Two hours

UNIVERSITY OF MANCHESTER
SCHOOL OF COMPUTER SCIENCE

AI and Games

Date: Thursday 17th May 2018

Time: 09:45 - 11:45

Please answer all Questions.

Use a SEPARATE answerbook for each SECTION

This is a CLOSED book examination

The use of electronic calculators is permitted provided they are not programmable and do not store text

Section A

Semester 1

This text appears only in the exam paper version

1. Normal Form Games

a) What follows is the pay-off for a two-player, zero-sum game with perfect information and no chance. Player 1 has five actions: A – E. Player 2 also has five actions: a – e. The table shows the pay-off to Player 1; the pay-off to Player 2 is minus this. For example, if Player 1 plays B and Player 2 plays a, Player 1 gets 4 and Player 2 gets -4. The two players choose their moves simultaneously.

		Player 2				
		a	b	c	d	e
Player 1	A	5	4	0	3	6
	B	4	4	-2	5	2
	C	8	7	4	6	80
	D	104	108	-2	50	8
	E	6	5	4	4	2

i) Find a (Nash) equilibrium using the minimax approach. Justify your answer. (3 marks)

ii) The concept of *dominance* can be used to reduce the size of the pay-off table in certain instances. Find a (Nash) equilibrium by repeatedly appling dominance to shrink the pay-off table as much as possible. (3 marks)

b) This question involves the following game.

		Bob	
		a	b
Alice	A	$(1,1)$	$(-1,-1)$
	B	$(-1,-1)$	$(2,2)$

i) What type of game is this? Name all the categories of game this game is in. (3 marks)

ii) Find the pure strategy equilibrium. If there is more than one, find all of them. For every equilibrium found, justify that it is a Nash equilibrium. (3 marks)

iii) Is there a fully mixed strategy equilibrium? If so, find it. If not give a convincing argument that one does not exist. (3 marks)

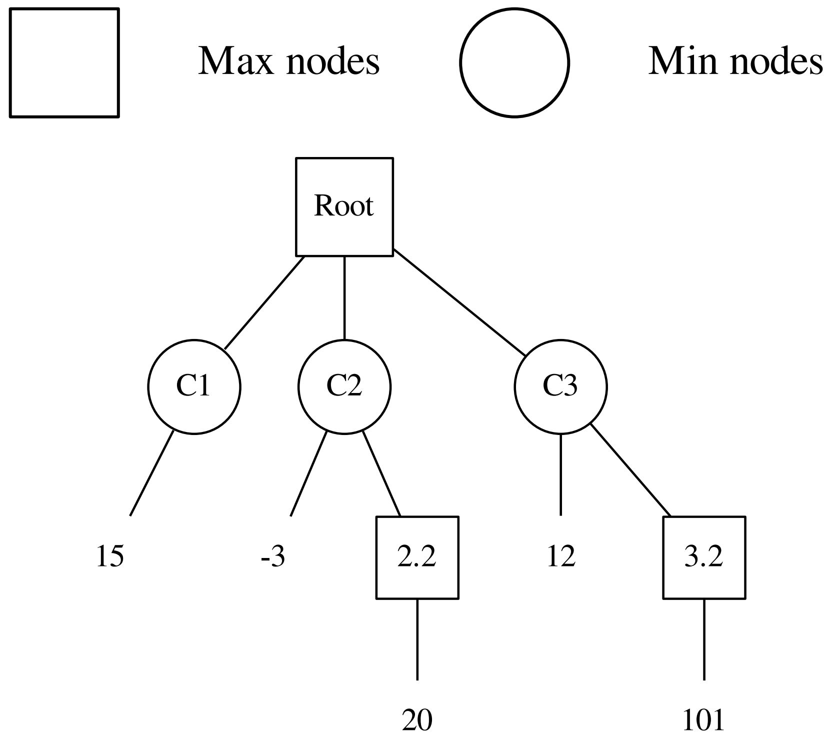

Figure 1: A game tree for Question 2a. Square nodes are Player 1 (MAX) decision nodes. Circle nodes are Player 2 (MIN) decision nodes. Numbers without shapes are payoffs. The numbers within the shapes are labels.

[Next page →]

2. Extensive Form Games

   a) Figure 1 shows a game tree in extensive form. Apply alpha-beta pruning to this tree to answer the following questions:

      i) After the algorithm is completed, what are the alpha, and beta values of the root node and its three children? (4 marks)

      ii) What nodes are *not* evaluated? (2 marks)

   b) This question concerns the game Nim-N introduced in lectures. The game starts with N matches. Each player alternates taking 1, 2, or 3 matches. *The player taking the last match loses.* When there are less than 3 matches left, the player who is on turn can take between 1 and all of the remaining matches.

      Consider N=4, so the game starts with 4 matches. Apply Win-Loss pruning to the Nim-4 game tree, where the children are ordered from the most matches taken to 1 match taken. (E.g. the first child is the result of the root node taking three matches, and the last child is the result of 1 match taken.) How much of the tree is built before the value of the root node and the appropriate action from the root node has been determined. Assume the root node is the MAX player. Justify you answer.

      (3 marks)

   c) These questions are about heuristics.

      i) Describe three heuristics for Kalah. (3 marks)

      ii) The stronger your heuristic, the stronger your player will be. Describe the most effective technique you know of to build a strong heuristic. Describe also how you would compare heuristics to determine which are the most effective ones. (3 marks)

Section B

Semester 2

3. a) Find a Stackelberg strategy for the following Stackelberg pricing game:
There are two players of which L is the leader and F is the follower. The leader's strategy (i.e., price) is u_L and the follower's strategy (i.e., price) is u_F. Assume that the demand functions (i.e., price-sale relationships) for the leader and follower are $S_L(u_L, u_F) = 10 - 4.2u_L + 2u_F$ and $S_F(u_L, u_F) = 3 + 0.6u_L - 3u_F$ respectively. Then the profit functions to be maximised for the leader and follower are

$$J_L(u_L, u_F) = (u_L - c_L)S_L(u_L, u_F) = (u_L - c_L)(10 - 4.2u_L + 2u_F)$$
$$J_F(u_L, u_F) = (u_F - c_F)S_F(u_L, u_F) = (u_F - c_F)(3 + 0.6u_L - 3u_F)$$

where the unit costs and strategy spaces for the leader and follower are $c_L = c_F = 1$ and $U_L = U_F = [1, +\infty)$ respectively. (7 marks)

Note. For the above question, 1 mark is for the correct answer and the other marks are for the step by step calculation and reasoning.

b) Find ALL Stackelberg strategies for the following Stackelberg game on discrete strategy spaces:
There are two players of which L is the leader and F is the follower. The leader's strategy is u_L and the follower's strategy is u_F. Further the payoff functions for the leader and follower are

$$J_L(u_L, u_F) = u_L + u_F$$
$$J_F(u_L, u_F) = u_F - 2u_L u_F$$

respectively, and the strategy spaces are $U_L = U_F = \{0, 1, 2\}$ (that is, each player's strategy can only take the value as 0, 1, or 2). (5 marks)

Note. For the above question, you need to give both your answer and the step by step calculation.

c) In a two-person Stackelberg game with imperfect information, please list or propose two reasonable options or solutions that the leader can use to play the game. (3 marks)

Note. The two solutions you propose can be under the different scenarios or assumptions. For example, one solution can be used under the assumption that there are historical data available,whereas another solution can be used under the assumption that there is no historical data.

4. a) Assume that the recursive least square method is used to learn the reaction function in a two-person Stackelberg game with the estimation model as

$$\hat{y}(t) = \hat{R}[X(t), \theta_t]$$

Please give the recursive formula to update parameter θ_T when a new pair of data $[X(T+1), y(T+1)]$ becomes available, and explain the meaning of each item in the formula. (3 marks)

Note. You do not need to give those formulas to calculate the adjusting factor in your formula.

b) Consider the following single-parameter multi-item auction game:

i) There are k items for sale, in which item i includes α_i units of the same product, that is, $Item_i = \{\alpha_i\}$ $(i = 1,2,...,k)$. Further it is assumed that $Item_1 > Item_2 > ... > Item_k$ (that is, $\alpha_1 > \alpha_2 > ... > \alpha_k$);

ii) There are n bidders and $n > k$. Bidder i has his private valuation v_i to per unit of the product and bidder i' objective is to maximise his utility function: $u_i = v_i x_i - p_i$ $(i = 1,2,...,n)$, where x_i is the number of the product with an item if bidder i wins the item and 0 otherwise, and p_i is the price bidder i pays if he/she wins an item and 0 otherwise. Further it is assumed that $v_1 > v_2 > ... > v_n$;

iii) There is a mechanism designer (seller) whose objective is to maximise the social surplus function

$$F(x_1, x_2, ..., x_n) = v_1 x_1 + v_2 x_2 + ... + v_n x_n$$

The mechanism designer's task is to design the allocation rule and payment rule to form a sealed-bid auction with the three steps: 1) Collect bids $b = (b_1, b_2, ..., b_n)$; 2) Apply the allocation rule; 3) Apply the payment rule.

For the above multi-item auction game, answer the following questions:

i) In what sense can an allocation be regarded as feasible? (2 marks)

ii) What is an allocation rule? (2 marks)

iii) What is a payment rule? (2 marks)

iv) Under the assumption that the true valuation of each bidder is known by the mechanism designer, please give an allocation rule which maximises the social surplus function. (2 marks)

v) Prove that the allocation rule you give above maximises the social surplus function. (4 marks)

END OF EXAMINATION