CS 188 Fall 2015 | Introduction to Artificial Intelligence | Section Handout 8 HMMs, VPI

Q1. HMMs: Tracking a Jabberwock

You have been put in charge of a Jabberwock for your friend Lewis. The Jabberwock is kept in a large tugley wood which is conveniently divided into a 10×10 grid. It wanders freely around the 100 possible cells. At each time step $t = 1, 2, 3, \ldots$, the Jabberwock is in some cell $X_t \in \{1, \ldots, 10\}^2$, and it moves to cell X_{t+1} randomly as follows: with probability 0.5, it chooses one of the (up to 4) valid neighboring cells uniformly at random; with probability 0.5, it uses its magical powers to teleport to a random cell uniformly at random among the 100 possibilities (it might teleport to the same cell). It always starts in $X_1 = (1, 1)$.

(a) Compute the probability that the Jabberwock will be in $X_2 = (2, 1)$ at time step 2. What about $\Pr(X_2 = (4, 4))$?

$P(X_2 = (2, 1)) =$ $1/2 \cdot 1/2 + 1/2 \cdot 1/100 = 0.255$

$P(X_2 = (4, 4)) =$ $1/2 \cdot 1/100 = 0.005$

(b) At each time step t, you dont see X_t but see E_t, which is the row that the Jabberwock is in; that is, if $X_t = (r, c)$, then $E_t = r$. You still know that $X_1 = (1, 1)$. Suppose we see that $E_1 = 1$, $E_2 = 2$ and $E_3 = 10$. Fill in the following table with the distribution over X_t after each time step, taking into consideration the evidence. Your answer should be concise. *Hint*: you should not need to do any heavy calculations.

t	$P(X_t \mid e_{1:t-1})$	$P(X_t \mid e_{1:t})$
1	(1,1) : 1.0, (others) : 0.0	(1,1): 1.0, (others): 0.0
2	(1,2), (2,1): 51/200, (others): 1/200	(2,1): 51/60, (2,2..10): 1/60

(c) These images correspond to probability distributions of the Jabberwock's location. Match them with the most appropriate probabilities from this list $P(X_1), P(X_1|e_1), P(X_2|e_1), P(X_2|e_1, e_2), P(X_3|e_1, e_2), P(X_3|e_1, e_2, e_3)$

P(X 2|e 1)

P(X 1)

P(X 3|e 1,e 2,e 3)

P(X 2|e 1,e 2)

P(X 3|e 1,e 2)

P(X 1|e 1)

Q2. VPI: Used Car Purchase

[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 buyers estimate is that c has 70% chance of being in good shape.

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

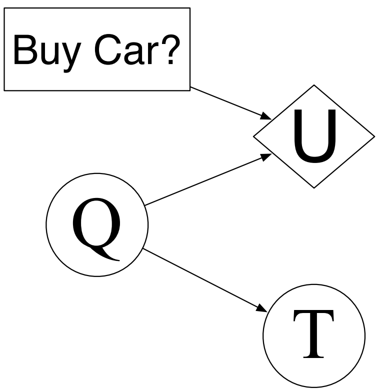

2. Calculate the expected net gain from buying car c, given no test.

$$\begin{aligned} EU(\text{buy}) &= P(Q = +q) \cdot U(+q, \text{buy}) + P(Q = \neg q) \cdot U(\neg q, buy) \\ &= .7 \cdot 500 + 0.3 \cdot -200 = 290 \end{aligned}$$

3. 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:

$$P(T = \text{pass}|Q = q) = 0.9$$
$$P(T = \text{pass}|Q = \neg q) = 0.2$$

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.

$$\begin{aligned}
P(T = \text{pass}) &= \sum_q P(T = \text{pass}, Q = q) \\
&= P(T = \text{pass}|Q = +q)P(Q = +q) + P(T = \text{pass}|Q = \neg q)P(Q = \neg q) \\
&= 0.69 \\
P(T = fail) &= 0.31 \\
P(Q = +q|T = \text{pass}) &= \frac{P(T = \text{pass}|Q = +q)P(Q = +q)}{P(T = \text{pass})} \\
&= \frac{0.9 \cdot 0.7}{0.69} \approx 0.91 \\
P(Q = +q|T = \text{fail}) &= \frac{P(T = \text{fail}|Q = -q)P(Q = -q)}{P(T = \text{fail})} \\
&= \frac{0.1 \cdot 0.7}{0.31} \approx 0.22
\end{aligned}$$

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

$$\begin{aligned}
EU(\text{buy}|T = pass) &= P(Q = +q|T = \text{pass})U(+q, \text{buy}) + P(Q = \neg q|T = \text{pass})U(\neg q, \text{buy}) \\
&\approx 0.91 \cdot 500 + 0.09 \cdot (-200) \approx 437 \\
EU(\text{buy}|T = \text{fail}) &= P(Q = +q|T = \text{fail})U(+q, \text{buy}) + P(Q = \neg q|T = \text{fail})U(\neg q, \text{buy}) \\
&\approx 0.22 \cdot 500 + 0.78 \cdot (-200) = -46 \\
EU(\neg\text{buy}|T = pass) &= 0 \\
EU(\neg\text{buy}|T = fail) &= 0
\end{aligned}$$

Therefore: $MEU(T = \text{pass}) = 437$ (with buy) and $MEU(T = \text{fail}) = 0$ (using $\neg$buy)

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

$$\begin{aligned}
VPI(T) &= (\sum_t P(T = t)MEU(T = t)) - MEU(\phi) \\
&= 0.69 \cdot 437 + 0.31 \cdot 0 - 290 \approx 11.53
\end{aligned}$$

You shouldn't pay for it, since the cost is \$50.

6. The value of the information in this problem depends greatly on the prior probability $P(Q = q)$. What do you think to the VPI as your as you vary $P(Q = q)$? What happens when $P(Q = q)$ approaches 1? Approaches 0? Approaches 0.5?

Intuitively, the more unsure you are of something, the more value you stand to gain by getting "perfect information" related to it. Thus, we expect that near 1 and near 0, there is little value in the information, but somewhere in the middle there is higher value. Utilities mess with this intuition: if you have an action that is low risk (not buying) and one that is high risk (buying), the midpoint might shift. Thus, as you might see in the next part, the critical range where getting the test is useful lies somewhere between 0.18 and 0.52.

7. If you still have time, try calculating some VPI's for different values of $P(Q = q)$ from the previous part. (Using Python or Excel is a good idea here!) Where is/are the break-even point(s)? Where is the maximum?

The break-even points are somewhere around 0.52 and 0.18. The maximum is around 0.3.