

Foundations of Artificial Intelligence

Dr. J. Boedecker, Prof. Dr. W. Burgard, PD Dr. M. Ragni
J. Aldinger, J. Boedecker, C. Dornhege, M. Krawez
Summer Term 2016

University of Freiburg
Department of Computer Science

Exercise Sheet 6

Due: Wednesday, July 13, 2016, before the lecture

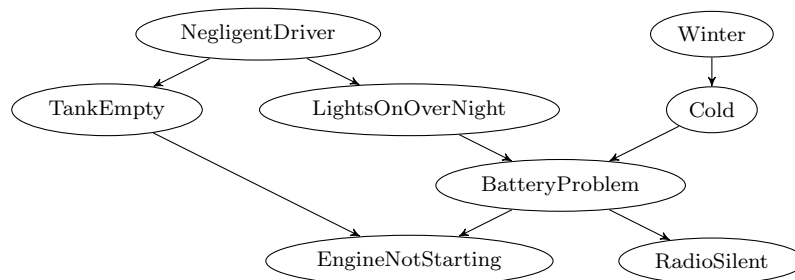
Exercise 6.1 (Conditional Independence)

Suppose you are given a bag containing n unbiased coins, out of which $n - 1$ are normal, with heads on one side and tails on the other, whereas one coin is a fake, with heads on both sides.

- Suppose you reach into the bag, pick out a coin uniformly at random, toss it, and get a head. What is the (conditional) probability that the coin you chose is the fake coin?
- Suppose you continue flipping the coin for a total of k times after picking it and see k heads. Now what is the conditional probability that you picked the fake coin?
- Suppose you wanted to decide whether the chosen coin was fake by flipping it k times. The decision procedure returns FAKE if all k flips come up heads, otherwise it returns NORMAL. What is the (unconditional) probability that this procedure makes an error?

Exercise 6.2 (Bayesian Networks)

Consider the following Bayesian network:



- Determine which of the following conditional independence statements follow from the structure of the Bayesian network ($Ind(U, V | W)$ denotes that U is conditionally independent of V given W , and $Ind(U, V)$ denotes unconditional independence of U and V):
 - $Ind(Cold, Winter)$
 - $Ind(Winter, NegligentDriver)$
 - $Ind(Winter, RadioSilent | BatteryProblem)$
 - $Ind(Winter, EngineNotStarting | BatteryProblem)$
 - $Ind(Cold, NegligentDriver | RadioSilent)$

- (b) Compute $P(\text{EngineNotStarting}|\text{NegligentDriver}, \neg\text{Cold})$. The relevant entries in the conditional probability tables are given below:

$$\begin{aligned}
P(\text{LightsOnOverNight}|\text{NegligentDriver}) &= 0.3 \\
P(\text{LightsOnOverNight}|\neg\text{NegligentDriver}) &= 0.02 \\
P(\text{TankEmpty}|\text{NegligentDriver}) &= 0.1 \\
P(\text{TankEmpty}|\neg\text{NegligentDriver}) &= 0.01 \\
P(\text{BatteryProblem}|\text{Cold}, \text{LightsOnOverNight}) &= 0.9 \\
P(\text{BatteryProblem}|\text{Cold}, \neg\text{LightsOnOverNight}) &= 0.2 \\
P(\text{BatteryProblem}|\neg\text{Cold}, \text{LightsOnOverNight}) &= 0.8 \\
P(\text{BatteryProblem}|\neg\text{Cold}, \neg\text{LightsOnOverNight}) &= 0.01 \\
P(\text{EngineNotStarting}|\text{BatteryProblem}, \text{TankEmpty}) &= 0.9 \\
P(\text{EngineNotStarting}|\text{BatteryProblem}, \neg\text{TankEmpty}) &= 0.7 \\
P(\text{EngineNotStarting}|\neg\text{BatteryProblem}, \text{TankEmpty}) &= 0.8 \\
P(\text{EngineNotStarting}|\neg\text{BatteryProblem}, \neg\text{TankEmpty}) &= 0.05
\end{aligned}$$

Exercise 6.3 (Value iteration algorithm)


Consider the following grid world. The u values specify the utilities after convergence of the value iteration and r is the reward associated with a state. Assume that $\gamma = 1$ and that an agent can perform four possible actions: **North**, **South**, **East** and **West**. With probability 0.7 the agent reaches the intended state, with probability 0.2 it moves to the right of the intended direction, and with probability 0.1 to the left.

$u = 8$	$u = 15$	$u = 12$
$u = 2$	$r = 2$	$u = 10$
$u = 7$	$u = 16$	$u = 11$

Which is the best action an agent can execute if he is currently in the center state of the grid world? Justify your answer. Which utility does the center state have?

Exercise 6.4 (Policy iteration algorithm)

Let $\gamma = 0.5$ and let there be only the actions **East** and **West**. With probability 0.9 the agent reaches the intended state (or stays where he was, if the action would move him out of the grid), and with probability 0.1 he moves in the opposite direction. The reward in the three western states is -0.05 each.



Perform one step of the policy iteration algorithm. The initial policy is given by the arrows in the states. Give the linear system of equations for the first policy evaluation, a solution to the system as well as the first improved policy π_1 .

Exercise 6.5 (Decision Tree Learning)

Two candidates O and M who appeal to different parts of the population run for a political office. The following table shows the preferences of seven voters of different age, income and educational background.

No.	Age	Income	Education	Candidate
1	≥ 35	High	Highschool	O
2	< 35	Low	University	O
3	≥ 35	High	College	M
4	≥ 35	Low	Highschool	M
5	≥ 35	High	University	O
6	< 35	High	College	O
7	< 35	Low	Highschool	M

- (a) Use the learning algorithm from the lecture to compute a minimum-size decision tree correctly classifying all examples wrt the preferred candidate based on the attributes *age*, *income*, and *education*. For the root node, give the information gains associated with all candidate attributes.
- (b) Deduce from the decision tree a logical formula which is satisfied iff candidate O is preferred.

The exercise sheets may and should be worked on in groups of three (3) students. Please write all your names and the number of your exercise group on your solution.