Universität des Saarlandes
Fachrichtung 6.2 – Informatik
J. Siekmann, S. Autexier, K. Fischer, E. Melis, C. Zinn
M. Schiller, E. Schulz, R.Vollmann

11th Theoretical Assignment in Artificial Intelligence (WS 2006/2007)
Solutions

Exercise 11.1

In a galaxy far, far away, 90% of the taxicabs are green and 10% are blue. An accident involving a taxicab occurs; we presume that the accident rate for green taxicabs is the same as for blue taxicabs. A court considers the accident, and a newspaper reporter who was at the scene says, "The cab involved at the scene was blue". This reporter is usually reliable; in fact, his statements are correct 80% of the time. That is, if the taxicab involved in the accident was in fact blue (or green), the probability that our witness would say "blue" (or "green") is 0.8. What is the probability that the taxicab involved in the accident was blue, given the reporter's statement? (from Nils J. Nilsson: Artificial Intelligence, 1998)

Solution:

- *We identify the prior probabilities that a taxi is* `green` *or* `blue`*:*

$$P(green) = 0.9 \text{ and } P(blue) = 0.1$$

- *We also know the conditional probabilites for the reports given by the observer:*

$$P(\text{observer-says-blue}|blue) = 0.8 \text{ and } P(\text{observer-says-green}|green) = 0.8 \text{ and}$$

$$P(\text{observer-says-blue}|green) = 0.2 \text{ and } P(\text{observer-says-green}|blue) = 0.2$$

- *The probability that the taxi was indeed blue if it is reported to be blue is given by*

$$P(blue|\text{observer-says-blue}) = \frac{P(\text{observer-says-blue}|blue)P(blue)}{P(\text{observer-says-blue})} \quad \text{(Bayes' rule)}$$

- *Now there are two possible ways to calculate the results.*

Option 1:

– *We calculate*

$$\begin{aligned} P(\text{observer-says-blue}) &= P(\text{observer-says-blue}\&green) + P(\text{observer-says-blue}\&blue) \\ &= P(\text{observer-says-blue}|green) * P(green) \\ &\quad + P(\text{observer-says-blue}|blue) * P(blue) \\ &= 0.2 * 0.9 + 0.8 * 0.1 \\ &= 0.26 \end{aligned}$$

– *We conclude:*

$$P(blue|\text{observer-says-blue}) = \frac{P(\text{observer-says-blue}|blue)P(blue)}{P(\text{observer-says-blue})} \quad \textit{(Bayes' rule)}$$
$$= \frac{0.8 * 0.1}{0.26}$$
$$= 0.3077$$

Option 2:

– *We also calculate*

$$P(\neg blue|\text{observer-says-blue}) = \frac{P(\text{observer-says-blue}|negblue)P(blue)}{P(\text{observer-says-blue})}$$
$$= \frac{0.2 * 0.9}{P(\text{observer-says-blue})}$$

– *We know that*

$$P(blue|\text{observer-says-blue}) + P(\neg blue|\text{observer-says-blue})$$
$$= \frac{0.08}{P(\text{observer-says-blue})} + \frac{0.18}{P(\text{observer-says-blue})}$$
$$= 1$$

Therefore, we can calculate that our normalization constant α is 3.846 (without having to calculuate $P(\text{observer-says-blue})$). Therefore, $P(\neg blue|\text{observer-says-blue}) = 3.846 * 0.08 = 0.3077$

- *Note: Even though the observer is pretty reliable, the chance that he is correct in this example is quite low (30%) – because the prior probability that a taxi is blue is very low.*

Exercise 11.2

Three prisoners A, B and C are locked up in their cells. Everybody knows that one of them will be put to death the next morning and the two others will be pardoned. The decision who will be put to death has already been made by the governor. The prisoner A asks the guard a favour: "Please ask the governor who will be put to death tomorrow and tell one of the other two prisoners that he will be pardoned."

Assume that the guard accepts and returns to tell A to whom of his fellow prisoners he passed the pardon message. To represent that situation, we use to variables $Execute$ and $GuardSays$. The possible values of $Executed$ are $\{A, B, C\}$, the possible values for $GuardSays$ are $\{B, C\}$. The joint probability distribution for the situation is

	$Execute = A$	$Execute = B$	$Execute = C$
$GuardSays = B$	$\frac{1}{6}$	0	$\frac{1}{3}$
$GuardSays = C$	$\frac{1}{6}$	$\frac{1}{3}$	0

(a) Why are there two values $\frac{1}{6}$ in each row of the column for $Execute = A$, while for the other two columns there is one row with $\frac{1}{3}$ and the other is 0? (Hint: Compute $P(GuardSays = B|Execute = A)$ and $P(GuardSays = C|Execute = A)$).

Solution:

In case A is executed, the guard has a choice between telling B is pardoned and telling C is pardoned. Hence, the probability of $\frac{1}{3}$ of A being executed is split equally between these.

Computing the two values resulted in $\frac{1}{2}$, which was supposed to give you the indication that the probabilities are split equally.

$$\begin{aligned} P(GuardSays = B|Execute = A) &= \frac{P(GuardSays = B \wedge Execute = A)}{P(Execute = A)} \\ &= \frac{\frac{1}{6}}{\frac{1}{3}} \\ &= \frac{1}{2} \end{aligned}$$

(b) What is the conditional probability of A being executed the next day given B is not executed the next day? (justify your answer mathematically!)

Solution:

$$\begin{aligned} P(Execute = A|\neg(Execute = B)) &= \frac{P(Execute = A \wedge \neg(Execute = B))}{P(\neg(Execute = B))} \\ &= \frac{\overbrace{\frac{1}{6} + \frac{1}{6}}^{Firstcolumn}}{\underbrace{\frac{1}{6} + \frac{1}{6}}_{FirstColumn} + \underbrace{\frac{1}{3}}_{Thirdcolumn}} \\ &= \frac{\frac{1}{3}}{\frac{2}{3}} \\ &= \frac{1}{2} \end{aligned}$$

(c) The guard returns to tell A that he passed the pardon message to prisoner B. What is the conditional probability of A being executed the next day if the guard says B is not executed the next day? (justify your answer mathematically!)

Solution:

$$
\begin{aligned}
P(Execute = A|GuardSays = B)) &= \frac{P(Execute = A \wedge GuardSays = B)}{P(GuardSays = B)} \\
&= \frac{\overbrace{\frac{1}{6}}^{First\ column, First\ row}}{\underbrace{\frac{1}{6} + \frac{1}{3}}_{FirstRow}} \\
&= \frac{\frac{1}{6}}{\frac{3}{6}} \\
&= \frac{1}{3}
\end{aligned}
$$

Exercise 11.3

From Nils J. Nilsson: Artificial Intelligence, 1998: Orville, the robot juggler, drops balls quite often when the battery is low. In previous tests, it has been determined that the probability that it will drop a ball when its battery is low is 0.9. Whereas when its battery is not low, the probability that it drops a ball is only 0.01. The battery was recharged not so long ago, and our best guess (before looking at Orville's latest juggling record) is that the odds that the battery is low are 10 to 1 against. A robot observer, with a somewhat unreliable vision system, reports that Orville dropped a ball. The reliability of the observer is given by the following probabilities:

$$p(\text{observer says that Orville drops}|\text{Orville does drop}) = 0.9$$

$$p(\text{observer says that Orville drops}|\text{Orville does not drop}) = 0.2$$

- Draw the Bayes network (and indicate the conditional probability tables)!

Solution:

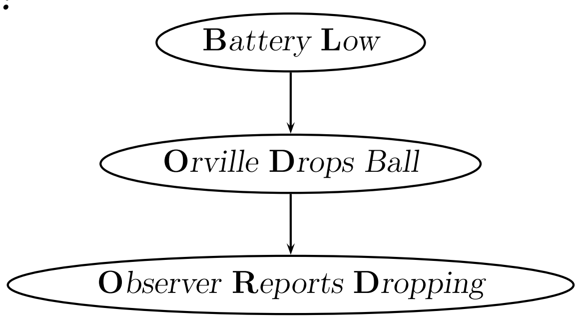

P(BL)
.0909

BL	**P(OD)**
T	.900
F	.001

OD	**P(ORD)**
T	.9
F	.2

- *Calculate that the battery is low given that the observer reports that Orville drops the ball.*

Solution:

$$
\begin{aligned}
P(BL \mid ORD) &= \frac{P(ORD \mid BL)P(BL)}{P(ORD)} \\
P(ORD \mid BL) &= P(ORD \mid OD)P(OD \mid BL) + P(ORD \mid \neg OD)P(\neg OD \mid BL) \\
P(ORD) &= P(ORD \mid OD)P(OD) + P(ORD \mid \neg OD)P(\neg OD) \\
P(OD) &= P(OD \mid BL)P(BL) + P(OD \mid \neg BL)P(\neg BL) \\
P(\neg OD) &= 1 - P(OD)
\end{aligned}
$$

Now substituting in the above from the values in the tables, we get:

$$
\begin{aligned}
P(OD) &= .9 \times .0909 + .001 \times .9091 \\
&= .0827191 \\
P(\neg OD) &= 1 - .0827191 \\
&= .9172809 \\
P(ORD) &= .9 \times .0827191 + .2 \times .9172809 \\
&= .25790337 \\
P(ORD \mid BL) &= .9 \times .9 + .2 \times .001 \\
&= .8102 \\
P(BL \mid ORD) &= \frac{.8102 \times .0909}{.25790337} \\
&= .28556
\end{aligned}
$$

- *Consider the following network (with probabilities omitted). Is it a correct Bayes network for the problem? Do you think that any other aspect of this network might be unsatisfactory?*

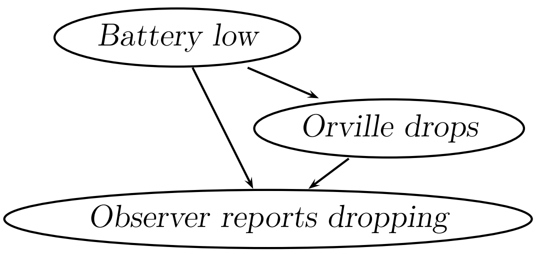

Solution:
The network is a correct one, but not as compact as it could be. Remember that Bayes networks are constructed such that $\mathbf{P}(X_i|X_{i-1}, ..., X_1) = \mathbf{P}(X_i|Parents(X_i))$*. In our example, the probability that the observer reports that Orville drops the ball depends on the event that Orville does or does not drop the ball. However, knowing in addition that the battery is low does not give any additional information. Therefore, the node* `Battery low` *should not necessarily be a parent for* `Observer reports dropping`*. The network that is displayed has the unpleasant proberty that the CPT (conditional probability*

table) for `Observer reports dropping` *needs to contain four entries if it depends both on* `Battery low` *and* `Orville drops`*, but only needs two if the link between* `Battery low` *and* `Observer reports dropping` *is removed, which makes the network smaller (and computation easier).*

Exercise 11.4

In this exercise we consider the example introduced in the lecture. There, the following Bayes net was presented:

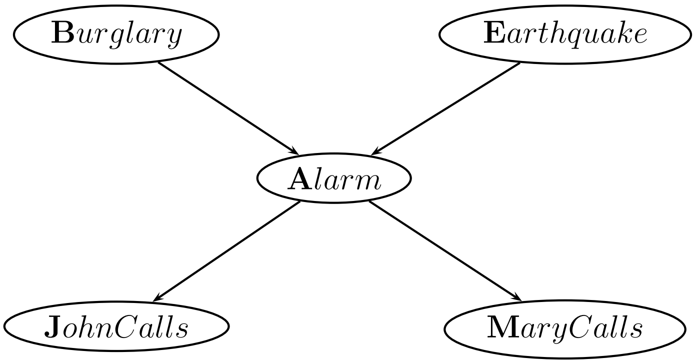

The *a priory* probabilities and conditional probabilites are:

P(B)
.001

P(E)
.002

B	**E**	**P(A)**
T	T	.95
T	F	.94
F	T	.29
F	F	.001

A	**P(J)**
T	.90
F	.05

A	**P(M)**
T	.70
F	.01

Compute the following conditional probabilities with the technique for exact inference presented in the lecture (also see Russell/Norvig, p. 504 ff.). Indicate all computations!

1. $\mathbf{P}(\mathbf{J} \mid \mathbf{B}) \ = ?$
2. $\mathbf{P}(\mathbf{B} \mid \mathbf{A}) \ = ?$
3. Explain the state of affairs that are considered in 1. and 2. in (ordinary) English or German sentences!

Solution:

a. **[Causal Inference]** $\mathbf{P}(\mathbf{J} \mid \mathbf{B}) \ = ?$

Solution:

$$
\begin{aligned}
P(J \mid B) &= P(J \mid A)P(A \mid B) + \\
&\quad P(J \mid \neg A)P(\neg A \mid B) \\
&= P(J \mid A)[P(A \mid B, E)P(E) + P(A \mid B, \neg E)P(\neg E)] + \\
&\quad P(J \mid \neg A)[1 - P(A \mid B)] \\
&= [P(J \mid A) - P(J \mid \neg A)][P(A \mid B, E)P(E) + P(A \mid B, \neg E)P(\neg E)] + P(J \mid \neg A) \\
&= [.90 - .05] \times [.95 \times .002 + .94 \times .998] + .05 \\
&= .85 \times [.0019 + .93812] + .05 \\
&= .85 \times .94002 + .05 \\
&= .849017
\end{aligned}
$$

Therefore, the probability that John calls if there is a burglary is 84.9017%.

b. **[Diagnostic Inference] P(B | A)** $= ?$

Solution:

$$
\begin{aligned}
P(B \mid A) &= \frac{P(A \mid B)P(B)}{P(A)} \\
&= \frac{[P(A \mid B, E)P(E) + P(A \mid B, \neg E)P(\neg E)]P(B)}{P(A)} \\
&= \frac{P(A \mid B, E)P(B)P(E) + P(A \mid B, \neg E)P(B)P(\neg E)}{P(A \wedge (B \vee \neg B) \wedge (E \vee \neg E))} \\
&= \frac{P(A \mid B, E)P(B)P(E) + P(A \mid B, \neg E)P(B)P(\neg E)}{P((A \wedge B \wedge E) \vee (A \wedge B \wedge \neg E) \vee (A \wedge \neg B \wedge E) \vee (A \wedge \neg B \wedge \neg E))} \\
&= \frac{P(A \mid B, E)P(B)P(E) + P(A \mid B, \neg E)P(B)P(\neg E)}{P(A \wedge B \wedge E) + P(A \wedge B \wedge \neg E) + P(A \wedge \neg B \wedge E) + P(A \wedge \neg B \wedge \neg E)} \\
&= \frac{P(A \mid B, E)P(B)P(E) + P(A \mid B, \neg E)P(B)P(\neg E)}{\begin{array}{ll} P(A \mid B \wedge E)P(B)P(E) & + \\ P(A \mid B \wedge \neg E)P(B)P(\neg E) & + \\ P(A \mid \neg B \wedge E)P(\neg B)P(E) & + \\ P(A \mid \neg B \wedge \neg E)P(\neg B)P(\neg E) & \end{array}} \\
&= \frac{.95 \times .001 \times .002 + .94 \times .001 \times .998}{\begin{array}{ll} P(A \mid B \wedge E)P(B)P(E) & + \\ P(A \mid B \wedge \neg E)P(B)P(\neg E) & + \\ P(A \mid \neg B \wedge E)P(\neg B)P(E) & + \\ P(A \mid \neg B \wedge \neg E)P(\neg B)P(\neg E) & \end{array}} \\
&= \frac{.00094002}{\begin{array}{ll} .00094002 & + \\ .29 \times .999 \times .002 & + \\ .001 \times .999 \times .998 & \end{array}} \\
&= \frac{.00094002}{.00094002 + .001576422} \\
&= \frac{.00094002}{.002516442} \\
&= .37355123
\end{aligned}
$$

I.e. when the alarm rings, the probability that a burglary has taken place is 37.35%.