Algorithmic Game Theory
Problems 4

22.05.2022

Problem 1. *Consider a $2 \times n$ zero-sum game where the best response to every pure strategy is unique. Suppose that the top row and leftmost column define a pure-strategy equilibrium of the game.*

(a) Show that this equilibrium can be found by iterated elimination of strictly dominated strategies. [Hint: Consider first the case $n = 2$.]

(b) Does the claim in (a) still hold if some pure strategy has more than one best response?

Problem 2. *This problem exercises the basic concepts of game playing, using tic-tac-toe as an example. We define X_n as the number of rows, columns, or diagonals with exactly n $X's$ and no $O's$. Similarly, O_n is the number of rows, columns, or diagonals with just n $O's$. The utility function assigns $+1$ to any position with $X_3 = 1$ and -1 to any position with $O_3 = 1$. All other terminal positions have utility 0. For nonterminal positions, we use a linear evaluation function defined as $Eval(s) = 3X_2(s) + X_1(s) - (3O_2(s) + O_1(s))$.*

- *Approximately how many possible games of tic-tac-toe are there?*
- *Show the whole game tree starting from an empty board down to depth 2 (i.e., one X and one O on the board), taking symmetry into account.*
- *Mark on your tree the evaluations of all the positions at depth 2.*
- *Using the minimax algorithm, mark on your tree the backed-up values for the positions at depth 1 and 0, and sue those values to choose the best starting move.*
- *Circle the nodes at depth 2 that would not be evaluated if alpha-beta pruning were applied, assuming the nodes are generated in the optimal order for alpha-beta pruning.*

Problem 3. *Consider the two-player game tree shown below. The values of the max nodes are actual values for the associated states as given in the game description, not state utilities by game tree search.*

a *What is the state utility of the top of the tree (as determined by Minimax).*

b *Now consider a player using fixed-depth heuristic search with depth limit 1. How many max nodes are searched in evaluating the top node in this tree.*

c *Suppose the player uses a goal proximity heuristic with state reward as the heuristic value for non-terminal states. What is the minimal final reward for this player in this game?*

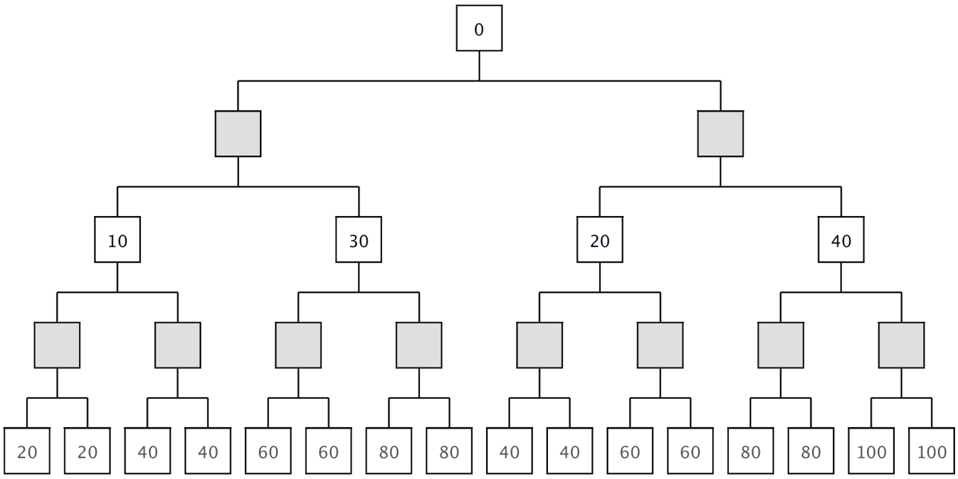

Problem 4. *(Bonus Exercise) Describe how the minimax and alpha-beta algorithms change for two-player,* ***non-zero-sum*** *games in which each player has his or her own utility function. You may assume that each player knows the other's utility function. If there are no constraints on the two terminal utilities, is it possible for any node to be pruned by alpha-beta?*