Co-financed by the Connecting Europe Facility of the European Union

Prof. Dr. techn. Wolfgang Nejdl
Franziska Schoger

November 1st, 2023

Advanced Topics in AI
Exercise 3

Question 1: Minimax Algorithm

1. Briefly, describe the main idea of the Minimax algorithm.
2. Discuss the Time Complexity, Space Complexity, Completeness and Optimality of the Minimax algorithm.
3. Which search algorithm processes the nodes in an identical order to the Minimax-Algorithm?
4. Define an utility function that can be used for the game "Tic-Tac-Toe".
5. From https://en.wikipedia.org/wiki/Reversi: *[Othello] is a strategy board game for two players, played on an 8x8 uncheckered board. There are sixty-four identical game pieces called disks (often spelled "discs"), which are light on one side and dark on the other. Players take turns placing disks on the board with their assigned color facing up. During a play, any disks of the opponent's color that are in a straight line and bounded by the disk just placed and another disk of the current player's color are turned over to the current player's color. The object of the game is to have the majority of disks turned to display your color when the last playable empty square is filled.*

   Define an utility function for the game of Othello.
6. Compute the Min and Max values for the following game tree (Max is starting at the root):

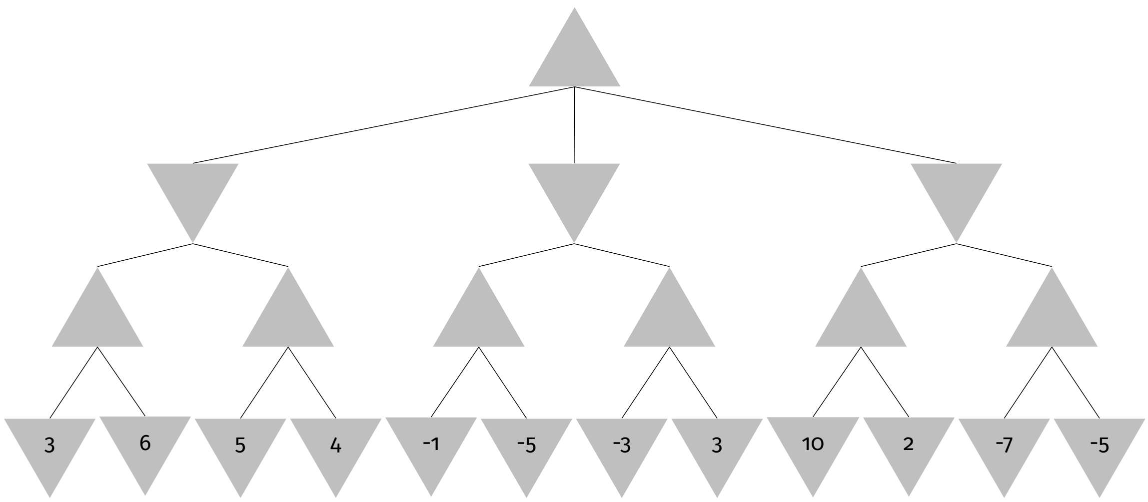

7. Which strategy should Max use?
8. What are disadvantages of the Minimax algorithm?

Question 2: α/β-pruning

1. Explain the main idea of α/β-pruning.
2. What does α and β stand for?
3. Figure 1 shows a game tree. Perfom α/β-pruning on it (Max is starting at the root).

Question 3: Nonzero-sum Games

1. Let's look at a non-zero-sum version of a game. In this formulation, player A's utility will be represented as the first of the two leaf numbers, and player B's utility will be represented as the second of the two leaf 3 numbers. Fill in this non-zero game tree assuming each player is acting optimally (A is starting at the root).

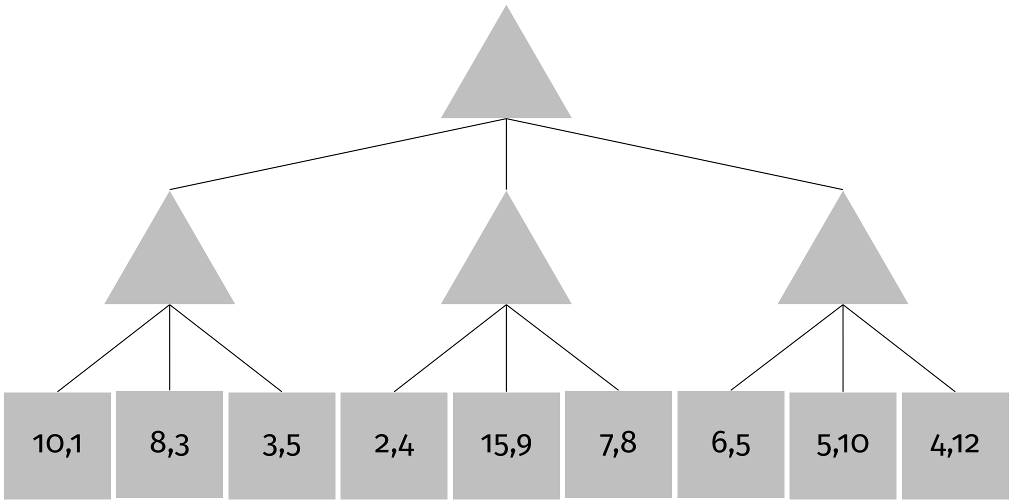

2. Which nodes can be pruned from the game tree above through alpha-beta pruning? If no nodes can be pruned, explain why not.

Question 4: Surrealist Pacman

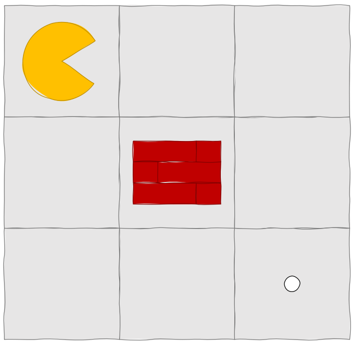

In the game of Surrealist Pacman, Pacman plays against a moving wall. On Pacman's turn, Pacman must move in one of the four cardinal directions, and must move into an unoccupied square. On the wall's turn, the wall must move in one of the four cardinal directions, and must move into an unoccupied square. The wall cannot move into a dot-containing square. Staying still is not allowed by either player. Pacman's score is always equal to the number of dots he has eaten.

The first game begins in the configuration shown. Pacman moves first.

a) Draw a game tree with one move for each player. Nodes in the tree represent game states (location of all agents and walls). Edges in the tree connect successor states to their parent states. Draw only the legal moves.

b) According to the depth-limited game tree you drew above what is the value of the game? Use Pacman's score as your evaluation function.

c) If we were to consider a game tree with ten moves for each player (rather than just one), what would be the value of the game as computed by minimax?

A second game is played on a more complicated board. A partial game tree is drawn, and leaf nodes have been scored using an (unknown) evaluation function *e*.

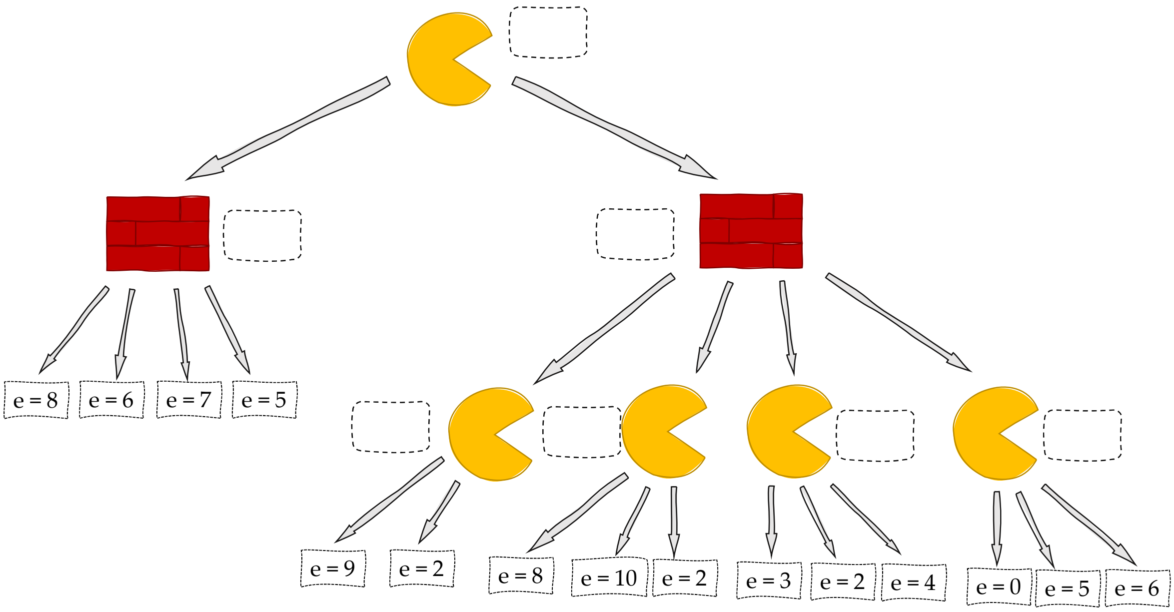

d) In the dashed boxes, fill in the values of all internal nodes using the minimax algorithm.

e) Cross off any nodes that are not evaluated when using alpha-beta pruning (assuming the standard left-to-right traversal of the tree).

Question 5: Minimax and Expectimax

In this problem, you will investigate the relationship between expectimax trees and minimax trees for zero-sum two player games. Imagine you have a game which alternates between player 1 (max) and player 2. The game begins in state s_0, with player 1 to move. Player 1 can either choose a move using minimax search, or expectimax search, where player 2's nodes are chance rather than min nodes.

1. Draw a (small) game tree in which the root node has a larger value if expectimax search is used than if minimax is used, or argue why it is not possible.
2. Draw a (small) game tree in which the root node has a larger value if minimax search is used than if expectimax is used, or argue why it is not possible.
3. Under what assumptions about player 2 should player 1 use minimax search rather than expectimax search to select a move?
4. Under what assumptions about player 2 should player 1 use expectimax search rather than minimax search?

5. Imagine that player 1 wishes to act optimally (rationally), and player 1 knows that player 2 also intends to act optimally. However, player 1 also knows that player 2 (mistakenly) believes that player 1 is moving uniformly at random rather than optimally. Explain how player 1 should use this knowledge to select a move. Your answer should be a precise algorithm involving a game tree search, and should include a sketch of an appropriate game tree with player 1's move at the root. Be clear what type of nodes are at each play and whose turn each play represents

Question 6: Utilities

1. Consider a utility function of $U(x) = 2x$. What is the utility for each of the following outcomes?
   a) 3
   b) $L(\frac{2}{3}, 3; \frac{1}{3}, 6)$
   c) -2
   d) $L(0.5, 2; 0.5, L(0.5, 4; 0.5, 6))$
2. Consider a utility function of $U(x) = x^2$. What is the utility for each of the following outcomes?
   a) 3
   b) $L(\frac{2}{3}, 3; \frac{1}{3}, 6)$
   c) -2
   d) $L(0.5, 2; 0.5, L(0.5, 4; 0.5, 6))$
3. What is the expected monetary value (EMV) of the lottery $L(\frac{2}{3}, \$3; \frac{1}{3}, \$6)$?
4. For each of the following types of utility function, state how the utility $U(L)$ of a lottery $L := L(p_1, x_1; \ldots; p_n, x_n)$ compares to the utility of the amount of money equal to the EMV of the lottery, $U(EMV(L))$. Write $\leq$, $\geq$, $=$, or ? for can't tell. (Assume that U is at least twice differentiable.)
   a) U is an arbitrary function.
   b) U is monotonically increasing and its rate of increase is increasing (its second derivative is positive).
   c) U is monotonically increasing and linear (its second derivative is zero).
   d) U is monotonically increasing and its rate of increase is decreasing (its second derivative is negative).

Question 7: Lotteries in Ghost Kingdom

Diverse Utilities. Ghost-King (GK) was once great friends with Pacman (P) because he observed that Pacman and he shared the same preference order among all possible event outcomes. Ghost-King, therefore, assumed that he and Pacman shared the same utility function. However, he soon started realizing that he and Pacman had a different preference order when it came to lotteries and, alas, this was the end of their friendship.
Let Ghost-King and Pacman's utility functions be denoted by U_{GK} and U_P respectively. Assume both U_{GK} and U_P are guaranteed to output non-negative values.

1. Which of the following relations between U_{GK} and U_P are consistent with Ghost King's observation that U_{GK} and U_P agree, with respect to all event outcomes but not all lotteries?
   a) $U_P = aU_{GK} + b \quad (0 < a < 1, b > 0)$
   b) $U_P = aU_{GK} + b \quad (a > 1, b > 0)$
   c) $U_P = U_{GK}^2$
   d) $U_P = \sqrt{U_{GK}}$
2. In addition to the above, Ghost-King also realized that Pacman was more risk-taking than him. Which of the relations between U_{GK} and U_P are possible?
   a) $U_P = aU_{GK} + b \quad (0 < a < 1, b > 0)$
   b) $U_P = aU_{GK} + b \quad (a > 1, b > 0)$
   c) $U_P = U_{GK}^2$
   d) $U_P = \sqrt{U_{GK}}$

Guaranteed Return. Pacman often enters lotteries in the Ghost Kingdom. A particular Ghost vendor offers a lottery (for free) with three possible outcomes that are each equally likely: winning \$1, \$4, or \$5.
Let $U_P(m)$ denote Pacman's utility function for \$m. Assume that Pacman always acts rationally.

1. The vendor offers Pacman a special deal - if Pacman pays \$1, the vendor will manipulate the lottery such that Pacman **always gets the highest reward possible**. For which of these utility functions would Pacman choose to pay the \$1 to the vendor for the manipulated lottery over the original lottery? (Note that if Pacman pays \$1 and wins \$m in the lottery, his actual winnings are \$m-1.)
   a) $U_P(m) = m$
   b) $U_P(m) = m^2$
2. Now assume that the ghost vendor can only manipulate the lottery such that Pacman **never gets the lowest reward** and the remaining two outcomes become equally likely. For which of these utility functions would Pacman choose to pay the \$1 to the vendor for the manipulated lottery over the original lottery?

a) $U_P(m) = m$

b) $U_P(m) = m^2$

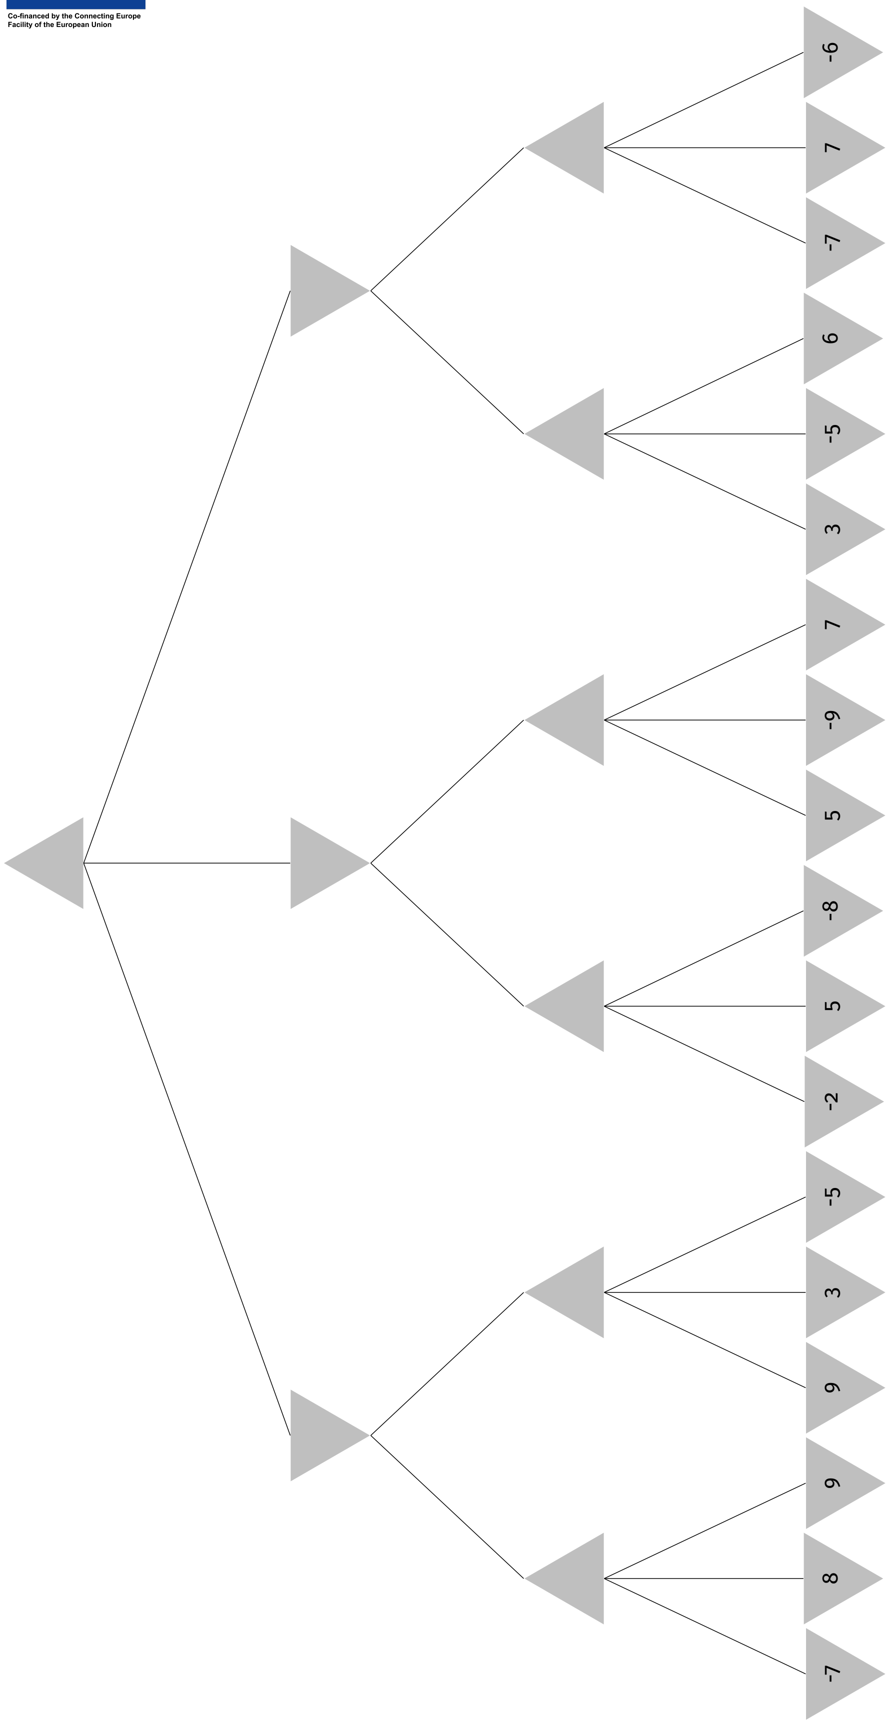

Figure 1: Game Tree for α/β-pruning