

CPSC 474/574 S18 Exam #2 Solutions

CPSC 474/574 - Spring 2018 - Exam #2

Write your name and NetID on this cover page and only on this cover page. Please indicate clearly your final answers to each question. Do not write within $\frac{1}{2}$ inch of the edges of the pages.

Problem 1 (12 points): Consider a two-player game played with n piles of stones each initially containing k stones. On each turn, the current player may take one stone from each remaining pile, or divide any pile containing at least 2 stones into 2 or 3 non-empty piles. For example if the sizes of the piles are $[3, 2]$ then the possible moves are to $[2, 1]$ (by taking one stone from each pile), to $[3, 1, 1]$ (by dividing the pile of 2 into two piles of 1), to $[2, 1, 1, 1]$ (by dividing the pile of 3 into three piles of 1), and to $[2, 2, 1]$ (by dividing the pile of 3 into a pile of 1 and a pile of 2). The player who makes the last move wins.

Find the Grundy numbers of every position reachable starting with one pile of 4 stones.

4	1
3 1	2
2 2	2
2 1 1	0
1 1 1 1	1
3	2
2 1	0
1 1 1	1
2	0
1 1	1
1	1
0	0

Problem 2 (15 points): Consider a solitaire game played on a h -row, w -column grid of positive numbers. You start with a marker on the bottom right corner of the grid (location $(h - 1, w - 1)$). On each turn you can choose to move left or up, and your marker moves to a position in that direction chosen uniformly randomly and the number in the resulting position is added to your score. If your marker is already in the top row or left column then you may not choose up or left respectively. The game ends when the marker is in the top left corner (position $(0, 0)$).

For example, in the pictured game the player chose up, left, up, and left and the randomly selected moves were 1, 2, 2, and 1 spaces respectively. The total score was 11 (note that it does not matter what is in the lower right corner since that value is never added to the score).

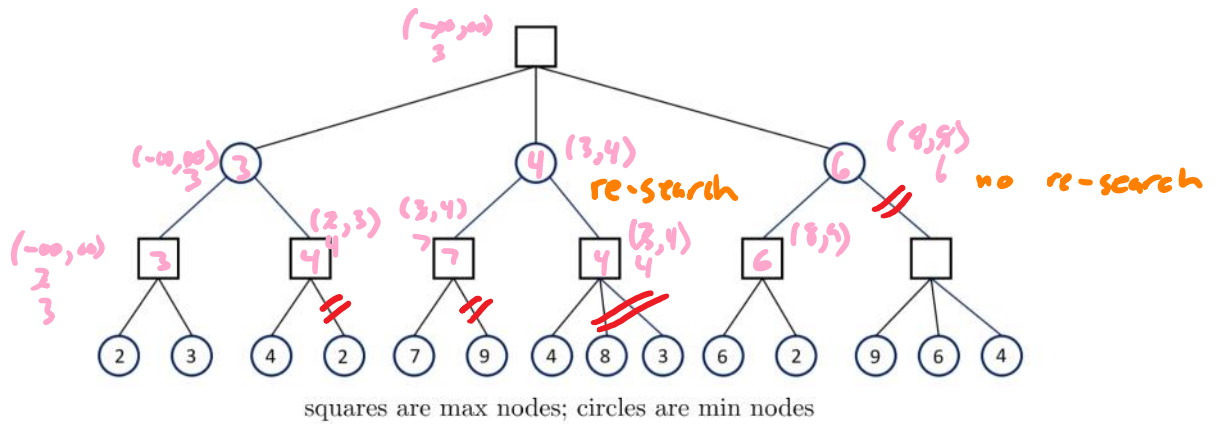
1	4	6	2
2	5	3	2
3	5	6	1
6	5	4	

Let $E(r, c)$ be the expected future winnings when following the optimal strategy starting at position (r, c) . Write a recurrence (including the base cases) for $E(r, c)$. Assume that $A[r, c]$ holds the value of the number at row r , column c .

$$E(r, c) = \begin{cases} 0 & \text{if } r = 0 \text{ and } c = 0 \\ \frac{1}{r} \cdot \sum_{i=0}^{r-1} (A[i, 0] + E[i, 0]) & \text{if } c = 0 \text{ and } r > 0 \\ \frac{1}{c} \cdot \sum_{j=0}^{c-1} (A[0, j] + E[0, j]) & \text{if } r = 0 \text{ and } c > 0 \\ \max \left(\frac{1}{r} \cdot \sum_{i=0}^{r-1} (A[i, c] + E[i, c]), \right. & \text{otherwise} \\ \left. \frac{1}{c} \cdot \sum_{j=0}^{c-1} (A[r, j] + E[r, j]) \right) & \end{cases}$$

Problem 3 (12 points): Illustrate the operation of Scout on the given game tree by showing

- any pruning that occurs during the initial calls to the children of the root (not any pruning that happens during any re-searches)
- the null windows that are passed initially to the second and third children of the root
- whether or not a re-search is required after the initial null-window calls to the second and third children of the root, and if so, the full window then passed to the child to start the re-search.



Problem 4: (12 points) Consider the constant sum game with the following payoff matrix.

$$\begin{pmatrix} 6 & 3 \\ 2 & 4 \end{pmatrix}$$

For each of the following pairs of mixed strategies, determine if it is a saddle point. If it is, then show that it is and find its value; if it is not then find player I's best response to the given strategy for player II (find the best response to Y_i).

(a) $X_1 = (\frac{1}{2} \ \frac{1}{2}), Y_1 = (\frac{1}{3} \ \frac{2}{3})$

(b) $X_2 = (\frac{2}{5} \ \frac{3}{5}), Y_2 = (\frac{1}{5} \ \frac{4}{5})$

$$E(X_2, Y_2) = \frac{2}{5} \cdot \frac{1}{5} \cdot 6 + \frac{2}{5} \cdot \frac{4}{5} \cdot 3 + \frac{3}{5} \cdot \frac{1}{5} \cdot 2 + \frac{3}{5} \cdot \frac{4}{5} \cdot 4 = \frac{50}{25} = \frac{10}{5}$$

$$E(X_2, 1) = \frac{2}{5} \cdot 6 + \frac{3}{5} \cdot 2 = \frac{18}{5} \quad E(1, Y_2) = \frac{1}{5} \cdot 6 + \frac{4}{5} \cdot 3 = \frac{16}{5}$$

$$E(X_2, 2) = \frac{2}{5} \cdot 3 + \frac{3}{5} \cdot 4 = \frac{18}{5} \quad E(2, Y_2) = \frac{1}{5} \cdot 2 + \frac{4}{5} \cdot 4 = \frac{18}{5}$$

X_2, Y_2 is a saddle point

$$E(X_1, Y_1) = \frac{1}{2} \cdot \frac{1}{3} \cdot 6 + \frac{1}{2} \cdot \frac{2}{3} \cdot 3 + \frac{1}{2} \cdot \frac{1}{3} \cdot 2 + \frac{1}{2} \cdot \frac{2}{3} \cdot 4 = \frac{22}{6} \neq \frac{10}{5}$$

not a saddle pt $E(1, Y_1) = \frac{1}{3} \cdot 6 + \frac{2}{3} \cdot 3 = \frac{10}{3} \quad E(2, Y_1) = \frac{1}{3} \cdot 2 + \frac{2}{3} \cdot 4 = \frac{10}{3}$
best response is (1 0)

Problem 5: (10 points) Consider a multi-armed bandit with three arms: one that pays 1 or 6, each with probability $\frac{1}{2}$; one that pays 0 or 8, each with probability $\frac{1}{2}$; and one that pays 2 or 3, each with probability $\frac{1}{2}$.

Consider a policy that plays each arm once, eliminates the arm with the lowest payoff during those plays, and then alternates playing the remaining two arms. Compute the exact value of the expected value of the limit of the average regret for this policy on this machine.

$$\begin{aligned}
 P(\text{1st arm eliminated}) &= P(1, 8, x) = \frac{1}{4} & \mu_1 &= \frac{7}{2} \\
 P(\text{2nd arm eliminated}) &= P(0, x, y) = \frac{1}{2} & \mu_2 &= 4 = \mu^* \\
 P(\text{3rd arm eliminated}) &= \frac{1}{4} & \mu_3 &= \frac{5}{2}
 \end{aligned}$$

$$\begin{aligned}
 \lim_{T \rightarrow \infty} \frac{1}{T} \cdot \sum_{i=1}^T R_T &= \frac{1}{4} \cdot \left(\frac{1}{2}(\mu^* - \mu_2) + \frac{1}{2}(\mu^* - \mu_3) \right) \\
 &\quad + \frac{1}{2} \cdot \left(\frac{1}{2}(\mu^* - \mu_1) + \frac{1}{2}(\mu^* - \mu_3) \right) \\
 &\quad + \frac{1}{4} \cdot \left(\frac{1}{2}(\mu^* - \mu_1) + \frac{1}{2}(\mu^* - \mu_2) \right) \\
 &= \frac{1}{4} \cdot \frac{3}{4} + \frac{1}{2} \cdot 1 + \frac{1}{4} \cdot \frac{1}{4} = \frac{3}{4}
 \end{aligned}$$

Problem 6: (16 points) Describe one way of incorporating domain knowledge in Monte Carlo Tree Search.

MAST - for each possible move, record number of times that move led to a win or loss. Bias payout towards moves with better win rate.

Problem 7: (8 points) The following code implements part of the tree policy for Monte Carlo Tree Search for a one-player game. $r[pos]$ is the average observed reward for a position, $n[pos]$ is the number of times a position has been reached by the search, and $moves(pos)$ returns a list of the positions reachable in one move from position pos (assume that the game tree is in fact a tree, so there is at most one path from the root to any given position).

```

BEST-CHILD(p)
  max = -infinity
  move = null

  for c in moves(p)

    exploit =  $r[p]$  if  $c.next() == p.next()$  else  $-r[c]$ 

    explore =  $\sqrt{2 * \log(n[p]) / n[c]}$ 

    ucb = exploit + explore

    if ucb > max
      max = ucb
      move = c

  return move

```

Modify the above code to work for a two-player, zero-sum combinatorial game, assuming that


- for a given position p , the next player to move at some positions in $moves(p)$ could be the same as the player to move at position p and could be different for other positions in $moves(p)$, and
- $r[p]$ is the observed reward at position p from the point of view of the next player at p : positive if that player is likely to win and negative otherwise.

You may assume that there is a function $next(p)$ that returns the index of the player making the next move at position p .

Problem 8: (15 points) Describe the inputs you might use for an artificial neural network that, for a position in 2-player Yahtzee, estimates the probability that the first player wins (a value network for 2-player Yahtzee).

one input per category on each scoresheet - 0 for unused, 1 for used
 two inputs per scoresheet for Yahtzee category - 00 for unused, 01 for 0, 10 for 50
 one input per scoresheet for upper total - $\frac{\text{upper total}}{63}$
 one input for score margin, normalized to range -1 to 1
 one input for current turn - 0 for P1, 1 for P2
 one input for rolls left on current turn - $\frac{\text{rolls}}{2}$
 ? inputs for current roll - (figure this out for Assignment #6)

Problem 7 (10 points): Illustrate the operation of Alpha-Beta pruning on the given game tree by showing the (α, β) window passed to each non-leaf and which branches are pruned (you need not indicate the values returned).



From SIC Exam #1

Problem 8 (10 points): Illustrate the operation of Scout on the given game tree by showing the null window that is passed to the right child of the root and then the results of the search with that null window: show the (α, β) windows passed to nodes in the right subtree of the root and which branches are pruned in that subtree. Indicate whether the right child of the root will be re-searched as a result of the initial null-window search. Note that while the left subtree is the same as the previous two, the right subtree is different.

