

Name:

Birth Date:

Matriculation Number:

## Exam Artificial Intelligence 2

Feb. 14., 2019

	To be used for grading, do not write here												
prob.	1.1	1.2	1.3	1.4	2.1	2.2	2.3	3.1	3.2	4.1	4.2	Sum	grade
total	4	4	7	7	4	10	12	3	10	10	4	75	
reached													

Exam Grade:

Bonus Points:

Final Grade:

The “solutions” to the exam/assignment problems in this document are supplied to give students a starting point for answering questions. While we are striving for helpful “solutions”, they can be incomplete and can even contain errors.

If you find “solutions” you do not understand or you find incorrect, discuss this on the course forum and/or with your TA and/notify the instructors.

In any case, grading student’s answers is not a process of simply “comparing with the reference solution”.

In the course Artificial Intelligence I/II we award 5 bonus points for the first student who reports a factual error (please report spelling/formatting errors as well) in an assignment or old exam and 10 bonus points for an alternative solution (formatted in L<sup>A</sup>T<sub>E</sub>X) that is usefully different from the existing ones.

# 1 Bayesian Reasoning

## Problem 1.1 (Bayesian Rules)

Name four of the basic rules in Bayesian inference and explain each with a short sentence and formula. \_\_\_\_\_ 4 pt

**Solution:** \_\_\_\_\_ 4 min

1. Bayes rule (compute  $P(A|B)$  from  $P(B|A)$ ),
  2. Normalization (Fixing evidence  $e$ , updating the probabilities of all other events using a normalization constant  $\alpha$ ),
  3. Marginalization ( $P(A) = \sum_y P(A, y)$ ),
  4. Chain rule ( $P(A_1, \dots, A_n) = P(A_n|A_{n-1}, \dots, A_1) \cdot P(A_{n-1}|A_{n-2}, \dots, A_1) \cdot \dots$ )
  5. Product rule ( $P(A, B) = P(A|B)P(B)$ )
  6. Conditional Independence (Not really bayesian inference, but rather bayesian networks, but we'll be lenient)
-

**Problem 1.2 (Causal and Diagnostic)**

State the difference between *causal* and *diagnostic* edges in (e.g.) a Bayesian network. Use no more than four sentences. 4 pt  
4 min

**Problem 1.3 (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: 7 pt  
7 min

- *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 1.4 (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  $F_1$  (call that event  $\neg W_1$ ); in fact, he is 90% sure. He thinks the Wumpus is probably in field  $F_2$  (call that event  $W_2$ ) 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.

7 pt  
7 min

Unfortunately, when Robbie approaches  $F_1$ , he notices a stench (call that event  $S_1$ ).

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  $F_2$ ? *Do not compute actual values* - a formula how to compute them is sufficient!
2. Assume the updated values from the previous subexercise given. Just to be sure, Robbie takes a slight detour to  $F_2$  and notices that it stinks there as well (call that event  $S_2$ ). 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  $F_1$ ? (Again, a formula is sufficient!)
3. Which random variables in this example are conditionally independent given which other random variable?

---

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

1.

$$\begin{aligned}P(S_1) &= P(S_1 | \neg W_1, P(\neg W_1)) + P(S_1 | W_1) \cdot P(W_1) \\ &= 0.2 \cdot 0.9 + 1 \cdot 0.1 = 0.28 \\ P(\neg W_1 | S_1) &= \frac{P(S_1 | \neg W_1, P(\neg W_1))}{P(S_1)} = \frac{P(S_1 | \neg W_1, P(\neg W_1))}{0.28} \\ &= \frac{0.2 \cdot 0.9}{0.28} \approx 64\% \\ P(W_2 | S_1) &= \frac{P(S_1 | W_2) \cdot P(W_2)}{P(S_1)} = \frac{0.20 \cdot 0.6}{0.28} \approx 43\%\end{aligned}$$

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

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

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

---

## 2 Decision Theory

### Problem 2.1 (Expected Utility)

What is the formal(!) definition of *expected utility*? Explain every variable in the defining equation. 4 pt

---

**Solution:** The expected utility  $EU$  is defined as 4 min

$$EU(a|e) = \sum_{s'} P(R(a) = s'|a, e) \cdot U(s')$$

where

1.  $a$  is the action for which we want to find out the expected utility, given the evidence  $e$ .
  2.  $U(s')$  is the utility of a state  $s'$ .
  3.  $R(a)$  is the result of the action  $a$ .
-



**Problem 2.2 (Textbook Decisions)**

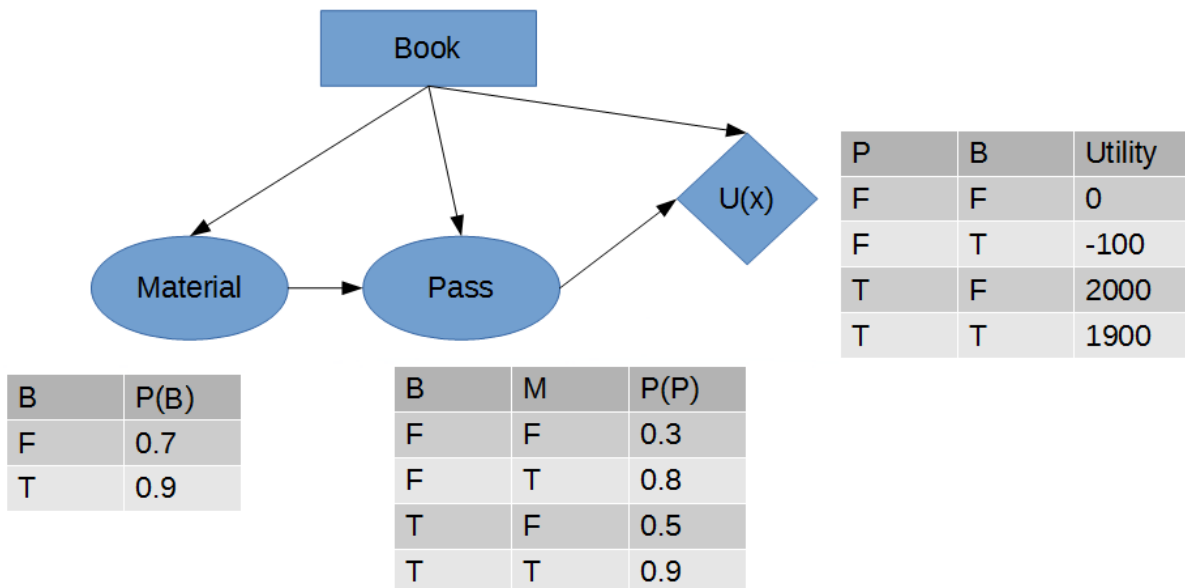
Abby has to decide whether to buy Russell&Norveig for 100\$. There are three boolean variables involved in this decision:  $B$  indicating whether Abby buys the book,  $M$  indicating whether Abby knows the material in the book perfectly anyway and  $P$  indicating that Abby passes the course. Additionally, we use a utility node  $U$ . 10 pt  
10 min

Abby's utility function is additive, so  $U(B) = -100$ . Furthermore, she evaluates passing the course with a utility of  $U(P) = 2000$ . The course has an open book final exam, so  $B$  and  $P$  are not independent given  $M$ . Assume the following conditional probabilities as given:

- $P(P|B, M)$
- $P(P|B, \neg M)$
- $P(P|\neg B, M)$
- $P(P|\neg B, \neg M)$
- $P(M|B)$
- $P(M|\neg B)$

1. Draw the decision network for this problem.
2. Explain formally how to compute the expected utility of buying the book and of not buying it.

**Solution:**(Note that the numbers are wrong/outdated)



Assuming  $B$ :

$$\begin{aligned}P(B,P,M) &= P(P \mid M,B) \cdot P(B) \cdot P(M \mid B) = 0.95 \cdot 1 \cdot 0.85 = 0.8075 \\P(B,P,\neg M) &= 0.45 \cdot 1 \cdot 0.15 = 0.0675 \\P(B,\neg P,M) &= 0.05 \cdot 1 \cdot 0.85 = 0.0425 \\P(B,\neg P,\neg M) &= 0.55 \cdot 1 \cdot 0.15 = 0.0825\end{aligned}$$

$$\begin{aligned}P(B,P) &= 0.8075 + 0.0675 = 0.875 \\P(B,\neg P) &= 0.0425 + 0.0825 = 0.125 \\\rightarrow U(B) &= 0.875 \cdot 1900 + 0.125 \cdot -100 = 1650\end{aligned}$$

Assuming  $\neg B$ :

$$\begin{aligned}P(\neg B,P,M) &= P(P \mid M,\neg B) \cdot P(\neg B) \cdot P(M \mid \neg B) = 0.9 \cdot 1 \cdot 0.6 = 0.54 \\P(\neg B,P,\neg M) &= 0.2 \cdot 1 \cdot 0.4 = 0.08 \\P(\neg B,\neg P,M) &= 0.1 \cdot 1 \cdot 0.6 = 0.06 \\P(\neg B,\neg P,\neg M) &= 0.8 \cdot 1 \cdot 0.4 = 0.32\end{aligned}$$

$$\begin{aligned}\rightarrow P(\neg B,P) &= 0.54 + 0.08 = 0.62 \\P(\neg B,\neg P) &= 0.06 + 0.32 = 0.38 \\\rightarrow U(\neg B) &= 0.62 \cdot 2000 + 0.38 \cdot 0 = 1240\end{aligned}$$

So buying the book has the higher utility.

---

### Problem 2.3 (Markov Decision Procedures)

12 pt  
12 min

1. How do Markov decision procedures differ from (simple) decision networks?
2. How does the value iteration algorithm work? (Give an actual equation and explain its role in the algorithm)
3. What is the disadvantage of value iteration that is “fixed” by policy iteration?
4. How can we reduce *partially observable Markov decision procedures* to normal MDPs?

---

#### Solution:

1. In Markov decision procedures, the probabilistic model is a Markov Process (i.e. random variables are indexed over time, Markov Properties)
2. We assign a random utility to each state and update them using the Bellman equation:

$$U(s) = R(s) + \gamma \cdot \max_a \left( \sum_{s'} U(s') \cdot T(s, a, s') \right)$$

Once this iteration has converged, we can compute the “best” action for each state by considering the expected utilities of all possible actions.

3. The policy resulting from value iteration can be stable long before the individual utilities have converged to their precise values.
  4. By introducing belief states representing the probability distribution over the physical state space (i.e. the belief state space has one dimension for each physical state).
-

### 3 Markov Models

#### Problem 3.1 (Stationary)

Define what it means for a Markov model to be *stationary*, and why we are interested in stationarity. 3 pt

---

**Solution:** A [Markov process](#) is called [stationary](#) if the [transition model](#) is independent of time, i.e.  $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$  is the same for all  $t$ . 3 min

We like [stationary Markov processes](#), since they are finite.

---

### Problem 3.2 (Sleeping Patterns Predictions)

Your room mate tends to keep you up by blasting music whenever they are awake. Notably, they tend to sleep a lot less when they are stressed (binary variable  $St$ ), but since you don't talk to each other you never know when they are. You only observe whether they sleep a lot ( $Sl$ ) or little ( $\neg Sl$ ). Stress seems to come in phases and last for a couple of days, so if they are stressed at day  $t$ , they will more likely be stressed at day  $t + 1$  as well (and analogously for  $\neg St$ ).

10 pt  
10 min

1. Model this situation as a Markov Model and explain what the *prediction*, *filtering* and *smoothing* algorithms compute in this scenario.
2. Give the underlying equations for the first two of these algorithms and explain what each variable in the equation represents.

---

#### Solution:

1. **Prediction** Given the amount of sleep up to time  $t_0$ , compute the probability of them being stressed at time  $t_1 > t_0$

**Filtering** Given the of sleep up to now, compute the probability of them being stressed right now

**Smoothing** Given the amount of sleep up to  $t_0$ , compute the probability that they were stressed at an earlier point  $t_1 < t_0$ .

- 2.

$$P(St_{t+1}|Sl_{1:t+1}) = \alpha P(Sl_{t+1}|St_{t+1}) \cdot \sum_{st_t} P(St_{t+1}|st_t, Sl_{1:t}) \cdot P(st_t|Sl_{1:t})$$

---

## 4 Learning

### Problem 4.1 (Tennis Trees)

Consider the following decisions on whether or not to go play tennis. The target is “PlayTennis”. 10 pt  
10 min

Outlook	Temperature	Humidity	Wind	PlayTennis
Sunny	Hot	High	Weak	No
Overcast	Hot	High	Weak	Yes
Rain	Mild	High	Weak	Yes
Rain	Cool	Normal	Strong	No
Rain	Mild	Normal	Weak	Yes

Explain how to apply decision tree learning to this table. In particular, define the notion of *information entropy*.

---

**Solution:** <http://courses.cs.tamu.edu/choe/17spring/633/lectures/slide05.pdf>

$$E := I(\langle \frac{3}{5}, \frac{2}{5} \rangle) = -\frac{3}{5} \log_2(\frac{3}{5}) - \frac{2}{5} \log_2(\frac{2}{5})$$

$$\text{Gain}(\text{Outlook}) = E - \frac{1}{5} I(\langle 0, 1 \rangle) - \frac{1}{5} I(\langle 1, 0 \rangle) - \frac{3}{3} I(\langle \frac{2}{3}, \frac{1}{3} \rangle)$$

$$\text{Gain}(\text{Temperature}) = \dots$$

---

Pick whichever attribute has the highest information gain, split there and build subtrees, iterate...

**Problem 4.2 (Overfitting)**

Explain what *overfitting* means and why we want to avoid it.

4 pt

**Solution:** *Overfitting* is a modeling error that occurs when a function is too closely fit to a limited set of data points. Overfitting the model generally takes the form of making an overly complex model to explain idiosyncrasies in the data under study rather than the underlying mechanisms.

4 min