Foundations of Artificial Intelligence

Prof. Dr. J. Boedecker, Prof. Dr. W. Burgard, Prof. Dr. F. Hutter, Prof. Dr. B. Nebel
M. Krawez, T. Schulte
Summer Term 2018

University of Freiburg
Department of Computer Science

Exercise Sheet 6

Due: Wednesday, July 11, 2018, before 12:00

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.

(a) 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?

(b) 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?

(c) 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 (Bayes Rules)

In Freiburg 80% of all cars are red. You see a car at night that does *not* appear red to you. You know that you can correctly identify a red car only in 70% of all cases, given the car is red. However you identify a non-red car correctly in 90% of the cases.

(a) List all conditional and non-conditional probabilities that you can determine directly from the task description. Note: Differentiate between the statement that a car *is red* and the statement that you have *seen a red car*.

(b) Compute the probability that the car is actually red, wenn you perceive a car as red in Freiburg at night.

Exercise 6.3 (Bayesian Networks)

Consider the following Bayesian network:

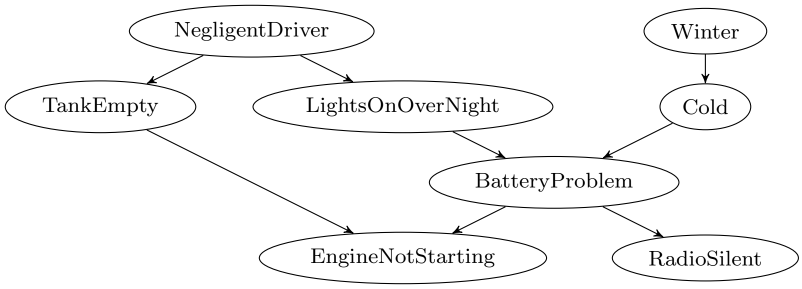

(a) Determine which of the following conditional independence statements follow from the structure of the Bayesian network ($Ind(U, V \mid W)$ denotes that U is conditionally independent of V given W, and $Ind(U, V)$ denotes unconditional independence of U and V):

  (i) $Ind(Cold, Winter)$

  (ii) $Ind(Winter, NegligentDriver)$

  (iii) $Ind(Winter, RadioSilent \mid BatteryProblem)$

  (iv) $Ind(Winter, EngineNotStarting \mid BatteryProblem)$

  (v) $Ind(Cold, NegligentDriver \mid RadioSilent)$

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

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

(c) List all nodes in the Markov blanket for node $LightsOnOverNight$.

Exercise 6.4 (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 a discount $\gamma = 1$. The agent can perform four possible actions: **North**, **South**, **East** und **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 6.5 (Policy iteration algorithm)

Let $\gamma = 0.5$ be the discount and **East** and **West** the only actions. 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.

s_0	s_1	s_2	s_3
$\longleftarrow$	$\longleftarrow$	$\longleftarrow$	$r = +1$

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.

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