Exam for IN4010 Artificial Intelligence Techniques

19 January 2011

This exam will test your knowledge and understanding of the material discussed in the second period of the course Artificial Intelligence Techniques. Using the book, notes, or slides during the examination is *not* allowed. You will have 3 hours (from 14 till 17) to complete the exam. It has 5 questions, for a total of 80 points. Please don't include irrelevant information: you will be marked down for this. Before you hand in your answers, please check that you have put your name and student number on top of every sheet you hand in.

Question 4 *10 points*

Consider the following Markov decision process (MDP) environment. The numbers in the cells represent the reward received in that cell. The cells with boxed rewards are terminal states. The player starts in the cell at the lower left corner of the environment. The player can (try to) move in all directions. The moves are subject to chance. The intended outcome occurs with probability 0.8, but with probability 0.2 the agent moves at right angles to the intended direction. A collision with a wall results in no movement.

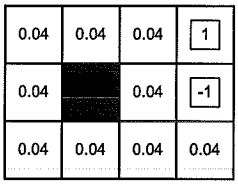

(a) (5 points) Present the policy with optimal expected utility in the form of arrows. Motivate your answer.

(b) (5 points) Explain with examples why the reward values are essential for the outcome of your answer in the previous question.

Question 5 *30 points*

Consider the following scenario: General Arnie and general Bernie have their enemy pinned down in the valley below them. War in their world is a profession from which money can be made or lost. They agreed to attack tomorrow at sunrise. Each of them realizes that they can only win the battle convincingly if they attack at the same time. Such an attack will cost them 100 soldiers of each of their own armies in casualties, and 50 of each army's noblemen will be captured. Both armies will kill 100 soldiers from the enemy army and capture 100 noblemen.

However, if one army (say X) attacks about an hour later than the other army (say Y) then the enemy is expected to kill 150 of the army of Y, and capture 100 of their troupes. The army of Y will only kill 50 and capture 50 of the enemy army. Of the army of X that attacked later, the enemy will kill 50 soldiers and capture 30 noblemen. The army of X will kill 100 and capture 130 of the enemy army.

Making captives will mean income later when the family of the captives will buy them free. Having your own people captured is a costly affair because captives will have to be bought free later on, or exchanged for enemy captives that you made. (You may assume that the ransom will be about the same for all captives.) The dead have to be buried, but this is a relatively cheap affair, that you may ignore. Both armies have a practically infinite supply of new soldiers since there are so many poor people and the army would provide a steady income for them. So the number of casualties does not influence the chances of the army in the future. If Arnie and Bernie both wait, they face the costs of maintaining their armies and make no profit.

The main question is, will Arnie and Bernie stick to their agreement to attack at the same time? For this you have to analyse their pay-off matrices. General Arnie asked his consultant, who came up with the following:

	B: attack	B: wait
A: attack	A: 20 B: 30	A: 50 B: -10
A: wait	A: 0 B: 50	A: -10 B: -10

Obviously the consultant is wrong, so you will have to do better.

(a) (5 points) Construct a pay-off matrix for this problem that reflects the information above. Motivate your calculations.

However brilliant your answer to question (a), for the rest of the questions you are to work with the poor calculations of the consultant.

(b) (5 points) What is the definition of a dominant strategy? Does one exist for the pay-off matrix given by the consultant?

(c) (5 points) Which of the outcomes are Pareto optimal? Motivate your answer.