English

Norges teknisk-naturvitenskapelige universitet
Institutt for datateknikk og informasjonsvitenskap

TRIAL*examination*	TDT4171 Artificial Intelligence Methods
Date	
Antall timer	4
Antall studiepoeng	7,5
Antall sider:	3
Tillatte hjelpemidler	Ingen
Sensurdato	

All 5 questions, including all sub-questions, should be answered.

Question 1 – Bayesian Networks and Influence diagrams (30%)

a) Describe the syntax and semantics of a Bayesian Network.

b) Model the following problem using a Bayesian network (only the *structure* is required, the quantitative part should not be given):

If my car does not start, it is eithe due to a problem with the starting motor or that I have run empty on gas. I can read off the level of gas in the tank by looking at the gauge (but the gauge does not always work). The most common reason for the starter motor too fail is that the battery is empty. I can check the battery by looking at the head lights: If they shine the battery is OK, if not it is usually because the battery is flat (but there may also be other reasons).

c) Explain the syntax of an influence diagram. What is the *maximum expected utility* principle, and how does it relate to an influence diagram?

d) Extend the Bayesian Network from part (b) to incorporate the following decision problem:

When my car does not start, I have two strategies to get to work: I can walk to the gas station and buy a can of gas and try to fill the tank. This takes 10 minutes; alternatively I can walk to the office, which takes 25 minutes. Driving to office takes 5 min. I wonder whether I should walk should go to the gas station, get gas, and try to get the car to work or walk directly to the office, when I want to get there as quickly as possible.

Question 2 - Learning (15%)

a) What are the main reasons for letting an agent learn from its environment (as opposed to hardcode all the agent's reasoning by hand)

b) Describe the DT learning algorithm informally. What is *Occam's razor*, how does it come into play when learning decision trees?

c) *Gradient descent* is a powerful, general-purpose algorithm, which amongst other things can be used for learning the weights of a perceptron network. Describe the main steps, and the strengths and weaknesses of the learning algorithm.

Question 3 – Case-based reasoning (15%)

Describe the four steps in the CBR cycle: Give their names, explain what happens in each step, and explain how general domain knowledge is/can be used in each of the steps.

Question 4 – Markov Decision Processes (30%)

a) What assumptions are used when modelling *Markov Decision Processes* (also called MDPs)

b) Consider a robot manoeuvring in the grid-world in Figure 1. The robot receives an immediate reward of $100 if it arrives at the top-left corner of the grid, there are no cost involved in moving between the cells of the grid, but penalties of $50 and $10 are given for entering cells (2,2) and (2,3), respectively. In any cell, the robot can choose between the actions **`up`**, **`down`**, **`right`**, and **`left`**. There is a probability *0.7* for the robot to correctly implement the action it actually decides to, and probability *0.1* for doing each of the actions it chose not to implement. Hence, if the robot decides to move **`left`** we have

P(Doing action **`left`** | Decided to do **`left`**) = 0.7,
P(Doing action **`right`** | Decided to do **`left`**) = 0.1,
P(Doing action **`up`** | Decided to do **`left`**) = 0.1,
P(Doing action **`down`** | Decided to do **`left`**) = 0.1,

… and similarly when deciding to do the other action.

Show the first two iterations of the *value iteration* scheme to solve this MDP. Use discount factor 0.1.

	Column 1	*Column 2*	*Column 3*
Row 1			**Gold state: $100**
Row 2		**Trap: -$50**	**Trap: -$10**
Row 3			

Question 5 – Mixed questions (10%)

a) Information Retrieval systems are typically evaluated using two criteria. Define the two criteria. Why is it insufficient to use only one of them to evaluate an IR system?

b) What is "*The Turing Test*", and how does it relate to the term "*Artificial Intelligence*"?