Künstliche Intelligenz 2 – SS 2019

Michael Kohlhase
Informatik, FAU Erlangen-Nürnberg
For Course Purposes Only

June 17, 2019

Contents

1 Assignment 1 (Conditional Probabilities) – Given April 26., Due May 05.

Some Basic Notions Assume we have a set of elementary events Ω.

- The expression $P(A|B)$ means "The probability of A *given* B" (conditional probabilities), i.e. the probability of event A under the assumption that the event B has occured.

  We can define this as $P(A|B) = \frac{P(A \wedge B)}{P(B)}$.

- One of the central tools in probability theory is Bayes' theorem:

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

- Bayes' theorem is a formal theorem telling us how to *update our beliefs* in a *hypothesis* when confronted with new *evidence*. Let $\mathcal{H}$ the "space of all competing hypotheses", $H \in \mathcal{H}$ and e a new piece of evidence. Then Bayes theorem tells us, that

$$P(H|e) = \frac{P(e|H) \cdot P(H)}{P(e)}$$

  We call $P(H)$ the *prior* (or *a priori* probability); the probability you *thought* H was true before encountering e. We call $P(e|H)$ the *likelihood*; the probability with which your hypothesis H *predicted* the evidence e (before encountering it). We call $P(H|e)$ the *a posteriori* probability that H is true.

  Note, that $P(e)$ can be decomposed as $P(e) = P(e|H) \cdot P(H) + P(e|\neg H) \cdot P(\neg H) = \sum_{h \in \mathcal{H}} P(e|h) \cdot P(h)$. In other words, the probability of encountering a piece of evidence e *at all* is the sum of individual priors·likelihoods over all competing hypotheses.

- To update my belief in a hypothesis H given evidence e is to compute the a posteriori probability of H using Bayes' theorem, and to accept the result as my new prior going forward.

Axiom of Bayesian Rationality: Whenever you learn from evidence, the degree to which you learn *correctly* is precisely the degree to which your update of subjective probabilities corresponds to the result of applying Bayes' theorem.

The following exercise has no "correct" solution. Its intention is to make you think with and about conditional probabilities, and to model "real world" situations in terms of formal probabilities.

Problem 1.1 (Bayesian Epistemology)
Consider the following sayings. How can we express them in the form of conditional 50pt

probabilities? Give actual mathematical formulas and explain how they relate to the sayings.

Are they (as statements about probabilities) actually true or under which assumptions can they be?

- *Occam's Razor:* The simplest explanation is always the best.
- Extraordinary claims require extraordinary evidence.
- Absence of evidence is not evidence of absence.

Solution: Open question, but here are some ways to think about these things:

- Obviously, we have a problem to quantify what *simple* exactly means. So let's take one step back and think about what it means for one hypothesis A to be simpler than some hypothesis B.

  One way to answer that question would be to say: A is *simpler* than B if the set of propositions entailed by A is a proper subset of those entailed by B. In this case we can say $A \equiv P_1 \wedge \ldots \wedge P_n$ and $B \equiv P_1 \wedge \ldots \wedge P_n \wedge P_{n+1} \wedge \ldots \wedge P_m$.

  Naturally, we have $P(A) \geq P(A \wedge B)$ for any propositions A, B, hence under this interpretation the claim is true.

- Let's assume "extraordinary" means that the prior probability (in the absence of any evidence) is rather small, i.e. $P(A) \approx 0$[1]. If we want the claim A to be *likely*, we need to find some evidence e such that $P(A \mid e) \approx 1$. By Bayes' Theorem:

$$1 \approx P(A \mid e) = \frac{P(e \mid A) \cdot \overbrace{P(A)}^{\approx 0}}{P(e)}$$

  It is immediately obvious that for this to hold, $P(e)$ needs to be highly unlikely, i.e. extraordinary.

- Let's assume that e is evidence *for* A iff $P(A \mid e) > P(A)$, i.e. observing e actually makes the proposition A more likely. I claim: *If e is evidence for A, then $\neg e$ is evidence for $\neg A$*, making the claim false:

[1] I'll write $P(x) \approx 0$ resp. ≈ 1 simply for "is very unlikely" and "is very likely"

$$
\begin{aligned}
&P(A \mid e) = \frac{P(e \mid A)P(A)}{P(e)} > P(A) \\
\Rightarrow &P(e \mid A)P(A) > P(e)P(A) \\
\Rightarrow &(1 - P(\neg e \mid A))P(A) > (1 - P(\neg e))P(A) \\
\Rightarrow &\underbrace{P(\neg e \mid A)P(A)}_{=P(\neg e, A)} < P(\neg e)P(A) \\
\Rightarrow &\overbrace{P(A \mid \neg e){\color{red}P(\neg e)}} < {\color{red}P(\neg e)}P(A) \\
\Rightarrow &P(A \mid \neg e) < P(A) \\
\Rightarrow &1 - P(\neg A \mid \neg e) < 1 - P(\neg A) \\
\Rightarrow &P(\neg A \mid \neg e) > P(\neg A)
\end{aligned}
$$

□

---

2 Assignment 2 (Conditional Probabilities) – Given May 03., Due May 12.

Problem 2.1 (AFT Tests)

Trisomy 21 (*Down syndrome*) is a genetic anomaly that can be diagnosed during pregnancy using an amniotic fluid test. 40pt

The probability of a fetus having Down syndrome is strongly correllated with the age of the pregnant parent. For 25 year olds the probability is one in 1250, for 43 year old parents it increases to one in fifty (we only consider those two age groups).

However, diagnostic tests are never perfect. We distinguish two kinds of errors:

- **Type I Error (False Positive)**: The test result is positive even though the child is healthy.
- **Type II Error (False Negative)**: The test result is negative even though the child has trisomy 21.

The probabilities of Type I and Type II Errors are both merely 1% for amniotic fluid tests for Down syndrome.

10 pt

1. Express all of the above in the form of conditional probabilities. Use the random variable F with Domain $\{Age_{25}, Age_{43}\}$ for the age of the pregnant person and the boolean random variables Pos and $Down$ for the propositions "*The amniotic fluid test is positive*" and "*The child has Down syndrome*" respectively.

30 pt

2. Assume that we have a 25 year old pregnant person. Using Bayes' theorem, express and compute the probability that their child has Down syndrome, given that the amniotic fluid test is positive. What can we conclude from the result?

Solution:

1. $P(Down \mid F = Age_{25}) = 0.0008$, $P(Down \mid F = Age_{43}) = 0.02$, $P(Pos \mid \neg Down) = 0.01$, $P(\neg Pos \mid Down) = 0.01$.

2. We normalize to $F = Age_{25}$ and compute:

$$
\begin{aligned}
P(Down \mid Pos) &= \frac{P(Pos \mid Down) \cdot P(Down)}{P(Pos)} = \frac{P(Pos \mid Down) \cdot P(Down)}{P(Pos \wedge Down) + P(Pos \wedge \neg Down)} \\
&= \frac{P(Pos \mid Down) \cdot P(Down)}{P(Pos \mid Down) \cdot P(Down) + P(Pos \mid \neg Down \cdot P(\neg Down))} \\
&= \frac{(1 - P(\neg Pos \mid Down)) \cdot P(Down)}{(1 - P(\neg Pos \mid Down) \cdot P(Down)) + P(Pos \mid \neg Down) \cdot (1 - P(Down))} \\
&= \frac{0.99 \cdot 0.0008}{0.99 \cdot 0.0008 + 0.01\ 0.9992} \approx 0.07
\end{aligned}
$$

So, even with a positive test result, the probability of the child actually having Down syndrome is still only 7%, simply due to Down syndrome being relatively rare in young parents. Consequently, there is little point in applying this particular test without exceptional cause for concern.

Problem 2.2 (Disjunctive Random Variables)
We know that given boolean random variables A and B we have 20pt

$$P(A \vee B) = P(A) + P(B) - P(A \wedge B)$$

Extend this formula to the case of three random variables $P(A \vee B \vee C)$. Draw a Venn diagram to "prove" your formula.

Solution: $P(A \vee B \vee C) = P(A)+P(B)+P(C)-P(A \wedge B)-P(B \wedge C)-P(C \wedge A)+P(A \wedge B \wedge C)$

Problem 2.3 (Chained Production Elements)
An apparatus consists of six element A, B, C, D, E, F. The apparatus works if and only if at least A and B are operational, C and D are operational, or E and F are operational. 40pt

1. Assume the probabilities $P(X)$, that element X breaks down, are all stochastically independent, with $P(A) = 5\%$, $P(B) = 10\%$, $P(C) = 15\%$, etc. What is the probability the apparatus works? 20 pt
2. Unfortunately, the elements A and C, D and F and B and E are pairwise linked; such that if either of them breaks, then the linked element is not operational either. What is the probability that the apparatus works now? 20 pt

(**Note** that we deliberately differentiate between *not being operational* and *being broken*! If an element breaks, it is not operational; if an element is not operational, either it or the linked element broke)

Solution: Let W be the random variable stating that the apparatus works.

1. $$\begin{aligned}
P(W) &= P((A \wedge B) \vee (C \wedge D) \vee (E \wedge F)) \\
&= 1 - P(\underbrace{\neg(A \wedge B) \wedge \neg(C \wedge D) \wedge \neg(E \wedge F)}_{\text{all events are independent}}) \\
&= 1 - P(\neg(A \wedge B)) \cdot P(\neg(C \wedge D)) \cdot P(\neg(E \wedge F)) \\
&= 1 - P(\neg A \vee \neg B) \cdot P(\neg C \vee \neg D) \cdot P(\neg E \vee \neg F) \\
&= 1 - (P(\neg A) + P(\neg B) - P(\neg A) \cdot P(\neg B)) \cdot (P(\neg C) + P(\neg D) - P(\neg C) \cdot P(\neg D)) \\
&\cdot (P(\neg E) + P(\neg F) - P(\neg E) \cdot P(\neg F)) \\
&= 1 - (0.05 + 0.1 - (0.05 \cdot 0.1)) \cdot (0.15 + 0.2 - (0.15 \cdot 0.2)) \cdot (0.25 + 0.3 - (0.25 \cdot 0.3))
\end{aligned}$$

2. Using the exclusion principle:
$$\begin{aligned}
P(W) &= P((A \wedge B) \vee (C \wedge D) \vee (E \wedge F)) \\
&= P(A, B) + P(C, D) + P(E, F) - P(A, B, C, D) - P(A, B, E, F) - P(C, D, E, F) \\
&+ P(A, B, C, D, E, F)
\end{aligned}$$

Due to the links, this is equivalent to:

$$
\begin{aligned}
P(W) &= P(A,C,B,E) + P(A,C,D,F) + P(B,E,D,F) - P(A,B,C,D,E,F) - P(A,B,C,D,E,F) \\
&\quad - P(A,B,C,D,E,F) + P(A,B,C,D,E,F) \\
&= P(A,C,B,E) + P(A,C,D,F) + P(B,E,D,F) - 2P(A,B,C,D,E,F) \\
&= P(A)P(B)P(C)P(E) + P(A)P(C)P(D)P(F) + P(B)P(D)P(E)P(F) \\
&\quad - 2P(A)P(B)P(C)P(D)P(E)P(F) \\
&= 0.5450625 + 0.4522 + 0.378 - 2 \cdot 0.305235 \approx 76\%
\end{aligned}
$$

---

3 Assignment 3 (Bayesian Networks) – Given May 10., Due May 19.

Problem 3.1 (Stochastic Wumpus)
Robby lives in Wumpus world and wants to visit field F_1. He is pretty confident, that the Wumpus is not in field F_1; in fact, he is 90% sure. He thinks the Wumpus is probably in field F_2 with 60% confidence. Robby also thinks, that places without a Wumpus should rarely stink (in only 20% of cases), whereas every field with a Wumpus stinks. 70pt

Unfortunately, when Robbie approaches F_1, he notices a stench.

1. Give that F_1 stinks, how should Robbie update his belief that the Wumpus is not in F_1? How does the probability change, that it is in F_2? 25 pt
2. Just to be sure, Robbie takes a slight detour to F_2 and notices that it stinks there as well. Given this new piece of information, how should Robbie update his beliefs, that a) the Wumpus is in F_2 and b) he is not in F_1? 35 pt
3. Which random variables in this example are conditionally independent given which other random variable? 10 pt

Note that F_1 and F_2 are **not** the only fields that exist; so e.g. If the Wumpos isn't in F_1,that does not imply that it is in F_2.

Solution: Let S_i, W_i be the random variables expressing that it stinks / the Wumpus is in Field i respectively.

1.

$$
\begin{aligned}
P(S_1) &= P(S_1 \mid \neg W_1)P(\neg W_1) + P(S_1 \mid W_1) \cdot P(W_1) \\
&= 0.2 \cdot 0.9 + 1, 0.1 = 0.28 \\
P(\neg W_1 \mid S_1) &= \frac{P(S_1 \mid \neg W_1)P(\neg W_1)}{P(S_1)} = \frac{P(S_1 \mid \neg W_1)P(\neg W_1)}{0.28} \\
&= \frac{0.2 \cdot 0.9}{0.28} \approx 64\% \\
P(W_2 \mid S_1) &= \frac{P(S_1 \mid W_2) \cdot P(W_2)}{P(S_1)} = \frac{0.2\ 0.6}{0.28} \approx 43\%
\end{aligned}
$$

2. We normalize to S_1 (by using the probabilites from 1.) and compute:

$$\begin{aligned}
P(S_2) &= P(S_2 \mid W_1)\ P(W_1) + P(S_2 \mid W_2)\ P(W_2) + P(S_2 \mid \neg W_1 \wedge \neg W_2)\ P(\neg W_1 \wedge \neg W_2) \\
&= 0.2 \cdot (1 - 0.64) + 1 \cdot 0.43 + 0.2 \cdot (1 - ((1 - 0.64) + 0.43)) = 0.544 \\
P(W_2 \mid S_2) &= \frac{P(S_2 \mid W_2) \cdot P(W_2)}{P(S_2)} = \frac{P(S_2 \mid W_2) \cdot P(W_2)}{0.544} \\
&= \frac{1 \cdot 0.43}{0.544} \approx 79\% \\
P(\neg W_1 \mid S_2) &= 1 - P(W_1 \mid S_2) = 1 - \frac{P(S_2 \mid W_1) \cdot P(W_1)}{P(S_2)} \\
&= 1 - \frac{0.2 \cdot (1 - 0.64)}{0.544} \approx 87\%
\end{aligned}$$

3. S_1 and S_2 are conditionally independent given W_1 (or W_2)

Problem 3.2 (Nuclear Test)

Assume it is your responsibility to monitor the Nuclear Test Ban treaty. You receive data from two different stations (seismometers), S_1 and S_2. Each S_i is modeled as a Boolean variable where "true" stands for "I detected a Nuclear test" and "false" stands for "I did not detect a Nuclear test". The seismometers are not fully reliable, however; they may not detect a Nuclear test even though there was one, and they may mistake an earthquake for a Nuclear test. We model this situation with two additional Boolean variables: N for Nuclear test, and E for Earthquake. 30pt

Use the algorithm from the lecture (Slide 648) to construct a Bayesian network for these 4 variables. Do so for the following two variable orders:

(a) $X_1 = N$, $X_2 = E$, $X_3 = S_1$, $X_4 = S_2$.

(b) $X_1 = S_1$, $X_2 = S_2$, $X_3 = E$, $X_4 = N$.

For each of these orders, draw the resulting Bayesian network. Justify your design, i.e., for each variable X_i added to the network explain why the set of parents you give X_i are needed, and why they are sufficient.

Solution:

(a) With this variable order, we get the following network:

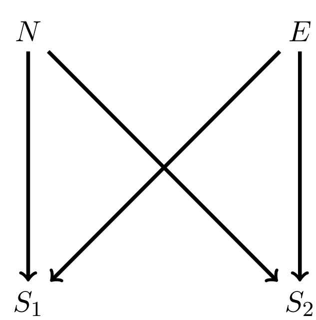

$X_2 = E$ does not need $X_1 = N$ as a parent because Earthquakes are independent from Nuclear tests. $x_3 = S_1$ needs both $X_1 = N$ and $X_2 = E$ as parents because each of these may influence the measurement; same for $X_4 = S_2$, i.e., here we also need the parents $X_1 = N$ and $X_2 = E$. However, given the values of N and E, the measurements of $X_3 = S_1$ and $x_4 = S_2$ are independent. So $X_4 = S_2$ does not require the parent $X_3 = S_1$.

(b) With this variable order, we get the following network:

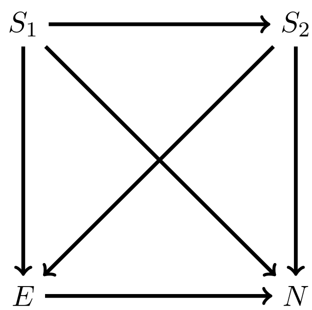

$X_2 = S_2$ needs $X_1 = S_1$ as a parent because, if S_1 detects a seismic phenomenon, then chances are higher S_2 will detect one as well. $X_3 = E$ needs each of $X_1 = S_1$ and $X_2 = S_2$ as parents because, if a station detects a seismic phenomenon, then chances are higher there was an earthquake; same for $X_4 = N$, i.e., here we also need the parents $X_1 = S_1$ and $X_2 = S_2$ because measurements indicate Nuclear tests as well. Finally, say we already know that S_1 and S_2 are true; then the value of E still has an influence on the value of N: If there was an earthquake, then there is a chance that the seismic measurements were caused by the earthquake rather than a Nuclear test. Thus N is *not* conditionally independent of E given S_1 and S_2, and we need $X_3 = E$ as a parent of $X_4 = N$ as well.

---

4 Assignment 4 (Inference in Bayesian Networks) – Given May 17., Due May 26.

Problem 4.1 (Medical Bayesian Network)

Dyspnea (shortness of breath) can be caused by several medical conditions; among them lung cancer, tuberculosis and bronchitis. Tuberculosis and cancer lead to abnormal x-ray results. Lung cancer and bronchitis can be caused by smoking, tuberculosis occurs more often in asia. We use the following random variables for some given patient: 50pt

- $Asia$: The patient recently visited asia.
- $Smoke$: The patient is a smoker.
- TBC: The patient has tuberculosis.
- LC: The patient has lung cancer.
- $Bron$: The patient has bronchitis.
- $Xray$: The patient's X-ray result is abnormal.
- $Dysp$: The patient is short of breath.

Model the dependencies stated above as a bayesian network, choosing a suitable(!) ordering of the variables. **Justify** your choices.

Solution:

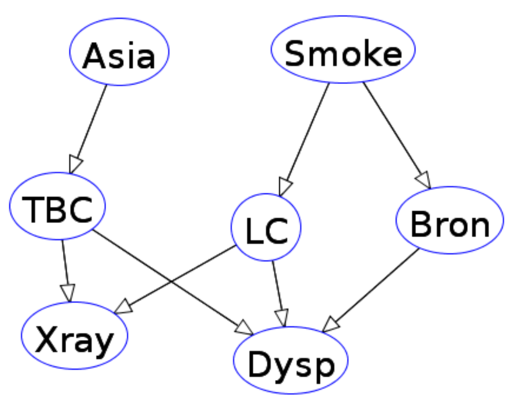

Problem 4.2 (Medical Bayesian Network 2)

Both Malaria and Meningitis can cause a measured high body temperature; the latter is an indicator of fever. We consider the following random variables for a given patient: 50pt

- Mal: The patient has malaria (probability 5%).
- Men: The patient has meningitis (probability 20%).

- HBT: The patient has a high body temperature:

Mal	Men	HBT
$\top$	$\top$	95%
$\top$	$\bot$	90%
$\bot$	$\top$	90%
$\bot$	$\bot$	0.5%

- F: The patient has a fever:

HBT	F
$\top$	95%
$\bot$	0.2%

1. Draw the corresponding Bayesian network for the above data for a suitable ordering of the variables.
2. What is the probability some patient has malaria and not meningitis, given that he has a fever? Indicate which variables are hidden variables, query variables and observed events, use the inference algorithm for bayesian networks from the lecture and compute the result by hand.

Solution:

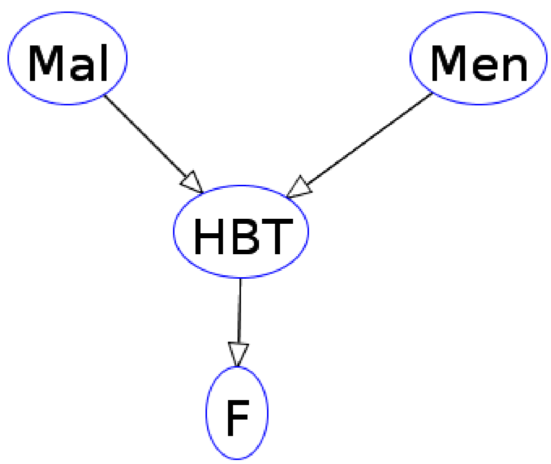

We need to compute all four cases $\{\pm Mal, \pm Men\}$:

$$\begin{aligned}
\langle P(v_{Mal}, v_{Men}|F)\rangle &= \alpha \left\langle \sum_{v_{HBT}} P(v_{HBT}, F, v_{Men}, v_{Mal}) \right\rangle \\
&= \alpha \left\langle \sum_{v_{HBT}} (P(F|v_{HBT}) \cdot P(v_{HBT}|v_{Men}, v_{Mal}) \cdot P(v_{Men}) \cdot P(v_{Mal})) \right\rangle \\
&= \alpha \left\langle P(v_{Men}) \cdot P(v_{Mal}) \cdot \sum_{v_{HBT}} \underbrace{(P(F|v_{HBT}) \cdot P(v_{HBT}|v_{Men}, v_{Mal}))}_{:=\sigma_{v_{HBT}, v_{Men}, v_{Mal}}} \right\rangle
\end{aligned}$$

, hence:

$$\begin{aligned}
P(Mal, Men|F) &= \alpha \cdot P(Men) \cdot P(Mal) \cdot (\sigma_{HBT,Men,Mal} + \sigma_{\neg HBT,Men,Mal}) \\
&= \alpha \cdot 0.2 \cdot 0.05 \cdot (0.95 \cdot 0.95 + 0.002 \cdot 0.05) = \alpha \cdot 0.009026 \\
P(Mal, \neg Men|F) &= \alpha \cdot P(\neg Men) \cdot P(Mal) \cdot (\sigma_{HBT,\neg Men,Mal} + \sigma_{\neg HBT,\neg Men,Mal}) \\
&= \alpha \cdot 0.8 \cdot 0.05 \cdot (0.95 \cdot 0.9 + 0.002 \cdot 0.1) = \alpha \cdot 0.034208 \\
P(\neg Mal, Men|F) &= \alpha \cdot P(Men) \cdot P(\neg Mal) \cdot (\sigma_{HBT,Men,\neg Mal} + \sigma_{\neg HBT,Men,\neg Mal}) \\
&= \alpha \cdot 0.2 \cdot 0.95 \cdot (0.95 \cdot 0.9 + 0.002 \cdot 0.1) = \alpha \cdot 0.162488 \\
P(\neg Mal, \neg Men|F) &= \alpha \cdot P(\neg Men) \cdot P(\neg Mal) \cdot (\sigma_{HBT,\neg Men,\neg Mal} + \sigma_{\neg HBT,\neg Men,\neg Mal}) \\
&= \alpha \cdot 0.8 \cdot 0.95 \cdot (0.95 \cdot 0.005 + 0.002 \cdot 0.995) = \alpha \cdot 0.0051224
\end{aligned}$$

Therefore:

$$\alpha\langle 0.009026, 0.034208, 0.162488, 0.0051224\rangle \approx \langle 4.3\%, 16.2\%, 77\%, 2.4\%\rangle$$

Hence, $P(Mal, \neg Men|F) \approx 16.2\%$

5 Assignment 5 (Bayesian Networks in Java) – Given May 24., Due June 02.

Problem 5.1 (Bayesian Networks in Java)

100pt

The goal of this exercise is to implement inference by enumeration (or aonther inference algorithm of your choice) in bayesian networks in java. Use the java classes in `bayes.zip` (linked on studon). Those are:

- `Node`: This class has two fields: A string `id` for the name, and a field of type `Node[]` called `parents` (for the parents of the node). To initialize the probabilities of a node, you can use the method `void setProb(Boolean[] given, double value)`. To access the probability of a node (given values for the parents), use `public double getProb(Boolean[] given)`.
- Nodes are collected in an object of the class `Network`, which has the two methods `public ArrayList<Node> getNodes()` and `public void addNode(Node n)`.

The final class is `Query`. Your taks is to implement an extension of this class, which requires you to implement its method

```
public double query(Network network, Node node, ArrayList<Pair<Node, Boolean>> evidence)
```

,which is supposed to compute the probability $P(\texttt{node}|\texttt{evidence})$ (where evidence is a list of pairs of the form $\texttt{node} = \top/\bot$).

You can test your network with the Test class, which contains an implementation of the Network from Exercise 4.1, which you are not required to solve.

If you prefer, you can use Scala as a programming language instead; for that case the file `ScalaQuery.scala` contains an abstract class `ScalaQuery` extending `Query` with convenience methods that take care of converting between java and Scala datatypes for you. Here you need to implement the method

```
queryScala(network:Network,node:Node,evidence:List[(Node,Boolean)]) : Double
```

instead.

Hand in your solution as a `.java` or `.scala` file containing a (uniquely named!) class extending `Query` in the namespace `info.kwarc.teaching.bayes.SS19`. Make sure your file contains your full name as a comment!

Solution:

6 Assignment 6 (Decisions and Utilities) – Given May 31., Due June 09.

Problem 6.1 (Decision Theory)

You are offered the following game: You pay x dollars to play. A fair coin is then tossed repeatedly until it comes up heads for the first time. Your payout is 2^n, where n is the number of tosses that occurred. 50pt

- Assume your utility function is exactly the monetary value. How much should you, as a rational agent, be willing to pay to play? Use the formal definition of "expected utility" from the lecture.
- Assume now, that your utility function for having k dollars is $U(k) = m \log_n k$ for some $m, n \in \mathbb{N}^+$. How does this change the result?
- What is "wrong" with the result from the first exercise? Which "implicit assumption" leads to the apparently nonsensical result? Can you think of a way to "repair" our utility function in a more "realistic" way than taking logarithms?

Hint: The series $\sum_{k=1}^{\infty} \frac{k}{2^k}$ is convergent with limit 2.

Solution:

- We have

$$EU(\mathsf{Play}) = \sum_{s'} (P(\mathsf{Payout} = s' | \mathsf{Play}) \cdot U(s') = \sum_{k \in \mathbb{N}^+} P(\mathsf{Payout} = k) \cdot 2^k$$
$$= \sum_{k \in \mathbb{N}^+} \frac{1}{2^k} \cdot 2^k = \sum_{k \in \mathbb{N}^+} 1 \to \infty$$

  We should be willing to pay any amount for a chance to play.

- $$EU(\mathsf{Play}) = \sum_{k \in \mathbb{N}^+} \frac{1}{2^k} \cdot m \log_n(2^k) = m \log_n(2) \sum_{k \in \mathbb{N}^+} \frac{k}{2^k} = 2m \log_n(2)$$

  Of course, this is wrong insofar as the utility, being logarithmic, is in particular not linear, i.e. the actual utility depends on our original capital K, but this just makes everything more complicated. The point is that using a logarithmic utility for money yields a finite result.

- This model assumes that the payout is potentially infinite (which is unrealistic), as well as that we have an unlimited amount of money at our disposal. If we add the (negative!) utility of the cost, we can set an "upper limit" of how much we can afford to pay in the first place; e.g. by setting $U(k) = -\infty$ for k being the amount in our bank account.

Problem 6.2 (Medical Decisions)

A patient shows symptoms of lung cancer (LC). The obvious way to check whether they do is to do a biopsy (B), however the symptoms could also be explained by the rare stab-me-and-I'll-burst lung disease (SMAIB). If the patient does have SMAIB instead, a biopsy could be highly fatal (for the patient, not you). Also, Biopsies are annoying to execute, and you've been in the medical profession long anough to have become jaded, so you don't really care about your patients anymore. 50pt

If lung cancer is confirmed (C) by a biopsy, there's a pretty good chance the patient can make a recovery (Liv) if you start treatment now. If you don't do a biopsy, that chance is significantly lower.

- $SMAIB = \neg LC$
- $P(LC) = 0.9$
- $P(Liv|C, LC) = P(Liv|C, B, LC) = P(Liv|C, \neg B, LC) = 0.95$
- $P(Liv|\neg C, LC) = P(Liv|\neg B, LC) = P(Liv|\neg C, B, LC) = P(Liv|\neg C, \neg B, LC) = P(Liv|\neg C, \neg B, \neg LC) = 0.3$
- $P(Liv|C, \neg LC) = P(Liv|B, \neg LC) = P(Liv|C, B, \neg LC) = P(Liv, |\neg C, B, \neg LC) = 0.05$
- $P(C|\neg B) = 0$
- $P(C|B, LC) = 0.8$
- $P(C|B, \neg LC) = 0.1$

Your utility function is additive. The plain utility of the patient surviving is 100, that of the patient dying -10, the plain utility of doing a biopsy is -50.

1. Draw the decision network for this problem. 10 pt
2. Compute whether it is "rational" to do a biopsy, given the above utility function. 40 pt

---

Solution:

$$\begin{aligned} EU(B) &= P(Liv|B)U(Liv) + P(\neg Liv|B)U(\neg Liv) - 50 \\ &= 100P(Liv|B) - 10P(\neg Liv|B) - 50 \\ EU(\neg B) &= P(Liv|\neg B)U(Liv) + P(\neg Liv|\neg B)U(\neg Liv) \\ &= 100P(Liv|\neg B) - 10P(\neg Liv|\neg B) \end{aligned}$$

Now we compute the conditional probabilities using inference by enumeration:

$$
\begin{aligned}
P(v_{Liv}|B) &= \left\langle \sum_{v_{LC}} \sum_{v_C} P(v_{Liv}, v_C, v_{LC}, B) \right\rangle \\
&= \left\langle \sum_{v_{LC}} \sum_{v_C} P(v_{Liv}|v_C, v_{LC}, B) \cdot P(v_C|v_{LC}, B) \cdot P(v_{LC}) \right\rangle \\
&= \left\langle \sum_{v_{LC}} P(v_{LC}) \sum_{v_C} P(v_{Liv}|v_C, v_{LC}, B) \cdot P(v_C|v_{LC}, B) \right\rangle \\
&= \langle 0.9(0.95 \cdot 0.8 + 0.3 \cdot 0.2) + 0.1(0.05 \cdot 0.1 + 0.05 \cdot 0.9)\ ,\ 0.05 \cdot 0.8 + 0.7 \cdot 0.2 + 0.95 \cdot 0.1 + 0.95 \cdot 0.9 \rangle \\
&= \langle 0.743, 0.257 \rangle = \langle 74.3\%, 25.7\% \rangle
\end{aligned}
$$

Hence $EU(B) = 74.3 - 2.57 - 50 = 21.37$.

Repeat for $\neg B$ and compare.

---

7 Assignment 7 (Markov Models) – Given June 07., Due June 16.

Problem 7.1 (Markov Mood Detection)

On any given day d, your flatmate Moody is in one of two states – either he is happy (H_d) or he is not. Usually when he's in a bad mood, it's because he had a fight with his spouse and those tend to go on for a couple of days, so $P(\neg H_{d+1}|\neg H_d) = 0.7$, but aside from that he's a cheery guy, so $P(H_{d+1}|H_d) = 0.85$. 100pt

Notably, you can hear his music blasting all day which tends to shift depending on his mood. On a good day he often listens to Jazz (i.e. $P(J_d|H_d) = 0.6$), on a bad day he usually blasts Death Metal at full volume ($P(\neg J_d|\neg H_d) = 0.85$). He has a limited taste in music, so it's always one of the two.

It's Friday now and and you can hear Jazz playing from his room, but Tuesday through Thursday you were constantly bombarded by Death Metal through your paper thin walls. You need to talk to him about his disregard for kitchen hygiene, but you would prefer to only do so if he is already in a good mood.

1. Assume you happen to know he was in a good mood on Monday. Which algorithm is adequate for computing the probability that he's in a good mood today? What is that probability? 40 pt
2. On second thought, maybe you'd rather wait until Sunday. Which algorithm is adequate for computing the probability that he'll be in a good mood then? What is that probability? 30 pt
3. Maybe it would have been better to talk to him on Wednesday? Which algorithm is best suited and what is that probability? 30 pt

---

Solution:

1. **Filtering:** We assume $P(H_0) = 1$. Let h_i our (hidden) variables and $e_1 = \neg J_1$, $e_2 =$

$\neg J_2 \wedge \neg J_1$, etc. Then:

$$
\begin{aligned}
\langle P(h_1|e_1)\rangle &= \alpha \langle P(\neg J_1|h_1)P(h_1|H_0)\rangle \\
&= \alpha \left\langle \underbrace{P(\neg J_1|H_1)}_{0.4}\underbrace{P(H_1|H_0)}_{0.85}, \underbrace{P(\neg J_1|\neg H_1)}_{0.85}\underbrace{P(\neg H_1|H_0)}_{0.15} \right\rangle \\
&= \alpha \langle 0.34, 0.1275\rangle \approx \langle 0.73, 0.27\rangle \\
\langle P(h_2|e_2)\rangle &= \alpha \left\langle P(\neg J_2|h_2)\sum_{h_1} P(h_2|h_1)P(h_1|e_1) \right\rangle \\
&= \alpha \left\langle \underbrace{P(\neg J_2|H_2)}_{0.4}\underbrace{\sum_{h_1} P(H_2|h_1)P(h_1|e_1)}_{=0.85\cdot 0.73+0.3\cdot 0.27=0.7015}, \underbrace{P(\neg J_2|\neg H_2)}_{0.85}\underbrace{\sum_{h_1} P(\neg H_2|h_1)P(h_1|e_1)}_{=0.15\cdot 0.73+0.7\cdot 0.27=0.2985} \right\rangle \\
&= \alpha \langle 0.2806, 0.253725\rangle \approx \langle 0.53, 0.47\rangle \\
\langle P(h_3|e_3)\rangle &= \alpha \left\langle P(\neg J_3|h_3)\sum_{h_2} P(h_3|h_2)P(h_2|e_2) \right\rangle \\
&= \alpha \left\langle \underbrace{P(\neg J_3|H_3)}_{0.4}\underbrace{\sum_{h_2} P(H_3|h_2)P(h_2|e_2)}_{=0.85\cdot 0.53+0.3\cdot 0.47=0.5915}, \underbrace{P(\neg J_3|\neg H_3)}_{0.85}\underbrace{\sum_{h_2} P(\neg H_3|h_2)P(h_2|e_2)}_{=0.15\cdot 0.53+0.7\cdot 0.47=0.4085} \right\rangle \\
&= \alpha \langle 0.2366, 0.347225\rangle \approx \langle 0.41, 0.59\rangle \\
\langle P(h_4|e_4)\rangle &= \alpha \left\langle P(J_4|h_4)\sum_{h_3} P(h_4|h_3)P(h_3|e_3) \right\rangle \\
&= \alpha \left\langle \underbrace{P(J_4|H_4)}_{0.6}\underbrace{\sum_{h_3} P(H_4|h_3)P(h_3|e_3)}_{=0.85\cdot 0.41+0.3\cdot 0.59=0.5255}, \underbrace{P(J_4|\neg H_4)}_{0.15}\underbrace{\sum_{h_3} P(\neg H_4|h_3)P(h_3|e_3)}_{=0.15\cdot 0.41+0.7\cdot 0.59=0.4745} \right\rangle \\
&= \alpha \langle 0.3153, 0.071175\rangle \approx \langle 0.82, 0.18\rangle
\end{aligned}
$$

Hence the probability he's happy today is 82%.

2. **Prediction:** Using the numbers from the previous exercise:

$$\begin{aligned}
\langle P(h_5|e_4)\rangle &= \left\langle \sum_{h_4} P(h_5|h_4)P(h_4|e_4) \right\rangle \\
&= \left\langle \underbrace{\sum_{h_4} P(H_5|h_4)P(h_4|e_4)}_{=0.85\cdot 0.82+0.3\cdot 0.18=0.751}, 1 - P(H_5|e_4) \right\rangle \\
&\approx \langle 0.75, 0.25\rangle \\
\langle P(h_6|e_4)\rangle &= \left\langle \sum_{h_5} P(h_6|h_5)P(h_5|e_4) \right\rangle \\
&= \left\langle \underbrace{\sum_{h_5} P(H_6|h_5)P(h_5|e_4)}_{=0.85\cdot 0.75+0.3\cdot 0.25=0.7125}, 1 - P(H_6|e_4) \right\rangle \\
&\approx \langle 0.71, 0.29\rangle
\end{aligned}$$

3. **Smoothing:** Using the numbers from Exercise 1:

$$
\begin{aligned}
\langle P(J_4|h_3)\rangle &= \left\langle \sum_{h_4} P(J_4|h_4)P(h_4|h_3) \right\rangle \\
&= \left\langle \underbrace{\sum_{h_4} P(J_4|h_4)P(h_4|H_3)}_{=0.6\cdot 0.85+0.15\cdot 0.15=0.5325}, \underbrace{\sum_{h_4} P(J_4|h_4)P(h_4|\neg H_3)}_{=0.6\cdot 0.3+0.15\cdot 0.7=0.285} \right\rangle \\
&\approx \langle 0.53, 0.29\rangle \\
\langle P(h_3|e_4)\rangle &= \alpha \left\langle P(h_3|e_3)\cdot P(J_4|h_3)\right\rangle \\
&= \alpha \left\langle \underbrace{P(H_3|e_3)}_{0.41}\cdot \underbrace{P(J_4|H_3)}_{0.53}, \underbrace{P(\neg H_3|e_3)}_{0.59}\cdot \underbrace{P(J_4|\neg H_3)}_{0.29} \right\rangle \\
&= \alpha\langle 0.2173, 0.1711\rangle \approx \langle 0.56, 0.44\rangle \\
\langle P(\neg J_3, J_4|h_2)\rangle &= \left\langle \sum_{h_3} P(\neg J_3|h_3)P(J_4|h_3)P(h_3|h_2) \right\rangle \\
&= \left\langle \underbrace{\sum_{h_3} P(\neg J_3|h_3)P(J_4|h_3)P(h_3|H_2)}_{=0.4\cdot 0.53\cdot 0.85+0.85\cdot 0.29\cdot 0.15=0.217175}, \underbrace{\sum_{h_3} P(\neg J_3|h_3)P(J_4|h_3)P(h_3|\neg H_2)}_{=0.4\cdot 0.53\cdot 0.3+0.85\cdot 0.29\cdot 0.7=0.23615} \right\rangle \\
&\approx \langle 0.22, 0.24\rangle \\
\langle P(h_2|e_4)\rangle &= \alpha \left\langle P(h_2|e_2)\cdot P(\neg J_3, J_4|h_2)\right\rangle \\
&= \alpha \left\langle \underbrace{P(H_2|e_2)}_{0.53}\underbrace{P(\neg J_3, J_4|H_2)}_{0.22}, \underbrace{P(\neg H_2|e_2)}_{0.47}\underbrace{P(\neg J_3, J_4|\neg H_2)}_{0.24} \right\rangle \\
&= \alpha\langle 0.1166, 0.1128\rangle \approx \langle 0.51, 0.49\rangle
\end{aligned}
$$

---

8 Assignment 8 (Markov Decision Procedures) – Given June 14., Due June 23.

Problem 8.1 (Markov Games)

We want to apply Markov Decision Procedures to two-player games as considered in KI1. 100pt
We assume two players A, B and model everything from A's point of view - in particular, we have a reward function $R(s)$ for our states and our *states* are *only those game states, where it's A's move*. Hence, the transition function $T(s, a, s')$ models the probability that, if player A does action a in state s, player B will play such that we end up in state s'.

Let $N(s, a)$ be the set of possible successor states of s after picking move a.

10 pt

1. Let $U(s)$ be the utility of state s for player A. Write down the Bellman equation defining U, assuming player B always moves randomly (and uniformly distributed) and such that the equation does *not* contain the transition function T anymore (i.e. use $N(s, a)$ instead).

20 pt

2. Obviously B doesn't move randomly, though. In fact, the higher my utility $U(s')$, the smaller the probability that he's going to pick a move resulting in s'.

   Let's model the easiest case for that: In state s, if A picks given action a, then the probability $P(s') = \frac{c}{U(s')}$ for some constant c.

   Compute c and write down the corresponding Bellman equation. Again, it should not contain T or c anymore.

20 pt

3. Give the Bellman equation with a transition function, reward function and discount factor such that the resulting policy is equivalent to the one resulting from the Minimax algorithm from last semester.

10 pt

4. Consider a game with four fields F_1, F_2, F_3, F_4. Player A starts in F_1, player B in F_4. Each turn, a player has to move to an adjacent field (i.e. if the player is in F_i, he can move to F_{i+1} or F_{i-1}). If an adjacent field contains the other player, one may jump over him. E.g. assume player A is in F_2 and player B in F_3, then player A may jump over B to end up in F_4.

   The game ends when one player reaches the opposite end of the board, i.e. A wins if he ends up in F_4, B wins if he ends up in F_1. player A starts. The reward for player A winning is 100, the reward for losing is 1 (as to avoid utility 0 anywhere).

   Draw the state space, showing the moves by A and B as arrows in different colors/styles.

   You will find it helpful to arrange the states (s_A, s_B) on a two-dimensional grid, using s_A and s_B as coordinates.

40 pt

5. Apply value iteration (starting with utility 50 everywhere) assuming an unkown discount factor $0 < \gamma < 1$ to derive a strategy for player A, using the Bellman equation

from the second exercise (you can disregard impossible states, of course). If you can not solve the second exercise, use the first Bellman equation instead.

Give the resulting strategy (as a sequence of actions) and how this derives from the computed utility function.

Note that in the lecture we alway assume the transition function $T(s, a, s')$ to be fixed, whereas for this exercise, the transition function itself depends on the utilities in the Bellman equations for Exercises 2 and 3, and hence will change during value iteration as well!