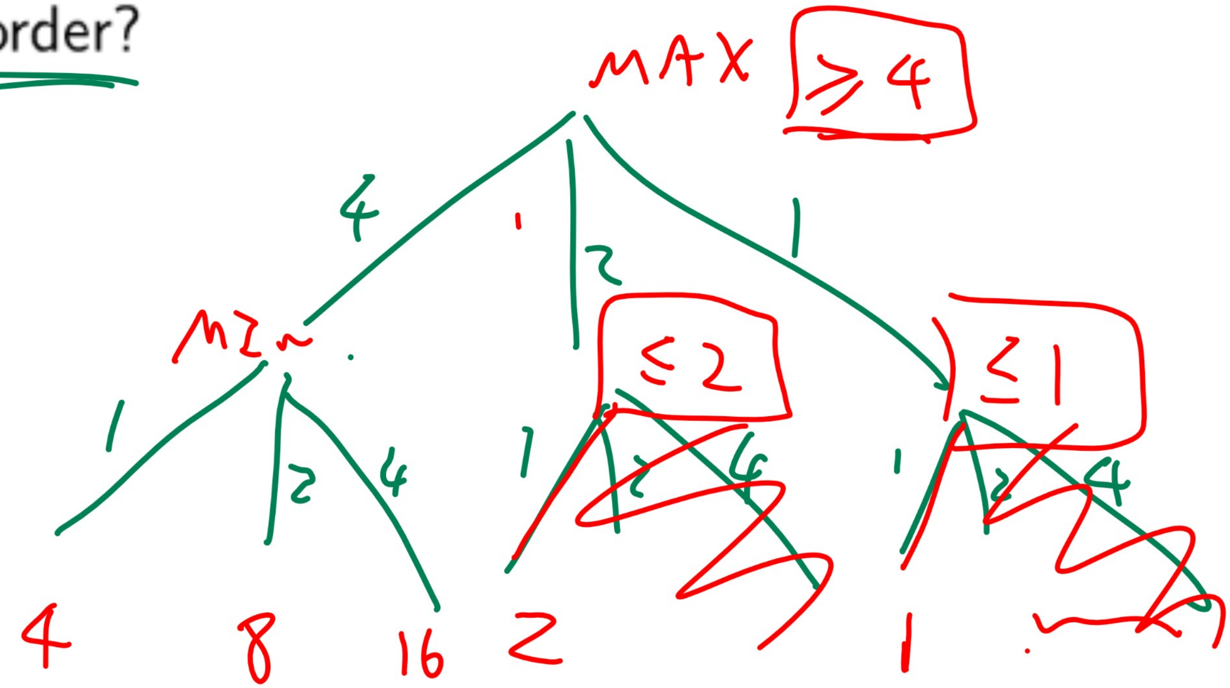

Alpha Beta Example 4

Quiz

Alpha Beta Example 4 Continued

Quiz

Alpha Beta Performance

Discussion

- In the best case, the best action of each player is the leftmost child.
- In the worst case, Alpha Beta is the same as minimax.

DFS IDS

Static Evaluation Function

Definition

- A static board evaluation function is a heuristics to estimate the value of non-terminal states.
- It should reflect the player's chances of winning from that vertex.
- It should be easy to compute from the board configuration.

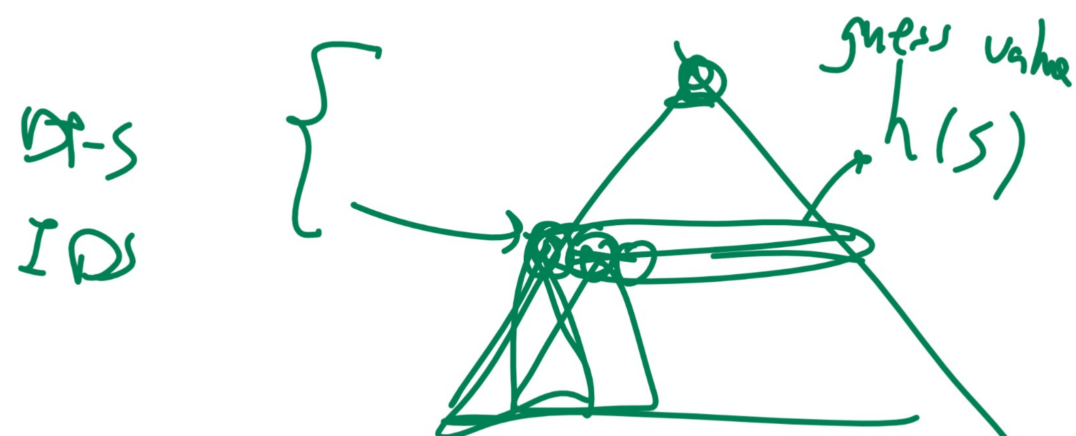

Linear Evaluation Function Example

Definition

- For Chess, an example of an evaluation function can be a linear combination of the following variables.

1. Material. x_1
2. Mobility. x_2
3. King safety. x_3
4. Center control. x_4

- These are called the features of the board.

$f(x_1, \ldots x_4) \approx$ prob of winning for MAX player.

Iterative Deepening Search

Discussion

- IDS could be used with SBE.
- In iteration d, the depth is limited to d, and the SBE of the non-terminal vertices are used as their cost or reward.

IDS with SBE Diagram

Discussion

Non Linear Evaluation Function

Discussion

- The SBE can be estimated given the features using a neural network.
- The features are constructed using domain knowledge, or a possibly a convolutional neural network.
- The training data are obtained from games between professional players.

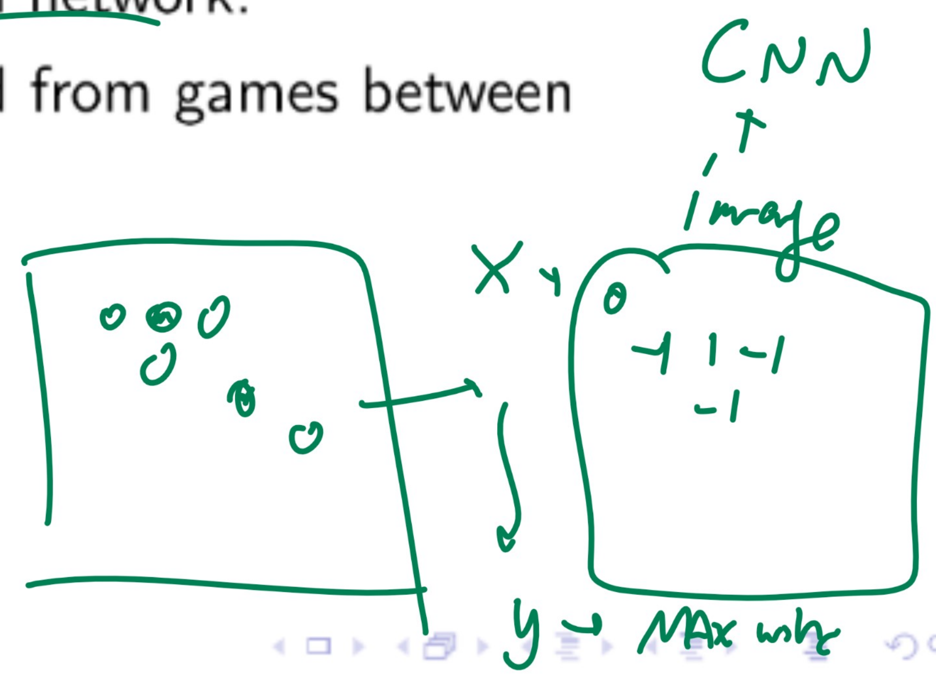

Monte Carlo Tree Search

Discussion

- Simulate random games by selecting random moves for both players.
- Exploitation by keeping track of average win rate for each successor from previous searches and picking the successors that lead to more wins.
- Exploration by allowing random choices of unvisited successors.

Monte Carlo Tree Search Diagram

Discussion

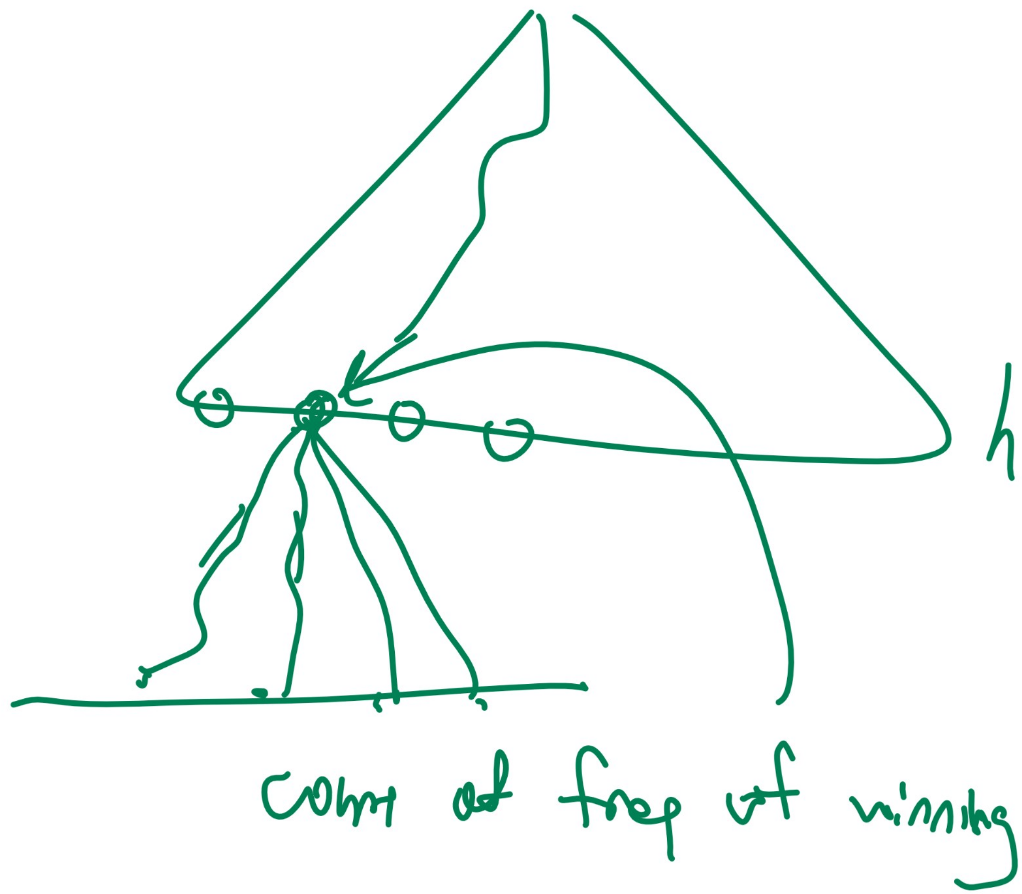

Alpha GO Example

Discussion

- MCTS with $> 10^5$ play-outs.
- Convolutional neural network to compute SBE.

Summary

Discussion

- Adversarial Search:

1. Sequential Move Games: Minimax $\rightarrow$ DFS on the game tree.
2. Sequential Move Games: Alpha-Beta Pruning $\rightarrow$ DFS to keep track α and β $\rightarrow$ prune the subtree with $\alpha \Rightarrow \beta$.
3. Simultaneous Move Games: Iterated Elimination of Strictly Dominated Strategies (Rationalizability).
4. Simultaneous Move Games: Nash Equilibrium.