Assignment Eight: Sequential Decisions

Due: 11:59pm, Monday 23 November 2020.

Solving following problems requires using the `http://aispace.org` belief and decision networks applet (you may need to download the jar file) and/or the Python code at `http://aipython.org/aipython_322.zip` (which now includes reasoning with uncertainty and planning with uncertainty).

This can be done in groups of size 1, 2 or 3. Working alone is not recommended. All members of the group need to be able to explain the group's answer.

Submit your answers in pdf using Canvas. Use proper sentences in your answers.

Ask questions on Canvas discussion board. Feel free to answer them too.

Question One

Students have to make decisions about how much to study for each course. Different students have different utilities. Suppose that the utility depends on their total effort and their final grade (A, B, C or F). Assume that total effort is a binary variable that measures their subjective effort.

(a) Give an appropriate utility function for a student who is lazy and just wants to pass (not get an F). The total effort here measures whether they (thought they) worked a lot or a little overall. (Use 100 for the best outcome and 0 for the worst outcome.) You can fill in the following table:

Grade	Total Effort	Utility
A	Lot	
A	Little	
B	Lot	
B	Little	
C	Lot	
C	Little	
F	Lot	
F	Little	

(b) Give an appropriate utility function for a student who doesn't mind working hard and really wants to get an A, and would be very disappointed with a B or lower. (Use 100 for the best outcome and 0 for the worst outcome. Fill in a table similar to part (a).)

Question Two

Some students choose to cheat on exams, and instructors want to make sure that cheating does not pay. A rational model would specify that the decision of whether to cheat depends on the costs and the benefits. Here we will develop and critique such a model.

Consider the decision network of Figure 1. This diagram models a student's decisions about whether to cheat at two different times. If students cheat they can be caught cheating, but they can also get higher grades. The punishment (either suspension, cheating recorded on the transcript, or none) depends on whether they get caught at either or both opportunities. Whether they get caught depends on whether they are being watched and whether they cheat. The utility depends on their final grades and their punishment.

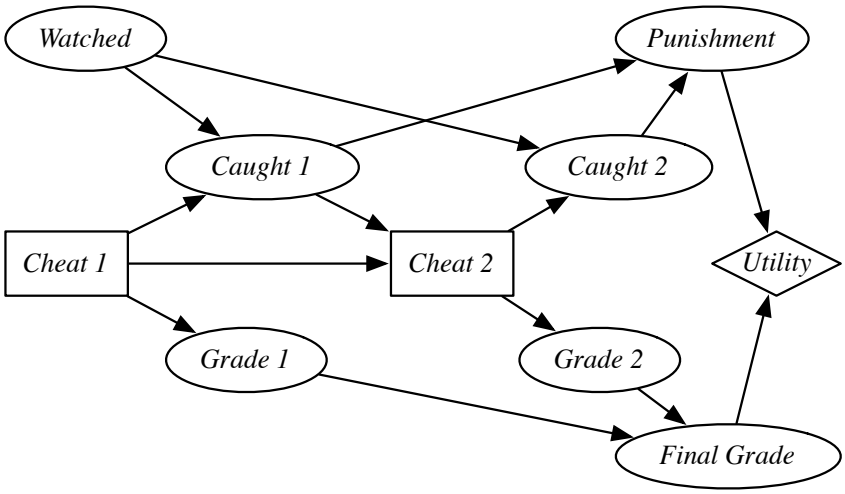

Figure 1: Cheat Decision Network

Consider the example `https://artint.info/code/aispace/cheat_decision.xml` for the AISpace Belief and Decision network tool. (Do a "Load from URL" in the "File" menu). This is also *cheat_dn* in decnNetworks.py in the Python distribution.

(a) What is an optimal policy? Give a description in English of an optimal policy. (The description should not use any jargon of AI or decision theory.) What is the value of an optimal policy?

(b) What happens to the optimal policy when the probability of being watched goes up? [Modify the probability of "Watched" in create mode.] Try a number of values. Explain what happens and why.

(c) What is an optimal policy when the rewards for cheating are reduced? Try a number of different parametrizations.

(d) Change the model so that the once a student has been caught cheating, they will be watched more carefully. [Hint: whether they are watched at the first opportunity needs to be a different variable than whether they are watched at the second opportunity.] Show the resulting model (both the structure and any new parameters), and give the policies and expected utilities for various settings of the parameters.

(e) Suppose the university decided to set up an honour system so that instructors do not actively check for cheating, but there is severe punishments for first offences if cheating is discovered. How could this be modelled? Specify a model for this and explain what decision it is optimal for the student to do (for a few different parameter settings).

Question Three

Students have to make decisions about how much to study for each course. The aim of this question is to investigate how to use decision networks to help them make such decisions.

Suppose students first have to decide how much to study for the midterm. They can study a lot, study a little, or not study at all. Whether they pass the midterm depends on how much they

study and on the difficulty of the course. As a first-order approximation, they pass if they study hard or if the course is easy and they study a bit. After receiving their midterm grade, they have to decide how much to study for the final exam. Again, the final exam result depends on how much they study and on the difficulty of the course. Their final grade depends on which exams they pass; generally they get an A if they pass both exams, a B if they only pass the final, a C if they only pass the midterm, or an F if they fail both. Of course, there is a great deal of noise in these general estimates.

As in Question 1, the depends on their total effort and their final grade. Suppose the total effort is obtained by combining the effort in studying for the midterm and the final. The total effort is a binary variable that measures their subjective effort.

Draw a decision network for a student decision based on the preceding story. (Show the random, decision and value nodes, but you don't need to specify the domains of the variables, or any of the probabilities or utilities.)

Question Four

For each question, specify how long you spend on it, and what you learned. How was the work in the team allocated? Was the question reasonable?