Probability, Bayes' Nets and Decision Networks

It is Monday night, and Bob is finishing up preparing for the CS188 Midterm II that is coming up on Tuesday. Bob has already mastered all the topics except one: Decision Networks. He is contemplating whether to spend the remainder of his evening reviewing that topic (*review*), or just go to sleep (*sleep*). Decision Networks are either going to be on the test ($+d$) or not be on the test ($-d$). His utility of satisfaction is only affected by these two variables as shown below:

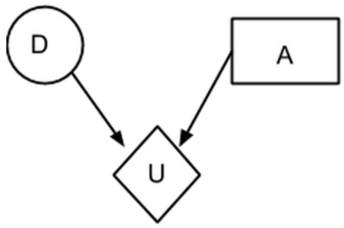

D	P(D)
+d	0.5
-d	0.5

D	A	U(D,A)
+d	*review*	1000
-d	*review*	600
+d	*sleep*	0
-d	*sleep*	1500

(a) Maximum Expected Utility

Compute the following quantities:

$EU(review) = P(+d)U(+d, review) + P(-d)U(-d, review) = 0.5 * 1000 + 0.5 * 600 = 800$

$EU(sleep) = P(+d)U(+d, sleep) + P(-d)U(-d, sleep) = 0.5 * 0 + 0.5 * 1500 = 750$

$MEU(\{\}) = max(800, 750) = 800$

Action that achieves $MEU(\{\}) = review$

This result notwithstanding, you should get some sleep.

(b) The TA is on Facebook

The TAs happiness (H) is affected by whether decision networks are going to be on the exam. The happiness (H) determines whether the TA posts on Facebook ($+f$) or doesn't post on Facebook ($-f$). The prior on D and utility tables remain unchanged.

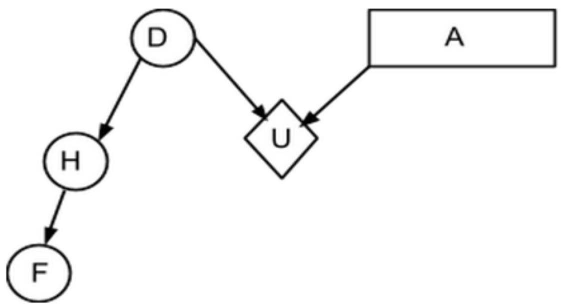

Decision network.

F	H	$P(F\|H)$
+f	+h	0.6
-f	+h	0.4
+f	-h	0.2
-f	-h	0.8

H	D	$P(H\|D)$
+h	+d	0.95
-h	+d	0.05
+h	-d	0.25
-h	-d	0.75

D	P(D)
+d	0.5
-d	0.5

D	A	U(D,A)
+d	*review*	1000
-d	*review*	600
+d	*sleep*	0
-d	*sleep*	1500

Tables that define the model are shown above.

H	$P(H)$
+h	0.6
-h	0.4

F	$P(F)$
+f	0.44
-f	0.56

D	F	$P(D\|F)$
+d	+f	0.666
-d	+f	0.334
+d	-f	0.370
-d	-f	0.630

F	D	$P(F\|D)$
+f	+d	0.586
-f	+d	0.414
+f	-d	0.300
-f	-d	0.700

D	H	$P(D\|H)$
+d	+h	0.79
-d	+h	0.21
+d	-h	0.06
-d	-h	0.94

Tables computed from the first set of tables. Some of them might be convenient to answer the questions below.

Compute the following quantities:

$EU(review|+f) = P(+d|+f)U(+d, review) + P(-d|+f)U(-d, review) = 0.666 * 1000 + 0.334 * 600 = 666 + 200.4 = 866.4$

$EU(sleep|+f) = P(+d|+f)U(+d, sleep) + P(-d|+f)U(-d, sleep) = 0.666 * 0 + 0.334 * 1500 = 501$

$MEU(\{+f\}) = \max(866.4, 501) = 866.4$

Optimal Action($\{+f\}$) = *review*

$EU(review|-f) = P(+d|-f)U(+d, review) + P(-d|-f)U(-d, review) = 0.370 * 1000 + 0.630 * 600 = 370 + 378 = 748$

$EU(sleep|-f) = P(+d|-f)U(+d, sleep) + P(-d|-f)U(-d, sleep) = 0.370 * 0 + 0.630 * 1500 = 0 + 945 = 945$

$MEU(\{-f\}) = \max(748, 945) = 945$

Optimal Action($\{-f\}$) = *sleep*

$VPI(\{F\}) = P(+f)MEU(\{+f\}) + P(-f)MEU(\{-f\}) - MEU(\{\}) = 0.44 * 866.4 + 0.56 * 945 - 800 = 110.416$

(c) VPI Comparisons

Now consider the case where there are n TAs. Each TA follows the same probabilistic models for happiness (H) and posting on Facebook (F) as in the previous question.

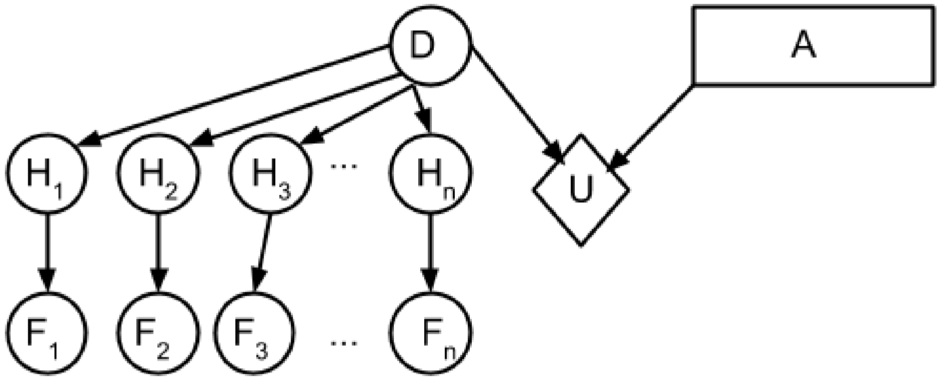

(i) ○True ●False $VPI(H_1|F_1) = 0$

Justify: F_1 is just a noisy version of H_1. Hence finding out H_1 gives us more information about D even when we have already observed F_1. This in turn will allow us to more often make the right decision between *sleep* and *review*.

(ii) ●True ○False $VPI(F_1|H_1) = 0$

Justify: The parent variable of the utility node, D, is conditionally independent of F_1 given H_1.

(iii) ○True ●False $VPI(F_3|F_2, F_1) > VPI(F_2|F_1)$

Justify: The F_i variables give us noisy information about D. The more F_i variables we get to observe, the better chance we end up being able to make the right decision. The more F_i variables we have already observed, however, the less an additiona observation of a new variable F_j will influence the distribution of D.

(iv) ●True ○False $VPI(F_1, F_2, \ldots, F_n) < VPI(H_1, H_2, \ldots, H_n)$

Justify: The F_i variables are noisy versions of the H_i variables, hence observing the H_i variables is more valuable.