

Intelligent Systems II

Exercise Sheet 2

Stuart Golodetz

February 9, 2006

1. Recall the Viterbi algorithm from the book and the notes (set 3):

$$m_{1:t+1} = \mathbf{P}(e_{t+1}|X_{t+1}) \max_{x_t} (\mathbf{P}(X_{t+1}|x_t)m_{1:t}) \quad (1)$$

where

$$m_{1:t} = \max_{x_1, \dots, x_{t-1}} \mathbf{P}(x_1, \dots, x_{t-1}, X_t | e_{1:t})$$

Assuming that probability distributions (such as $m_{1:t}$) are computed as S -component column vectors in the usual way (where S is the number of possible values of each X_t), write down a short piece of pseudo-code for computing the recurrence relation (1). (Hint: normal matrix operations will not suffice.) How should we initialise the recurrence?

Hence, or otherwise, go back to the Umbrella World example and compute the vectors $m_{1:n}$ for $1 \leq n \leq 4$.

Answer

First we define the operation $\overset{\max}{\times}$. If $A = (a_{i,j})$ is an $m \times n$ matrix, $B = (b_{i,j})$ is an $n \times p$ matrix and $C = (c_{i,j}) = A \overset{\max}{\times} B$ is an $m \times p$ matrix, then:

$$c_{i,j} = \max_{k=1}^n a_{i,k} b_{k,j}$$

Now if O_t and T are the sensor and transition models in matrix form, as defined in the textbook, then:

$$m_{1:t+1} = \alpha O_{t+1} \left(T \overset{\max}{\times} m_{1:t} \right)$$

We initialise the recurrence with $m_{1:1} = \mathbf{P}(X_1|e_1) = \alpha \mathbf{P}(e_1|X_1) \mathbf{P}(X_1)$.

Calculations

Renaming X as R and E as U as usual, we calculate as follows:

$$\begin{aligned}
 \mathbf{P}(R_1) &= \langle 0.5, 0.5 \rangle && \{\text{as usual}\} \\
 m_{1:1} &= \alpha \mathbf{P}(u_1 | R_1) \langle 0.5, 0.5 \rangle && \{U_1 = \text{true from data}\} \\
 &= \alpha \langle 0.9, 0.2 \rangle \langle 0.5, 0.5 \rangle \\
 &= \langle 0.8182, 0.1818 \rangle
 \end{aligned}$$

$$T = \begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix}$$

$$\begin{aligned}
 m_{1:2} &= \alpha O_2 \left(T^T \times^{\max} \langle 0.8182, 0.1818 \rangle \right) \\
 &= \alpha \begin{pmatrix} 0.1 & 0 \\ 0 & 0.8 \end{pmatrix} \left[\begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix} \times^{\max} \begin{pmatrix} 0.8182 \\ 0.1818 \end{pmatrix} \right] \quad \{O_2\text{'s that } \cdot : U_2 = \text{false}\} \\
 &= \alpha \begin{pmatrix} 0.1 & 0 \\ 0 & 0.8 \end{pmatrix} \begin{pmatrix} 0.57274 \\ 0.24546 \end{pmatrix} \\
 &= \alpha \begin{pmatrix} 0.057274 \\ 0.196368 \end{pmatrix} \\
 &= \begin{pmatrix} 0.2258 \\ 0.7742 \end{pmatrix}
 \end{aligned}$$

$$\begin{aligned}
 m_{1:3} &= \alpha O_3 \left(T^T \times^{\max} \langle 0.2258, 0.7742 \rangle \right) \\
 &= \alpha \begin{pmatrix} 0.9 & 0 \\ 0 & 0.2 \end{pmatrix} \left[\begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix} \times^{\max} \begin{pmatrix} 0.2258 \\ 0.7742 \end{pmatrix} \right] \quad \{O_3\text{'s that } \cdot : U_3 = \text{true}\} \\
 &= \alpha \begin{pmatrix} 0.9 & 0 \\ 0 & 0.2 \end{pmatrix} \begin{pmatrix} 0.2323 \\ 0.5419 \end{pmatrix} \\
 &= \alpha \begin{pmatrix} 0.2091 \\ 0.1084 \end{pmatrix} \\
 &= \begin{pmatrix} 0.6586 \\ 0.3414 \end{pmatrix}
 \end{aligned}$$

$$\begin{aligned}
 m_{1:4} &= \alpha O_4 \left(T^T \times^{\max} \langle 0.6586, 0.3414 \rangle \right) \\
 &= \alpha \begin{pmatrix} 0.9 & 0 \\ 0 & 0.2 \end{pmatrix} \left[\begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix} \times^{\max} \begin{pmatrix} 0.6586 \\ 0.3414 \end{pmatrix} \right] \quad \{O_4\text{'s that } \cdot : U_4 = \text{true}\} \\
 &= \alpha \begin{pmatrix} 0.9 & 0 \\ 0 & 0.2 \end{pmatrix} \begin{pmatrix} 0.4610 \\ 0.2390 \end{pmatrix} \\
 &= \alpha \begin{pmatrix} 0.4149 \\ 0.0478 \end{pmatrix} \\
 &= \begin{pmatrix} 0.8967 \\ 0.1033 \end{pmatrix}
 \end{aligned}$$

2. Consider the candy-wrapper example discussed in lectures (set 4), with the same starting conditions, namely

- 10% are h_1 : 100% C
- 20% are h_2 : 75% C + 25% L
- 40% are h_3 : 50% C + 50% L
- 20% are h_4 : 25% C + 75% L
- 10% are h_5 : 100% L

and again the bags of candy are assumed large enough that the probability of taking any type of candy from a bag does not alter as we start to empty the bag.

For a particular bag we extract the sequence of candies C, L, C.

- (a) What are now the probabilities of each hypothesis holding for this bag?
- (b) What is the probability that the next candy will be L?
- (c) What is the most likely hypothesis according to the maximum likelihood criteria?

Answer

- (a) We calculate as follows:

$$\begin{aligned}
 & P(h_1|\mathbf{d}) \\
 = & \alpha P(\mathbf{d}|h_1)P(h_1) \\
 = & \alpha \left(\prod_j P(d_j|h_1) \right) P(h_1) \\
 = & \alpha (P(C|h_1)P(L|h_1)P(C|h_1)) \times 0.1 \\
 = & \alpha \times 1 \times 0 \times 1 \times 0.1 \\
 = & 0
 \end{aligned}$$

$$\begin{aligned}
 & P(h_2|\mathbf{d}) \\
 = & \alpha (0.75 \times 0.25 \times 0.75) \times 0.2 \\
 = & \alpha \times 0.028125
 \end{aligned}$$

$$\begin{aligned}
 & P(h_3|\mathbf{d}) \\
 = & \alpha (0.5 \times 0.5 \times 0.5) \times 0.4 \\
 = & \alpha \times 0.05
 \end{aligned}$$

$$\begin{aligned}
 & P(h_4|\mathbf{d}) \\
 = & \alpha (0.25 \times 0.75 \times 0.25) \times 0.2 \\
 = & \alpha \times 0.009375
 \end{aligned}$$

$$\begin{aligned}
 & P(h_5|\mathbf{d}) \\
 = & 0
 \end{aligned}$$

Now we normalise:

$$\begin{aligned}
 1 &= \sum_{i=1}^5 P(h_i|\mathbf{d}) \\
 &= 0 + \alpha \times 0.028125 + \alpha \times 0.05 + \alpha \times 0.009375 + 0 \\
 &= \frac{7}{80}\alpha \\
 \Leftrightarrow \alpha &= \frac{80}{7}
 \end{aligned}$$

Whence:

$$\begin{aligned}P(h_1|\mathbf{d}) &= 0 \\P(h_2|\mathbf{d}) &= \frac{9}{28} \\P(h_3|\mathbf{d}) &= \frac{4}{7} \\P(h_4|\mathbf{d}) &= \frac{3}{28} \\P(h_5|\mathbf{d}) &= 0\end{aligned}$$

(b) We now want to calculate as follows:

$$\begin{aligned}\mathbf{P}(d_4 = L | \mathbf{d}) &= \sum_{i=1}^5 \mathbf{P}(d_4 = L | h_i) P(h_i|\mathbf{d}) \\&= \langle 0, 1 \rangle \times 0 + \langle 0.25, 0.75 \rangle \times \frac{9}{28} + \langle 0.5, 0.5 \rangle \times \frac{4}{7} + \langle 0.75, 0.25 \rangle \times \frac{3}{28} + \langle 1, 0 \rangle \times 0 \\&= \left\langle \frac{9}{112}, \frac{27}{112} \right\rangle + \left\langle \frac{4}{14}, \frac{4}{14} \right\rangle + \left\langle \frac{9}{112}, \frac{3}{112} \right\rangle \\&= \left\langle \frac{25}{56}, \frac{31}{56} \right\rangle\end{aligned}$$

$$\text{So } P(d_4 = L | \mathbf{d}) = \frac{25}{56}.$$

(c) We want h_{ML} maximising $P(\mathbf{d}|h_i)$. Well, we have:

$$\begin{aligned}P(\mathbf{d}|h_1) &= 1 \times 0 \times 1 = 0 \\P(\mathbf{d}|h_2) &= 0.75 \times 0.25 \times 0.75 = 0.140625 \\P(\mathbf{d}|h_3) &= 0.125 \\P(\mathbf{d}|h_4) &= 0.25 \times 0.75 \times 0.25 = 0.046875 \\P(\mathbf{d}|h_5) &= 0\end{aligned}$$

So the most likely hypothesis is h_2 . This is what we'd expect just by looking at the data, since we're getting 66% cherry and 33% lime, which is closest to h_2 .

3. A random process X generates integers according to a Poisson distribution, which generates an integer n with probability

$$P(n) = e^{-\nu} \frac{\nu^n}{n!}$$

for some unknown parameter ν (and where $n! = n(n-1)(n-2) \cdots 2 \times 1$). 100 samples are taken of the output of X , which generates the integers 0-9 according to the frequency distribution shown in the table in Figure 1. Derive the log-likelihood function for this distribution, and thus estimate the value of ν .

0	1	2	3	4	5	6	7	8	9
2	5	11	27	22	16	9	5	2	1

Figure 1: Distribution table for 100 samples from process X

Answer

The hypothesis here is h_ν , so the log-likelihood is:

$$L(\mathbf{d}|h_\nu) = \sum_{j=1}^{100} \log P(d_j|h_\nu)$$

Now:

$$\begin{aligned} P(X = n|h_\nu) &= e^{-\nu} \frac{\nu^n}{n!} \\ \Leftrightarrow \log P(X = n|h_\nu) &= \log e^{-\nu} \frac{\nu^n}{n!} \\ &= \log e^{-\nu} + \log \frac{\nu^n}{n!} \\ &= -\nu + \log \nu^n - \log n! \\ &= -\nu + n \log \nu - \log n! \\ \Leftrightarrow \frac{\partial \log P(X=n|h_\nu)}{\partial \nu} &= -1 + \frac{n}{\nu} \end{aligned}$$

So:

$$\begin{aligned} L(\mathbf{d}|h_\nu) &= 2 \log P(X = 0|h_\nu) + 5 \log P(X = 1|h_\nu) + 11 \log P(X = 2|h_\nu) + \\ & 27 \log P(X = 3|h_\nu) + 22 \log P(X = 4|h_\nu) + 16 \log P(X = 5|h_\nu) + \\ & 9 \log P(X = 6|h_\nu) + 5 \log P(X = 7|h_\nu) + 2 \log P(X = 8|h_\nu) + \\ & 1 \log P(X = 9|h_\nu) \end{aligned}$$

And:

$$\begin{aligned} \frac{\partial L(\mathbf{d}|h_\nu)}{\partial \nu} &= 2(-1) + 5\left(-1 + \frac{1}{\nu}\right) + 11\left(-1 + \frac{2}{\nu}\right) + 27\left(-1 + \frac{3}{\nu}\right) + 22\left(-1 + \frac{4}{\nu}\right) + \\ & 16\left(-1 + \frac{5}{\nu}\right) + 9\left(-1 + \frac{6}{\nu}\right) + 5\left(-1 + \frac{7}{\nu}\right) + 2\left(-1 + \frac{8}{\nu}\right) + 1\left(-1 + \frac{9}{\nu}\right) \\ &= -100 + \frac{390}{\nu} \end{aligned}$$

Setting this equal to 0 gives $\nu = \frac{390}{100} = 3.9$. If we now tabulate $100 \times P(n)$, this does seem to correlate fairly well with the data:

n	$100 \times P(n)$	$round(100 \times P(n))$
0	2.024191145	2
1	7.894345464	8
2	15.39397365	15
3	20.01216575	20
4	19.51186161	20
5	15.21925205	15
6	9.892513835	10
7	5.511543422	6
8	2.686877418	3
9	1.164313548	1

4. For the question, see the problem sheet.

Answer

Considering only the x direction (since the y direction is analogous), suppose we're at μ_x and the variance is σ_x^2 . The distribution for our position is $N(\mu[x]_t, \sigma[x]_t^2)$, that for the x update is $N(2, 0.3)$ and that for the sensor (if available) is $N(sonar[x]_{t+1}, 0.1)$. There are two cases to consider:

- If we don't have the sensor available, we just compute the basic update:

$$N(\mu[x]_t, \sigma[x]_t^2) + N(2, 0.3) = N(\mu[x]_t + 2, \sigma[x]_t^2 + 0.3)$$

(Intuitively, all we do is shift the distribution along by 2 and spread it out a bit.)

- If we do have it available, we combine it with the above, i.e.

$$\text{combine}(N(\mu[x]_t + 2, \sigma[x]_t^2 + 0.3), N(sonar[x]_{t+1}, 0.1))$$

(Here, by contrast, we have another observation from the sensor so we combine the two.)

The end result is another Gaussian distribution, call it $N(\mu[x]_{t+1}, \sigma[x]_{t+1}^2)$. In the first case, for example, we have that:

$$\begin{aligned}\mu[x]_{t+1} &= \mu[x]_t + 2 \\ \sigma[x]_{t+1}^2 &= \sigma[x]_t^2 + 0.3\end{aligned}$$

In the second case, we have that:

$$\begin{aligned}\mu[x]_{t+1} &= \frac{0.1 \times (\mu[x]_t + 2) + (\sigma[x]_t^2 + 0.3) \times sonar[x]_{t+1}}{\sigma[x]_t^2 + 0.3 + 0.1} \\ \sigma[x]_{t+1}^2 &= \frac{(\sigma[x]_t^2 + 0.3) \times 0.1}{\sigma[x]_t^2 + 0.3 + 0.1}\end{aligned}$$

The values for $t = 1$, then, are calculated as follows:

$$\begin{aligned} & N(\mu[x]_{t+1}, \sigma[x]_{t+1}^2) \\ &= \text{combine}(N(1, 0) + N(2, 0.3), N(2.90, 0.1)) \\ &= \text{combine}(N(3, 0.3), N(2.90, 0.1)) = N(2.925, 0.075) \end{aligned}$$

$$\begin{aligned} & N(\mu[y]_{t+1}, \sigma[y]_{t+1}^2) \\ &= \text{combine}(N(1, 0) + N(1, 0.3), N(2.08, 0.1)) \\ &= \text{combine}(N(2, 0.3), N(2.08, 0.1)) = N(2.06, 0.075) \end{aligned}$$

We hence fill in the rest of the table as follows:

$$\begin{aligned} N(\mu[x]_2, \sigma[x]_2^2) &= N(4.925, 0.375) \\ N(\mu[x]_3, \sigma[x]_3^2) &= N(6.925, 0.675) \\ N(\mu[x]_4, \sigma[x]_4^2) &= N(8.925, 0.975) \\ N(\mu[x]_5, \sigma[x]_5^2) &= N(10.925, 1.275) \\ N(\mu[x]_6, \sigma[x]_6^2) &= N(12.925, 1.575) \\ N(\mu[x]_7, \sigma[x]_7^2) &= \text{combine}(N(14.925, 1.875), N(14.75, 0.1)) = N(14.7589, 0.0949) \\ N(\mu[x]_8, \sigma[x]_8^2) &= \text{combine}(N(16.7589, 0.395), N(16.80, 0.1)) = N(16.7917, 0.0798) \end{aligned}$$

$$\begin{aligned} N(\mu[y]_2, \sigma[y]_2^2) &= \text{combine}(N(3.06, 0.375), N(3.12, 0.1)) = N(3.107, 0.0789) \\ N(\mu[y]_3, \sigma[y]_3^2) &= N(4.107, 0.3789) \\ N(\mu[y]_4, \sigma[y]_4^2) &= N(5.107, 0.6789) \\ N(\mu[y]_5, \sigma[y]_5^2) &= N(6.107, 0.9789) \\ N(\mu[y]_6, \sigma[y]_6^2) &= \text{combine}(N(7.107, 1.2789), N(7.30, 0.1)) = N(7.286, 0.0927) \\ N(\mu[y]_7, \sigma[y]_7^2) &= \text{combine}(N(8.286, 0.3927), N(8.35, 0.1)) = N(8.337, 0.0797) \\ N(\mu[y]_8, \sigma[y]_8^2) &= \text{combine}(N(9.337, 0.3797), N(9.28, 0.1)) = N(9.292, 0.0792) \end{aligned}$$

The displacement errors probably are as bad as they were estimated since the variances for the final position are quite small. However, it may not matter that much given the context, since the robot probably doesn't have to be precisely positioned. It's within a foot in each direction of where it was supposed to be and that might well be good enough for what we want it to do. (For an example: suppose the robot is a robot dog and has been told to go over to the other side of the room and wag its tail at someone...a foot each way would be close enough. By contrast, suppose the robot is really a surgery laser...if it was known to sometimes be a foot out, would *you* take the risk?!)