COMP 34120

Two hours

UNIVERSITY OF MANCHESTER
SCHOOL OF COMPUTER SCIENCE

AI and Games

Date: Thursday 25th May 2017

Time: 09:45 - 11:45

Please answer any THREE Questions from the FOUR Questions provided

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

[PTO]

Section A

Semester One

1. Normal Form Games

   a) What follows is the pay-off table for a two-player, zero-sum game. 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 the negative of this. For example, if Player 1 plays D and Player 2 plays e, Player 1 gets 80 and Player 2 gets -80. The two players choose their moves simultaneously.

		Player 2				
		a	b	c	d	e
	A	4	3	-1	2	5
	B	3	3	-3	4	1
Player 1	C	7	6	3	5	8
	D	104	100	-3	50	80
	E	5	4	3	3	1

   i) Assuming the game will be played only once and both players know the pay-off table, what action should Player 1 choose and what action should Player 2 choose? Justify your answers. (3 marks)

   ii) The concept of *dominance* can be used to reduce the size of the pay-off table in this problem. Using dominance, find the actions which need not be considered by the two players. Draw the resulting pay-off table. (3 marks)

b) Here is a variant of Rock-Paper-Scissors which I made up, called *Even and Odd* 3[1] The two players are the Even player and the Odd player. Rules of the game are these.

- First, a coin is flipped to determine which player is Even and which Odd.
- Next, like in Rock-Paper-Scissors, the players say, "one, two, three, go!" and on "go" show their hand, which can have one, two, or three fingers extended.
- If the sum of the extended fingers on the two players' hands is even, Even wins, otherwise Odd wins. (For example, if one player extends one finger and the other extends two fingers, the sum is three and Odd wins.)

i) Is this game symmetric, i.e. are the two players equivalent? (2 marks)

ii) Are there any equilibria involving only pure strategies? Justify your answer. (2 marks)

iii) Inspired by Rock-Paper-Scissors, the two players decide to play each of the three actions with equal probability, i.e. they choose to play one finger, two fingers, or three fingers each with a probability $1/3$. Is this a Nash equilibrium? Justify your answer. (3 marks)

iv) Now consider *Even and Odd* 2 in which each player can produce one finger or two fingers, but not three fingers. Find a Nash equilibrium for this game. (3 marks)

v) Observe that playing one finger and three fingers is exactly equivalent. That is, if either player switches from playing one to playing three, or vice versa, the pay-off to both players is exactly the same. Using this fact, and the result from the previous question (part 1(b)iv), find a Nash equilibrium for the original Even and Odd 3. What is the value of the game? If possible, find a second Nash equilibrium. (4 marks)

[1] The 3 denotes that each player can hold up $1, 2,$ or, 3 fingers.

2. Extensive Form Games

   a) We learned in this course that for certain times of games, it is possible to be said that the game is "solved".

      i) Give all the possible meanings to the statement that a particular game has been solved. (2 marks)
      ii) What are the conditions which the game has to satisfy for it to be guaranteed that a game has a solution? (2 marks)

   b) We consider an *extensive form* version of a game which was considered in the previous question, but will be fully described here.

      - Player 1 is called Even and goes first; Player 2 is called Odd and goes second.
      - Each player at its go chooses a number 1, 2, or, 3.
      - Each player has one go, then the game ends.
      - If the sum of numbers from the two players adds to an even number, Player 1 wins, otherwise Player 2 wins.

      Obviously, this is a trivial game, but we are going to use it to test your knowledge of pruning.

      i) Draw the game tree. (2 marks)
      ii) This game can be viewed as a Win-Lose game. The players win or they lose. In this view, nodes take values as W or L. Briefly describe the method for pruning in Win-Lose games, and list or show the links which are pruned. (3 marks)
      iii) Now consider doing alpha-beta pruning on this tree from the perspective of Player 1 (Even) who goes first. After the algorithm has completed and the root node's value has been determined, give the values of each node, and the alpha and beta values at each node. Which children of each node have been pruned? (4 marks)

   c) These questions are about heuristics.

      i) Describe three simple heuristics for Kalah. (3 marks)
      ii) You want to make as effective a heuristic as you possibly can. Describe how you would produce an effective heuristic and how you would compare heuristics to determine which are the most effective ones. (4 marks)

Section B

Semester Two

3. a) We are using the recursive least square algorithm to update the parameters of the follower's reaction function in a Stackelberg game, and the following information is known:

   i) The follower's reaction function learned from the historical data $\{u_L(t), u_F(t)\}$ $(t = 1, 2, ..., T)$ is $u_F = R[u_L, \theta_T] = a_T + b_T u_L$, where $\theta_T = [a_T, b_T]^\tau$ represents the parameter vector at time T and τ is the vector transpose.

   ii) $\{u_L(T+1), u_F(T+1)\}$ is a pair of new data received.

   iii) L_{T+1} is the adjusting factor in the recursive least square algorithm.

   Based on the above known information, please give the mathematical formula to calculate the new parameter vector θ_{T+1} in the recursive least square algorithm. (Note. You are not required to give the formula to calculate L_{T+1} but just assume that it is known.) (2 marks)

b) Consider a two-person Stackelberg game and answer the following questions:

   i) What is the definition of Stackelberg strategy (i.e., Stackelberg equilibrium)? (2 marks)

   ii) What is the definition of Stackelberg incentive strategy? (3 marks)

   iii) What is the main difference between the above two types of strategies? (2 marks)

c) Consider the following Stackelberg game:

   - There are two players in which L is the leader and F is the follower, and the strategy spaces for the leader and the follower are U_L and U_F respectively;
   - The payoff functions for the leader and the follower are

   $$J_L(u_L, u_F) = (u_L - 1)(20 - \tfrac{13}{6}u_L + u_F)$$
   $$J_F(u_L, u_F) = (u_F - 1)(9 + u_L - 3u_F)$$

   in which $u_L \in U_L$ is the leader's strategy and $u_F \in U_F$ is the follower's strategy.

   Assume that the follower's strategy space is a continuous one given as $U_F = [1, +\infty)$. Now for each of the following cases, find a Stackelberg strategy:

   i) The leader's strategy space $U_L = [1, 11]$. (7 marks)

   ii) The leader's strategy space $U_L = [1, 5]$. (4 marks)

   (Note. For each case above, 1 mark is for the correct answer and the other marks are for the step by step reasoning and process.)

[Next page →]

4. a) Consider the single-item auction game in which the auctioneer's objective is to maximize social surplus by using a right mechanism. Now answer the following questions:

   i) List three desired properties of the 2nd price auction. (3 marks)

   ii) Which of these three properties is also true for the 1st price auction? (2 marks)

   iii) which of these three properties is not true for the 1st price auction? (2 marks)

b) Suppose that a group of companies on an island is considering whether to build a bridge to connect to the main land. The project cost is C. Each company has its private valuation about the project as $v_i(i = 1, 2, ..., I)$ and each has to share c_i to cover the cost if the bridge is decided to be build. Therefore the utility function for company i is $v_i - c_i$. The goal of the mechanism design is to enable every company to disclose their true valuation and build the bridge if the sum of each company's true valuation is larger than the cost. The designed mechanism for the project is as follows: 1) Allocation rule: If and only if $\sum_{i=1}^{I} Vi > C$ (where Vi is the reported valuation from company i), then the bridge is built; 2) Pricing rule: $c_i = \frac{C}{I}$ (i.e., the cost is even shared between the companies). Now answer the question: Whether such a mechanism is incentive compatible and why? (4 marks)

c) Consider the following multi-item auction game: There are 3 items, Items 1, 2 and 3, which include 20, 10, and 5 identical products respectively. There are 4 bidders, Bidders 1, 2, 3, and 4, whose valuation per product is 10, 5, 2, and 1 respectively. Assume that the auctioneer's objective is to maximise the social surplus function and his designed allocation rule is to allocate the i-th (i=1,2,3) biggest item to the i-th highest bidder. Now answer the following questions:

   i) Based on the given allocation rule, please give the corresponding pricing (i.e., payment) rule which ensures DSIC (Dominant Strategy Incentive Compatible). (2 marks)

   ii) In the given game, under the given allocation rule and its corresponding DSIC pricing rule, which bidder will win which item and how much does each bidder need to pay? (2 marks)

d) Consider the single item auction problem with the goal to maximise the social surplus (i.e., the value of the winner). It is assumed that there are two bidders whose valuation are two random variables v_1 and v_2 with the uniform distribution on interval $[0, 1]$. Now prove that the best bidding strategy (in Bayes-Nash equilibrium sense) for each bidder is "half of his valuation". (Hint: A random variable v with the uniformed distribution on [0,1] has the cumulative distribution function $F(z) = Pr(v < z) = z$.) (5 marks)

END OF EXAMINATION