Question 1 (Class) – Value of Information

[Adapted from problem 16.11 in Russell & Norvig]

A used-car buyer can decide to carry out various tests with various costs (e.g., kick the tires, take the car to a qualified mechanic) and then, depending on the outcome of the tests, decide which car to buy. We will assume that the buyer is deciding whether to buy car *c* and that there is time to carry out at most one test which costs \$50 and which can help to figure out the quality of the car.

A car can be in good shape (of good quality Q=q) or in bad shape (of bad quality Q=~q), and the test might help to indicate what shape the car is in. There are only two outcomes for the test T: pass (T=pass) or fail (T=fail). Car *c* costs \$1,500, and its market value is \$2,000 if it is in good shape; if not, \$700 in repairs will be needed to make it in good shape. The buyer's estimate is that *c* has 70% chance of being in good shape.

a) Draw the decision network that represents this problem. (The only action is whether to buy or not the car).

*The decision network is in blue.

The rest is supplementary information.

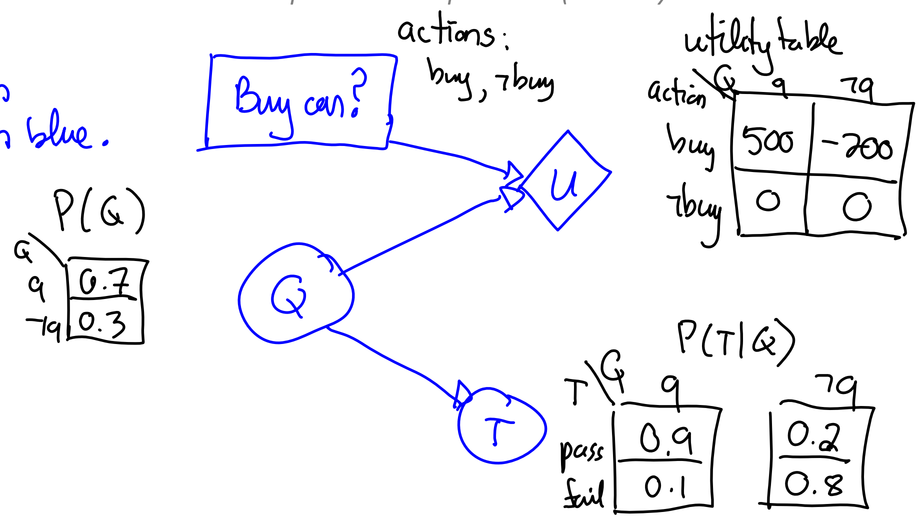

b) Calculate the expected net gain from buying car c, given no test.

$$EU(buy \mid \underset{\text{no info}}{\emptyset}) = P(Q=q) \cdot U(q, buy) + P(Q=\neg q)\, U(\neg q, buy)$$

$$= 0.7 \cdot 500 + 0.3 \cdot (-200) = 290$$

c) Tests can be described by the probability that the car will pass or fail the test given that the car is in good or bad shape. We have the following information:

P(T=pass | Q=q) = 0.9

P(T=pass | Q=~q) = 0.2

[Note that the numbers were changed from the book. It is also a useful exercise to redo this problem with the numbers from the book.]

$$P(Q=q \mid T=fail) = \frac{P(T=fail \mid Q=q)\,P(Q=q)}{P(T=fail)} = \frac{0.1 \cdot 0.7}{0.31} = \boxed{0.22}$$

Calculate the probability that the car will pass (or fail) its test, and then the probability that it is in good (or bad) shape given each possible test outcome [Hint: Bayes' theorem could come handy here.]

$$P(T=pass) = \sum_{v \in \{q, \neg q\}} P(T=pass, Q=v) = \sum_v P(T=pass \mid Q=v)P(Q=v) = P(T=pass \mid Q=q)P(Q=q) + P(T=pass \mid Q=\neg q)P(Q=\neg q)$$

$$= 0.9 \cdot 0.7 + 0.2 \cdot 0.3 = \boxed{0.69}$$

$$P(T=fail) = 1 - P(T=pass) = 1 - 0.69 = \boxed{0.31}$$

$$P(Q=q \mid T=pass) \underset{\text{[Bayes' thm]}}{=} \frac{P(T=pass \mid Q=q)P(Q=q)}{P(T=pass)} = \frac{0.9 \cdot 0.7}{0.69} = \boxed{0.91}$$

d) Calculate the optimal decisions given either a pass or a fail, and their expected utilities.

$$EU(buy \mid T=pass) = P(Q=q \mid T=pass)\,U(q, buy) + P(Q=\neg q \mid T=pass)\,U(\neg q, buy)$$
$$= 0.91 \cdot 500 + 0.09 \cdot (-200) = 437$$

$$EU(\neg buy \mid T=pass) = EU(\neg buy \mid T=fail) = 0$$

$$EU(buy \mid T=fail) = P(Q=q \mid T=fail)\,U(q, buy) + P(Q=\neg q \mid T=fail)\,U(\neg q, buy)$$
$$= 0.22 \cdot 500 + 0.78 \cdot (-200) = -46$$

so $\boxed{MEU(T=pass) = 437 \text{ (using buy) and } MEU(T=fail) = 0 \text{ (using } \neg buy)}$

e) Calculate the value of (perfect) information of the test. Should the buyer pay for a test?

$$VPI_{\emptyset}(T) = \sum_{t \in \{pass, fail\}} P(T=t \mid \emptyset)\,MEU(T=t) - MEU(\emptyset)$$

value of perfect information for variable T without previous evidence ($e=\emptyset$)

S

with rounding errors...

$$= 0.69 \cdot 437 + 0.31 \cdot (0) - 290 \approx \boxed{11.53}$$

The buyer shouldn't buy the test since its value is less than the cost...

Note: $P(s|e,e')\cdot P(e'|e) = P(s|e)$ like $P(s|e')\cdot P(e') = P(s,e')$

but more precisely
$$P(s|e,e')\cdot P(e'|e) = \frac{P(s,e,e')}{\cancel{P(e,e')}}\cdot\frac{\cancel{P(e',e)}}{P(e)} = P(s,e'|e)$$

Question 2 (Homework) – Properties of VPI

a) Prove that if the optimal action to do is not changed for any value of the new evidence E', then the value of information is 0.

[Use the formula for VPI from the slides of lecture 19]

Let a^* be the optimal action (which is unchanged by knowing $E'=e'$)

then $MEU(e,e') = \max_a \sum_s P(s|e,e')U(s,a) = \sum_s P(s|e,e')U(s,a^*)$

so
$$VPI_e(E') = \sum_{e'} P(e'|e)\, MEU(e,e') - MEU(e)$$
$$= \sum_{e'} P(e'|e)\Big(\sum_s P(s|e,e')\, U(s,a^*)\Big) - MEU(e)$$
interchange
$$= \sum_s U(s,a^*) \underbrace{\sum_{e'} \underbrace{P(e'|e)P(s|e,e')}_{P(s,e'|e)}}_{P(s|e)} - MEU(e) = MEU(e) - MEU(e) = 0$$
as required //

b) Prove that the value of perfect information is always non-negative (i.e. $VPE_e(E_j) \geq 0$ for any j and e).

[Hint: use the fact that the sum of maximums is always greater or equal than the maximum of a sum.]

$$VPE_e(E_j) = \sum_{e'} P(e'|e) \max_{a_{e'}} \sum_s P(s|e,e')U(s,a_{e'}) - MEU(e)$$

shorthand for $P(E_j=e'|E=e)$

in general, argmax could depend on e'

now, using the fact that $\sum_i a_i \max_{x_i} f_i(x_i) \geq \max_x \sum_i a_i f_i(x)$ for any a_i, f_i

we have that

interchange again
$$VPE_e(E_j) \geq \max_a \sum_{e'} P(e'|e) \sum_s P(s|e,e')U(s,a) - MEU(e)$$
$$= \max_a \sum_s \sum_{e'} \underbrace{P(e'|e)\, P(s|e,e')}_{P(s,e'|e) \text{ again}} U(s,a) - MEU(e)$$
$$= MEU(e) - MEU(e) = 0$$
so $VPE_e(E_j) \geq 0$ //

Question 3 (Homework) – HMM and particle filtering

You sometimes get colds, which make you sneeze. You also get allergies, which make you sneeze. Sometimes you are well, which doesn't make you sneeze. You decide to model the process using the following HMM, with hidden states *X* ∈ {*well, allergy, cold*} and observations *E* ∈ {*well, allergy, cold*}:

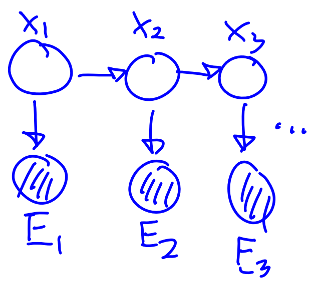

$P(X_1)$

well	1
allergy	0
cold	0

$P(X_t \mid X_{t-1} = well)$

well	0.7
allergy	0.2
cold	0.1

$P(X_t \mid X_{t-1} = allergy)$

well	0.6
allergy	0.3
cold	0.1

$P(X_t \mid X_{t-1} = cold)$

well	0.2
allergy	0.2
cold	0.6

Transitions

$P(E_t \mid X_t = well)$

quiet	1.0
sneeze	0.0

$P(E_t \mid X_t = allergy)$

quiet	0.0
sneeze	1.0

$P(E_t \mid X_t = cold)$

quiet	0.0
sneeze	1.0

Emissions

Note that colds are "stickier" in that you tend to have them for multiple days, while allergies come and go on a quicker time scale. However, allergies are more frequent. Assume that on the first day, you are well.

a) Imagine that you observe the sequence quiet, sneeze, sneeze. *What is the probability that you were well all three days and observed these effects?*

$P(X_1=well, X_2=well, X_3=well \mid E_1=quiet, E_2=sneeze, E_3=sneeze) = 0$

since it is proportional to joint which contains $P(E_2=sneeze \mid X_2=well)=0$ in the product

b) What is the posterior distribution over your state on day 2 (X_2) if E_1 = quiet, *E_2 =* sneeze*?*

This is filtering: $P(x_2 \mid e_{1:2}) \propto P(x_2, e_{1:2}) = \sum_{x_1} P(x_1, x_2, e_{1:2})$

$= \sum_{x_1} P(x_2 \mid x_1) P(x_1) P(e_1 \mid x_1) P(e_2 \mid x_2) = P(X_2 \mid X_1=well) \overbrace{P(X_1=well)}^{1} \overbrace{P(E_1=q \mid X_1=well)}^{1} P(E_2=s \mid X_2)$

since $P(X_1=cold) = P(X_1=allergy) = 0$

[see next page for cont.]

[cont. of b):]

so $P(X_2 \mid E_1=q, E_2=s] \propto \begin{array}{c|c} & X_2 \\ w & 0.7 \cdot 0 \\ a & 0.2 \cdot 1 \\ c & 0.1 \cdot 1 \end{array}$

i.e. $P(\; " \quad " \;) = \begin{array}{c|c} & X_2 \\ w & 0 \\ a & 2/3 \\ c & 1/3 \end{array}$

c) What is the posterior distribution over your state on day 3 (X_3) if E_1 = quiet, E_2 = sneeze, E_3 = sneeze?

from HMM filtering update:

$$p(x_3 \mid e_{1:3}) \propto \underbrace{p(E_3=s \mid x_3)}_{\text{new evidence}} \sum_{x_2} P(x_3 \mid x_2)\, \overbrace{p(x_2 \mid e_{1:2})\, \overbrace{p(e_{1:2})}^{\text{constant}}}^{p(x_2,\, e_{1:2})}$$

computed in b)

$$= \begin{array}{c} w \\ a \\ c \end{array}\begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \odot \begin{array}{c|ccc} x_3 \backslash x_2 & w & a & c \\ \hline w & 0.7 & 0.6 & 0.2 \\ a & 0.2 & 0.3 & 0.2 \\ c & 0.1 & 0.1 & 0.6 \end{array} \begin{bmatrix} 0 \\ 2/3 \\ 1/3 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \odot \begin{bmatrix} 0.47 \\ 0.27 \\ 0.27 \end{bmatrix} = \begin{bmatrix} 0 \\ 0.27 \\ 0.27 \end{bmatrix}$$

renormalize to get:

$$P(X_3 \mid e_{1:3}) = \begin{bmatrix} 0 \\ 0.5 \\ 0.5 \end{bmatrix}$$

d) What is the Viterbi (most likely) sequence for the observation sequence quiet, sneeze, sneeze, sneeze, quiet, quiet, sneeze, quiet, quiet? [Hint: you should not have to do extensive calculations.]

renormalize for convenience (*) renormalize

observations:	q	s	s	s	q	q	s	q	q
well	1	0	0	0	1	1	0	1	1
allergy	0	0.2	0.06	0.018	0	0	0.2	0	0
cold	0	0.1	0.06	0.036	0	0	0.1	0	0

entries

example of viterbi update: $\propto m_{1:4}(x_4) = \max_{x_{1:3}} P(x_{1:4} \mid e_{1:4}) \propto P(e_4 \mid x_4) \max_{x_3} P(x_4 \mid x_3) \cdot m_{1:3}(x_3)$

Imagine you are monitoring your state using the particle filtering algorithm, and on a given day you have 5 particles on *well*, 2 on *cold*, and 3 on *allergy* before making an observation on that day.

(*) we can renormalize without changing result since everything is defined within a proportionality constant

solution: (following back pointer from max of $m_{1:9}(x_9)$)

$x^*_{1:9} = (w, a, a, a, w, w, a, w, w)$

e) If you observe sneeze, *what weight will each of your particles have?*

weight → $w(x_t = \text{well}) = P(e_t = \text{sneeze} \mid x_t = \text{well}) = 0$

$w(x_t = \text{cold}) = P(\;" \mid x_t = \text{cold}) = 1$

$w(x_t = \text{allergy}) = P(\;" \mid x_t = \text{allergy}) = 1$

f) After resampling, what is the expected number of particles you will have on cold?

- total weight on cold was $w(x_t = \text{cold}) \cdot N_{cold} = 2$
- total weight of all particles was $0 \cdot 5 + 2 \cdot 1 + 3 \cdot 1 = 5$
- so prob. of sampling cold is $2/5$
- you sample 10 particles, so expected # of particles on cold is $\frac{2}{5} \cdot 10 = \boxed{4}$