

CS540 Introduction to Artificial Intelligence

Lecture 21

Young Wu

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

August 7, 2022

Lion Game Example

Quiz

Lion Game Example Diagram

Quiz

Summary

Discussion

Tic Tac Toe Example

Motivation

Nim Game Example

Motivation

Minimax Algorithm

Description

- Use DFS on the game tree.

Minimax Example

Quiz

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.

$$T = 1 + b + b^2 + \dots + b^d$$

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

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

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.

Alpha Beta Pruning

Description

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

Alpha Beta Example 1

Quiz

Alpha Beta Example 1 Continued

Quiz

Alpha Beta Example 2

Quiz

Alpha Beta Example 3

Quiz

Alpha Beta Example 4

Quiz

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.

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.
 - ② Mobility.
 - ③ King safety.
 - ④ Center control.
- These are called the features of the board.

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