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

Practice Final Exam

Solution

Probability

Question 1

Solution

(a) This one fits the story

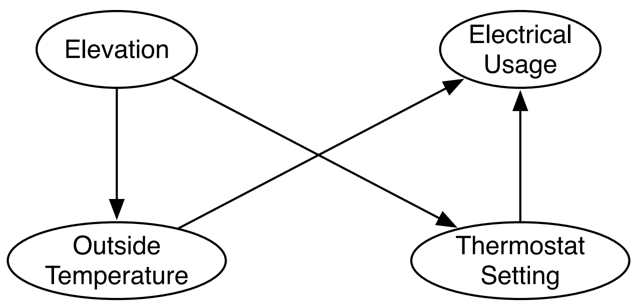

(b) Changing the thermostat setting affected the electrical usage, but does not change the elevation or the outside temperature. Observing the thermostat setting changes our belief in the elevation, and all of the other variables.

(c) There are too many scenes that look alike to be able to determine the current location from the view. And its difficult to find one corresponding to the current postion as they are only taken at 10m or so intervals. Steering is not accurate enough to be able to determine where one is without observing the world, and the errors accumulate. In an HMM, we estimate the current location from the previous location, the action and the sensors. There are far fewer possible images, and the image lets us adjust the position information so the errors don't accumulate.

(d) Importance sampling suffers from underflow. Most particles will have a probability too small to represent. Rejection sampling will end up with no particles; they will all be rejected. Particle filtering can work well.

Decisions

Question 2

Solution

(a) Whether the student is caught at the first time depends on whether they actually cheated at the first time.

(b) Action Cheat 1 occurs before Cheat 2, and the agent remembered what it did.

(c) *Watched*, *Punishment*, *Caught_2*, *Grade_1*, *Grade_2* and *Final_grade*. (All of the random variables except *Caught_1*).

(d) *Cheat_2* is eliminated first. The decision function specifies what the agent should do as a function of *Cheat_1* and *Caught_1*.

(e) The examiner could watch depending on whether the student cheated or not at the first opportunity.

(f) The student would know if they were watched when deciding whether to cheat at the second opportunity.

(g) It would either go up or stay the same. It can't go down.

(h) In the optimal policy, the agent will do something different at the second cheating opportunity depending on whether *Watched* is true. (E.g., the student may cheat only if they are not being watched).

Question 3

Solution

(a)

Look	See	value
true	true	23
true	false	56
false	true	28
false	false	22

(b)

Look	See	Run
true	true	yes
true	false	no
false	true	yes
false	false	no

That is, if the agent sees, it should run.

(c) The value of "contaminated specimen" for "discard sample" is 0. If we know the value of positive test, finding whether the specimen is contaminated gives no added value (because we don't change the action).

The value of "positive test" for "discard sample" is greater than 0. If we know the value of contaminated, finding whether test is postive gives added value because we change the decision (it is important in the description that there is a unique optimal policy).

Question 4

Solution

(a)

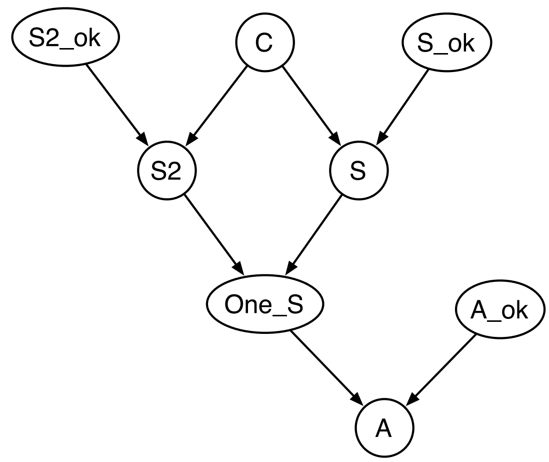

$P(S2 \mid S2\_ok, C)$ is the same as $P(S \mid S\_ok, C)$.
$P(S2\_ok)$ is the same as $P(S\_ok)$.
$P(One\_S = true \mid S = false, S2 = false) = 0$ and $P(One\_S = true \mid S, S2) = 1$ for other values of $S, S1$.

(b)

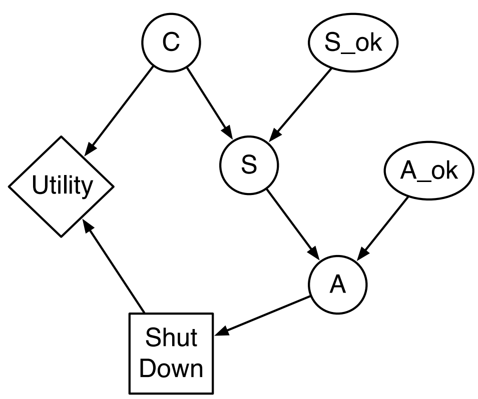

The only new table is the utility:

C	*ShutDown*	Value
true	true	$-c_s$
true	false	$-c_m$
false	true	$-c_s$
false	false	0

(c) The following is one possible elimination sequence (the first 4 eliminations can be done in any order):

Eliminate S_ok: multiply $P(S \mid C, S\_ok)$ and $P(S\_ok)$, and sum out S_ok giving $f_1(S, C)$.

Eliminate A_ok: multiply $P(A \mid S, A\_ok)$ and $P(A\_ok)$, and sum out S_ok giving $f_2(A, S)$.

Eliminate C: multiply $utility(C, ShutDown)$, $P(C)$ and $f_1(S, C)$, and sum out C, giving $f_3(S, ShutDown)$.

Eliminate S: multiply $f_2(A, S)$ and $f_3(S, ShutDown)$ and sum out S giving $f_4(A, ShutDown)$

Eliminate *ShutDown*: maximize *ShutDown* in $f_4(A, ShutDown)$, giving $f_5(A)$

Eliminate A: sum out A in $f_5(A)$

(d) When eliminating *Shutdown* record which values of *Shutdown* are maximal for each value of A.

The value obtained from summing out A is the expected utility of the optimal policy.

(e) Return the value of the optimal policy with an arc from C to shutdown, minus the value just computed above.