CPSC 322 **Artificial Intelligence I** **Fall 2020**

Practice Final Exam

The exam will be worth 120 marks and take 150 minutes. This means that we expect you should spend one minute per mark (and have 30 minutes at the end to check it over). Each question below has an associated mark, so you can calibrate your timing. Short answer questions require at most two sentences to explain each point, unless they specify otherwise.

You may bring in one sheet of letter-sized paper with anything written on it.

The final will cover (some of):

- AI and agents, including what is AI, definition of an agent, and the agent design dimensions.
- States and searching.
- CSPs: search, arc consistency with domain splitting, local search.
- Planning (forward, regression) and representations (STRIPS)
- Reasoning under Uncertainty: probability, conditioning, belief networks, variable elimination, sequential models, sampling
- Decision making: utility, decisions, policies, variable elimination, value of information and control, decision processes

It will cover sections of the textbook that are listed under the schedule tab of the course web page.

You can expect to have questions that arise from the assignments, so look at all of the questions and solutions. See also the practice questions from the two midterms. Here we give questions from after the second midterm.

Some important points (that students often forget):

- **Read the question and answer what is asked**. You will not get marks for writing things (whether they are true or not) that are not relevant to the question. Writing extra stuff that is not relevant cannot help you, and will result in a loss of marks if it is wrong.
- Use proper English in full sentences. You will not get marks if we cannot work out what you are saying.
- If a question asks about a particular instance of a problem, make sure your answer refers to that instance. Writing a general formula that you may have copied from the Internet, is not worth any marks. (The questions are usually asking to apply a formula to a particular case, to make sure you understand it).
- If the question asks for a definition, give the definition in your own words. Don't plagiarize from the textbook. (We are not testing your ability to copy from your notes.) If it doesn't ask for a definition, don't give one. Answer the question specifically to the case at hand, not a generic case. (If someone says "what is on the table?" you don't define what "on the table" means, nor do you tell them "all of the things in the room that are not on the floor except for the chair and the table", you tell them "there is a cup and two books").

Probability

Question 1

(a) [4 marks] Suppose Sam had a camper van (a mobile home) and liked to keep it at a comfortable temperature and noticed that the energy use depended on the elevation. Sam knows that the elevation affects the outside temperature. Sam likes the camper warmer at higher elevation for some quirky reason. [Note that not all of the variables directly affect electrical usage.] Show how this can be represented as a causal network, by adding arcs to the following Bayesian network:

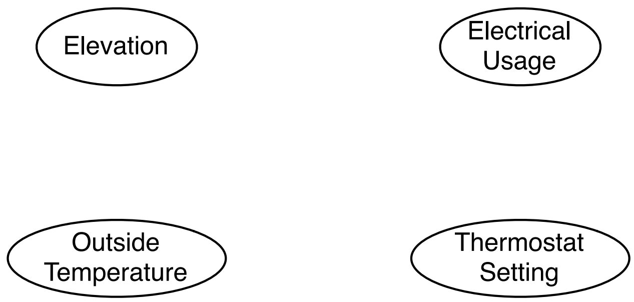

(b) [4 marks] Give an example where intervening has an effect different from conditioning for this network.

(c) [5 marks] In designing the Google car, the designers could use the images in Google Street View to find the current location by finding the best match to the current sensor information. They could instead, use knowledge of the starting location and the steering of the car to keep track of the current position. What is wrong with each of these? Explain how hidden Markov models can use both sorts of information and do better than either.

(d) [5 marks] Suppose that it is possible that the probability of a sequence of observations is less than $2 * 10^{-308}$ (the smallest number that can be represented on common computers), which is not unusual when there are multiple sensors readings at each time. Which of the following algorithms suffers from underflow (real numbers that are too small to be represented using double precision floats), and which will provide no samples in practice: rejection sampling, importance sampling, particle filtering? Explain why.

Decisions

Question 2

Consider the decision network (from assignment 8)

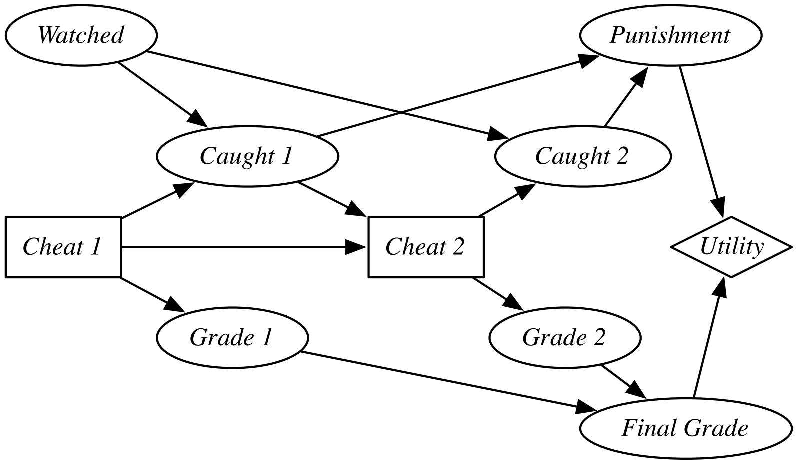

(a) [3 marks] Explain in English what the arc from "Cheat 1" to "Caught 1" means.

(b) [3 marks] Explain in English what the arc from "Cheat 1" to "Cheat 2" means.

(c) [4 marks] In using variable elimination to find the optimal policy, which variables are summed out before the first decision is eliminated?

(d) [4 marks] In using variable elimination to find the optimal policy, which decision variable is eliminated first? What does the decision function created specify?

(e) [2 marks] If there was an arc from "Cheat 1" to "Watched", what would that mean about the world? (What must be true about the world for this arc to be appropriate?)

(f) [2 marks] If there was an arc from "Watched" to "Cheat 2", what would that mean about the world? (What must be true about the world for this arc to be appropriate?)

(g) [2 marks] If an arc was added from "Watched" to "Cheat 2", how could the expected utility of the optimal policy be affected?

(h) [2 marks] What must be true about the optimal policy if the expected utility was affected by adding the arc from "Watched" to "Cheat 2"?

Question 3

Suppose that in a decision network, the decision variable *Run* has parents *Look* and *See*. Suppose that we are using variable elimination to find the optimal policy. Suppose that after eliminating all of the other variables, we have the factor

Look	See	Run	value
true	true	yes	23
true	true	no	8
true	false	yes	37
true	false	no	56
false	true	yes	28
false	true	no	12
false	false	yes	18
false	false	no	22

(a) [4 marks] What is the resulting factor after eliminating the decision variable *Run*?

(b) [4 marks] What is the optimal decision function for *Run*?

(c) [6 marks] Suppose that, in a decision network, there were arcs from random variables "Contaminated Specimen" and "Positive Test" to the decision variable "Discard Sample". You solved the decision network and discovered that there was a unique optimal policy:

contaminated specimen	positive test	discard sample
true	true	yes
true	false	no
false	true	yes
false	false	no

What can you say about the value of information in this case? [You will only get full marks for a precise statements contained in full sentences.]

Question 4

In a nuclear research submarine, a sensor measures the temperature of the reactor core. An alarm is triggered ($A = true$) if the sensor reading is abnormally high ($S = true$), indicating an overheating of the core ($C = true$). The alarm and/or the sensor can be defective ($A\_ok = false$, $S\_ok = false$) which can cause them to malfunction. The alarm system can be modelled by the following belief network (all variables are Boolean):

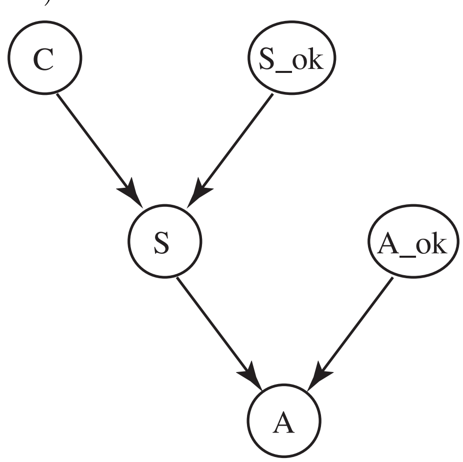

(a) [6 marks] Suppose we add a second, identical sensor to the system and trigger the alarm when either of the sensors reads a high temperature. The two sensors break and fail independently. Give the extended belief network. (Draw the graph and specify any new conditional probabilities).

(b) [6 marks] When an alarm is observed, a decision is made whether to shut down the reactor. Shutting down the reactor has a cost c_s associated with it (independent of whether the core was overheating), while not shutting down an overheated core incurs a cost c_m much higher than c_s. Draw the decision network modelling this decision problem for the original system (i.e., only one sensor). Specify any new tables that need to be defined (you should use the parameters c_s and c_m where appropriate in the tables). You can assume that the *utility* is the negative of *cost*.

(c) [6 marks] For the decision network in part (b), suppose you need to compute the optimal policy, and the value of the optimal policy using variable elimination. Show, for one legal elimination ordering, what variables are eliminated, what factors are combined, and what factors are created. [Just give the variables in the factors, not the tables of numbers. Specify whether the variables are summed out or maximized.]

(d) [3 marks] How can the optimal policy and its expected value be extracted from the elimination of part (c).

(e) [3 marks] Explain how to compute the value of the information of C for the decision of whether to shut down the reactor.