Two hours

UNIVERSITY OF MANCHESTER
SCHOOL OF COMPUTER SCIENCE

Artificial Intelligence and Games

Thursday 22nd January 2011

Time: 14:00 – 16:00

Marker's feedback version

Please answer all FOUR Questions provided.

Use a SEPARATE answerbook for each SECTION.

Marks will be awarded for reasoning and method as well as being correct.

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

> **Marker's feedback**
> In general, students did very well on Question 1. The average was 67.2% $\pm$ 18%. Almost a quarter of the students got 80% or above on this question, so many of you knew this material very well.

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
	A	5	4	0	3	6
	B	4	4	-2	5	2
Player 1	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)

> **Marker's feedback**
> The vast majority of the students understood that if the maximum of the minimum of the rows was equal to the minimum of the maximum of the columns, a pair of strategies which achieved that value would be a Nash equilibrium. They got full marks. Some wrote it out mathematically, but I did not require that. The few people who got this wrong did not know what they were doing at all, for example saying that Player 1 should choose row D because it has the biggest numbers in it. Totally wrong; no points.

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 applying dominance to shrink the pay-off table as much as possible. (3 marks)

Marker's feedback
If you apply dominance over and over, you end up with a table with only one row (C) and one column (c) which is the Nash equilibrium. There are many paths to this. For example, you can eliminate rows A,B, and E, then every column except c, and then row D. Some stopped before all other strategies were eliminated and did not say how the problem was fully solved. This cost a point.

b) This question involves the following game.

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

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

Marker's feedback
It is a two-player game with no chance in normal form. It is a general sum game; i.e. *not* a zero sum. It is *not* a game of perfect information. Each player has to make its decision without knowledge of the other player's strategy. Normal form games are never games of perfect information. I will have to emphasise this more next year. Many people said the game was zero sum and/or perfect information. They lost one/two point(s).

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)

Marker's feedback
There are two: (A,a) and (B,b). Many people got this correct and could explain that if either player changed its strategy, that player would do worse. A few people said that only (B,b) was Nash, because the payoff of (A,a) is worse. This is wrong, but I gave a point for getting one of the equilibrium correct. An answer which come up a few times was: (A,b) and (B,a). This is worth zero and I do not understand why some people thought it was the right answer.

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)

Marker's feedback
There is a fully mixed Nash equilibrium and it is Alice plays **A** with probability $3/5$ and Bob plays **a** also with probability $3/5$. Many people got this right or approached it correctly but made an arithmetic error. Some people just skipped it. A few people argued that there could be no mixed-strategy equilibria because 2 is the maximum payoff, any mixed strategy would have a lower value, and that is impossible. Unfortunately, that would be right in a zero-sum game, but it is not right in a generally sum game, where different equilibria can have different values. Such an answer got one mark if it were clearly argued, zero otherwise.

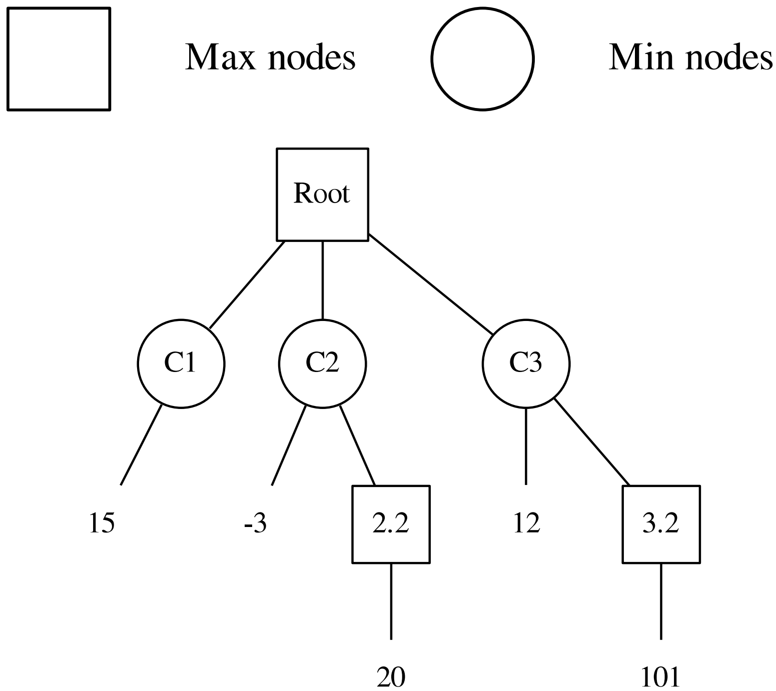

Figure 1: A game tree for Question 2a. Square nodes are Player 1 (MAX) decision nodes. Circle nodes are Player 2 (MIN) decision nodes. Numbers without shapes are payoffs. The numbers within the shapes are labels.

2. Extensive Form Games

> **Marker's feedback**
> Students did even better on this question. The mean was $71.5 \pm 19\%$ and 44% of the students made above 80%. Wow! Well done.

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)

> **Marker's feedback**
> Many students were able to get full marks on this. However, some strange values did seem to creep in. E.g. some said that 2.2 was pruned, but 20 came in as one of the alpha or beta values. Since there were eight values, each was worth $1/2$ mark unless the student seemed to devoid of any understanding and seemed to be guessing.
> A few people initialized [alpha beta] to [-3 101] since those are the minimum and maximal values that the nodes take. But, how could you possibly know that without exploring the entire tree, which completely defeats the purpose of alpha beta pruning, which is to *avoid* building the entire tree. So, that is totally wrong, and I could have given it 0. However, I just marked those values as wrong values.
> Most people seemed to know which children would not be evaluated, even when their alpha beta values would not lead to pruning.

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)

> **Marker's feedback**
> I tried to make this easy by having the first child immediately lead to a win. Many people identified that search could stop because a guaranteed winning path was found. Some did not justify it, which cost a point. Some people evaluated the tree in the wrong order, which required all children of the root to be evaluated, but some sub-trees could be pruned, which you needed to correctly identify to get marks.

c) These questions are about heuristics.

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

> **Marker's feedback**
> I accepted any reasonable heuristic. However, they needed to evaluate to numbers to get full credit. Some answers were more like concepts than heuristics. "have more stones on your side than the opponent's side"; "get an extra turn".

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)

> **Marker's feedback**
> Here, I was looking for two things. How would you determine which of two (say) heuristics was the better one? "Choose the better one" is not good enough. Second, if you were combining heuristics, how you change the relative amount of each one.

General Feedback to Part B (i.e., Question 3 and Question 4):

This part covers the teaching materials taught in the 2nd semester. The average mark for this part is 53%, which is slightly lower than 54.7% last year.

General Feedback to Question 3:

- 69 students took the exam and so answered this question.
- The average mark for this question is 64% (i.e., 9.6 marks out of 15). This is better than the last year, when the average mark for this question was 58.4%.
- 30 or 43 % students received a 1st class mark of 70% or better (i.e., 10.5 marks or more), which is much better than 28% of the last year. .
- 24 or 35% students received a 2nd class mark between 50%-69% (i.e., between 7.5 and 10 marks), which are lower than 48% of the last year.
- 9 or 13% students received a 3rd class mark between 40%-49% (i.e., between 6 and 7 marks), a little higher than 10.9 % of the last year.
- 6 or 8.7% students received a mark of less than 40% (i.e., 5 marks or fewer), this is lower than 12.5% last year.

Overall, the students' performance in this question is much better than the last year, as much more students get a 1st class mark and less students fails to get a pass mark. From the teaching outcome point of view, this is a better result than the one of the last year.

Detailed Feedback to Question 3:

- Question a). Except to a few students, most students know how to solve the given problem and answer this question well. The most common mistake is that no checking at the boundary strategies was done when finding the best global strategies. The second most common mistake is the incorrect calculation.
- Question b). This is a Stackelberg game problem on discrete strategy spaces, and most students (around 2/3) know how to answer this question. There are two common mistakes: One is using the derivative approach for the continuous strategy spaces to answer this question; and the other is to maximise the leader's payoff across the joint strategy spaces without taking into account the reaction mechanism in Stackelberg game.
- Question c). Almost all students have known the answer of this question, at least partially. The common mistake is that one option with the different algorithms is listed as different options. For example, learning the follower's reaction function and playing accordingly is just one option, but the detailed and different learning algorithms (such as the moving window and recursive least square algorithms) are just different methods to realise the learning and should not be regarded as different options.

General Feedback to Question 4:

- 69 students took the exam and so answered this question.
- The average mark for this question was 42% (i.e., 6.3 marks out of 15). This is worse than 51% of the last year.

- 12 or 17% students received a 1st class mark of 70% or better (i.e., 10.5 marks or more), in comparing with 26% in the last year.
- 18 or 26% students received a 2nd class mark between 50%-69% (i.e., between 7.5 and 10 marks), in comparing with 32% students in the last year.
- 10 or 15% students received a 3rd class mark between 40%-49% (i.e., i.e., between 6 and 7 marks), which is similar to 16% as the last year.
- 26 or 37.6% students received a mark of less than 40% (i.e., 5 marks or fewer), and mostly noticeable there are 11 students received 0 mark (largely empty answer), which is much worse than 26% students in the last year. Two possible reasons that so many students got 0 marks are 1) some students spent more time in other questions and left too little time for this question; 2) some students did not spent time to revise this part of materials.

General speaking, the students' performance in this question is not as good as expected and is worse than the last year. If taking off the 11 students with the 0 mark (in the case of the last year where the students were allowed to choose 3 out of 4 questions, these students were unlikely to choose to answer this question but this year every question is compulsory), the average this year is 50% which is largely similar to 51% last year. Despite such a similarity, this is certainly an area needed to improve in both teaching and exam paper setting.

Detailed Feedback to Question 4:

- Question a). More than half students answered this question reasonably well. The common mistake is the incompleteness or incorrect formula in their answer.
- Question b). There are 5 sub-questions in this question.
    - Question i): This sub-question is "In what sense that an allocation is regarded as feasible?" Most students failed to answer this question. This is less expected as this is a very basic conceptual type of question.
    - Questions ii)-iv): All these three sub-questions are conceptual type of questions and more than half of students are able to answer this question. The common mistakes are 1) the answers are incomplete; 2)the answers are inaccurate.
    - Question v). This is a proof type of question and was designed to be the difficult one. More than half students have some ideas how to prove this, but only a few students gave the complete proof.