Exam Performance Feedback Form

CS3191

2005/2006

General remarks: 118 students sat the exam. The result was very good this year, with an average mark of above 63%. In particular the first two questions, which can be practiced, received high average marks, and the marks for questions 4 and 5 were on a similar level as last year (which a slight dip in the latter due to the apparent suprise that I might ask something from the very last section of the notes).

There were a large number of first class marks this year, but also 16 fails, with the lowest taking just 8 marks out of sixty.

Question 1. All but three students attempted this question. All in all this question was well answered and achieved an average mark of above 64%. Only one student got a mark of five or lower, while 27 achieved a mark of 18 or higher.

Reasons why marks were lost typically were:

(a) Some students read the rules to say that if Alice goes low then the game is over. This is wrong. A few students did not give any pay-offs as part of the tree, or had some (or all!) pay-offs wrong. It seems to me that many of these mistakes were due to not reading the rules of the game carefully enough. Each of these mistakes resulted in the deduction of a mark, although I was generous this year with regard to information sets. A lot of students marked the draw, but did not put in which information Bob has before making a decision.

Very few students had the draw of the card *after* Alice and Bob had made their moves. This makes no sense and clearly does not agree with the description of the game.

(b) As in the past quite a few students neglected to describe proper strategies in that their strategies did not tell the player what to do in *all* the situations they might find themselves in.

There should be four strategies for Alice and four for Bob. In Alice's case, a strategy has to tell her what to do if she draws a red card, *and* what to do if she draws a black card. (So 'if a red card is drawn, go low' only describes *half* a strategy, not a full one. Some students treated the draw of a card as if Alice could decide whether she will get a black or a red card, which is clearly wrong. In Bob's case, the strategy has to tell him what to do if Alice goes high, *and* what to do if she goes low.

(c) People who had the wrong strategies typically had a much less complicated calculation to do for the matrix (probabilities didn't really feature there) so they lost a few marks here. The six marks were there because one had to calculate the 16 entries of the correct matrix—those who had fewer strategies for Bob and Alice only had a matrix half that size, and many of these did not require any calculations because students ignored the probabilities. More puzzling was that some students's numbers of strategies did not agree with the size of their matrix!

(d) The correct matrix has a mixed strategy equilibrium point which can be found using (pure strategy) dominance and solving the resulting symmetric 2×2 matrix. People who had wrong matrices with pure strategy equilibrium points lost marks here because they could read theirs off in 2 seconds, whereas those with the correct matrix had to do some calculations.

(f) Most students could determine the correct value for their game. However, a surprising number of students did not draw the correct conclusion about the fairness of the game (a two person zero-sum game is fair if and only if its value is 0, so that neither player wins or loses any money on average). More surprisingly, a lot of students did not want to play according to the equilibrium point strategy they had laboriously determined! All I was looking for here was a translation of the equilibrium point strategy for Alice back into the game, that is, on a red card go high or low with a probability of 1/2 each, and on a black card always go high.

Question 2. All but three students tackled this question. There were a few who had no idea what they were doing and got very low marks—9 students had a mark of 5 or lower. Given that there are 10 marks to be had for carrying out very simple algorithms (2a and 2c) which are *known* to appear in every exam I find this hard to explain. On the other hand, 44 students received a mark of 18 or higher, and the average mark for the question was 69.3%.

Reasons why some marks were lost typically were

(a) Some students forgot to give the value, losing a mark. A few missed out one or more of the four (pure strategy) equilibrium points for some reason, even fewer tried to solve this using dominance arguments—and as a result didn't find all pure strategy ones. (Note that I did not ask for mixed strategy equilibrium points and no marks were awarded for talking about those.) Some students didn't give the actual strategy pairs but only circled some entries in the matrix, which led to the deduction of a point, and some students were confused about the fact that in an equilibrium point for a 2-person zero-sum game the first strategy is always for the row player and the second for the column player. Four marks were available for the four equilibrium points, one for the correct value.

(b) Marks where lost when students did not give me a proof that a particular strategy could be eliminated, even if they picked the right one. A few students clearly were not able to deal with solving simultaneous inequalities, and removed strategies although they had not found a solution. They seemed to somehow convince themselves that they had. Four marks were given for the correct reduction of the matrix, one for the equilibrium point, and one for the value.

(c) I had meant to ask for the value as well but forgot to do so in the question, so all the marks were given for the equilibrium point. Those students who could solve the question by and large did so correctly. In a few cases marks were lost because

- Player 1 and Player 2 were confused;
- students made mistakes in solving the equations—a few of them received non-sensical results (a probability of 0, or a negative one) and still did not go back to check their work!

(d) There are five equilibrium points in this game, and many students did not find all of those. Two marks were given for finding all of them, one for finding at least some. As to the discussion, one mark was given if a sensible statement was made here, two marks if there was a sensible analysis of how all the players should behave.

Question 3. Only 6 students decided to answer this question, which was the hardest in the exam (and which had been announced as such a number of times). Three of these wrote some random things and received 2 marks or fewer, one didn't realize that he had only dealt with the simpler of two cases, one fell into a similar trap and made mistakes in that case. The remaining student solved the question almost completely, losing just two marks in the process.

In the serious attempts of solving the question the mistakes were

- not considering all the cases that the parameters in the 2×2 matrix might satisfy; one set of cases leads to a pure strategy equilibrium point, the other to a mixed strategy one.
- dividing by an expression and not making a special case for that expression being equal to 0.

Question 4. This question was attempted by 67 students. It had an average mark of almost 58%. Only one student received 18 marks or more, while five students got no more than 5 marks.

The one thing that future sitters of this exam should take to heart is that they should answer a question like this as if the reader didn't really know much about game playing programs. It certainly would help if students structured their writing, and if they could try to write legibly.

In part (a) very few students wrote anything at all about how the three parts work together, and some statements were wrong (the evaluation function does *not* evaluate moves, it evaluates positions). Quite a few students forgot that move generation goes together with the internal board representation. There were two marks to be had for each task (one of which was secured by merely naming it: board representation and move generation; evaluation function; alpha-beta search), and three marks for describing how it all fits together.

A lot of answers went into a lot of detail, but since there were only two marks to be had for each task this was a bit of a waste of time.

In part (b), some students only wrote down topics ('shape', 'mobility', etc) without explaining them, while quite a few explanations were plain wrong. Not many students bothered to say how these components might be put together, and those who believed that the evaluation function scores moves rather than positions only wrote half-truths for most of the items.

In part (c) I gave one mark for each sensible statement that was made regarding the game in question (until the four marks were used up). Some students did not write anything really specific to their chosen game, repeating what they had said under (b) and gaining no extra marks. Others tried to claim that all their points from (b) were really important and struggled to come up with sensible reasons for some of them.

Question 5. Fifty students attempted this question, and 6 of them had a mark of five or lower. Nine managed a mark of 18 or better. The average mark was 55.4%, which is disappointingly low.

Marks were typically lost for the following reasons.

(a) Not realizing that such systems are described in Section 6.6 of the notes, and not arriving a the correct matrix. which is

$$\begin{vmatrix} -40 & 40 \\ 0 & 10 \end{vmatrix}$$

(b) A lot of students did not realize that the answer is that a population consisting purely of Lions can be invaded by a single Lamb because a Lamb will receive 0 for each encounter, while the lions will lose 40 on average, and similarly for the other case.

(c) If the properties of being nice, retaliatory and forgiving were mentioned, a mark each was given. Another mark for a correct description for each. Mistakes were often made regarding niceness—this means never defecting first, not cooperating ‘at the start’, or on the first move. Similarly, being retaliatory does *not* require defecting immediately after the other side has defected, the player may choose to wait before dong so.

(d) This was a speculative question and as such marked generously. Two marks for an original choice of strategy, and three for giving good reasons. A lot of students chose variants of TitForTat, and in some cases their strategies didn’t differ apart from the last move. They ignored that it might not be known how many rounds there are, and so their submission would be TitForTat, which had been ruled out in the question. A lot of students wrote false statements into their justification (such as AlwaysD doing well in Axelrod’s tournaments, or claiming their strategies were nice when they weren’t).