Department of Computer Science
Examination paper for TDT4171 Artificial Intelligence Methods
Examination date: 20/05 - 2021
Examination time (from-to): 0900-1300

Permitted examination support material: A / All support material is allowed

Academic contact during examination: Helge Langseth
Use **Blacboard Collaborate** available from the course's Blackboard page

Technical support during examination: Orakel support services
Phone: 73 59 16 00

If you experience technical problems during the exam, contact Orakel support services as soon as possible before the examination time expires. If you don't get through immediately, hold the line until your call is answered.

OTHER INFORMATION

Make your own assumptions: If a question is unclear/vague, make your own assumptions and specify them in your answer. Only contact academic contact in case of errors or insufficiencies in the question set.

Cheating/Plagiarism: The exam is an individual, independent work. Examination aids are permitted, but make sure you follow any instructions regarding citations. During the exam it is not permitted to communicate with others about the exam questions, or distribute drafts for solutions. Such communication is regarded as cheating. All submitted answers will be subject to plagiarism control. *Read more about cheating and plagiarism here.*

Notifications: If there is a need to send a message to the candidates during the exam (e.g. if there is an error in the question set), this will be done by sending a notification in Inspera. A dialogue box will appear. You can re-read the notification by clicking the bell icon in the top right-hand corner of the screen. All candidates will also receive an SMS to ensure that nobody misses out on important information. Please keep your phone available during the exam.

Weighting: Please note the following about the exam and grading:

- The exam questions can give you up to 50 points, the remaining 50 points have been obtainable through the assignments. The sum of points obtained through the exam and the assignments will be the basis for the final grade.
- The mapping from point total to letter grade is based on NTNUs method with standard limits. To pass the course you need a total of at least 41 points.
- The question-set has 6 parts:
  - Question 1 gives up to 15 points. If the question has more than one part the score will be based on a combined evaluation over all parts.
  - Question 2 gives up to 5 points.
  - Question 3 has 2 parts, each giving up to 2.5 points. The five questions within one part count equally.
  - Question 4 has two parts, each giving either 0 or 2.5 points.
  - Question 5 has 15 parts, each giving either 0 or 1 point.
  - Question 6 gives up to 5 points.

- Note in particular that there is **no negative points for wrong answers** in the multiple choice questions.

ABOUT SUBMISSION

How to answer questions: All question types other than Upload assignment must be answered directly in Inspera. In Inspera, your answers are saved automatically every 15 seconds. **NB!** We advise against pasting content from other programs, as this may cause loss of formatting and/or entire elements (e.g. images, tables).

File upload:
The first two questions ("Modelling -1" and "Modelling-2") require you to upload your answers in files. You can only upload one file per question, therefore must answer **all parts** of a question in that same file.

When working in other programs because parts of/the entire answer should be uploaded as a file attachment – make sure to save your work regularly.

All files must be uploaded before the examination time expires.

The file types allowed are specified in the upload assignment(s). **Note in particular that Genie files are not supported.**

30 minutes are added to the examination time to manage the sketches/calculations/files. The additional time is included in the remaining examination time shown in the top left-hand corner.

NB! You are responsible to ensure that the file(s) are correct and not corrupt/damaged. Check the file(s) you have uploaded by clicking "Download" when viewing the question. All files can be removed or replaced as long as the test is open.

How to digitize your sketches/calculations
How to create PDF documents
Remove personal information from the file(s) you want to upload

Automatic submission: Your answer will be submitted automatically when the examination time expires and the test closes, if you have answered at least one question. This will happen even if you do not click "Submit and return to dashboard" on the last page of the question set. You can reopen and edit your answer as long as the test is open. If no questions are answered by the time the examination time expires, your answer will not be submitted. This is considered as "did not attend the exam".

Withdrawing from the exam: If you become ill, or wish to submit a blank test/withdraw from the exam for another reason, go to the menu in the top right-hand corner and click "Submit blank". This cannot be undone, even if the test is still open.

Accessing your answer post-submission: You will find your answer in Archive when the examination time has expired.

1 Part I:
Model the following domain using a Bayesian network. Make your model as simple and easy to understand as possible:

Milk from a cow may be infected with probability 0.01. There is a test to detect whether the milk is infected, which may give either a positive result (indicating milk is infected) or a negative result (indicating the milk is clean). The test is not perfect: It gives a positive result on clean milk with probability 0.05 and a negative result on infected milk with probability 0.02.

You are asked to give **both** the graphical structure **as well as** the conditional probability distributions here.

Part II:
A test is done to examine the milk, and comes back with the positive result, i.e., indicating that the milk is infected. Calculate the probability that the milk is indeed infected, given the test result. Use the probabilities in the model built above.

Part III:
Extend your model to include the following information:

It is known that the more milk is infected when the weather is humid. Further, the test is less reliable when the air-temperature is high.

You are asked to give **both** the graphical structure **and** the conditional probability distributions in this sub-question. If the question does not give sufficient information to define some probability table(s), you should *suggest relevant numbers yourself.*

Upload your file here. Maximum one file.

The following file types are allowed: **.jpg,.png,.pdf,.txt,.doc,.docx** Maximum file size is **50 GB**

Select file to upload

Maximum marks: 15

2 *The police wants a model for the number of traffic fatalities to expect during a single day. It is known that fatalities per day depends on the season and weather in one way or the other. Also, the danger-level of the average driver depends on road condition and the average speed of the drivers. Finally, the number of fatalities is assumed to be a function of this danger-level and the number of journeys that are taken on a given day.*

Build a Bayesian network model for this domain. Make your model as simple and easy to understand as possible.

What variables ("nodes") and edges ("links") do **you** think are relevant to include in the model in addition to those mentioned above?

Upload your file here. Maximum one file.

The following file types are allowed: **.jpg,.png,.pdf,.txt,.doc,.docx** Maximum file size is **50 GB**

Select file to upload

Maximum marks: 5

3(a) Choose one alternative for each of the following five statements. You get 0.5 points per correct answer. All statements relate to the following Bayesian network wit 6 variables (X1, X2, X3, X4, X5 and X6):

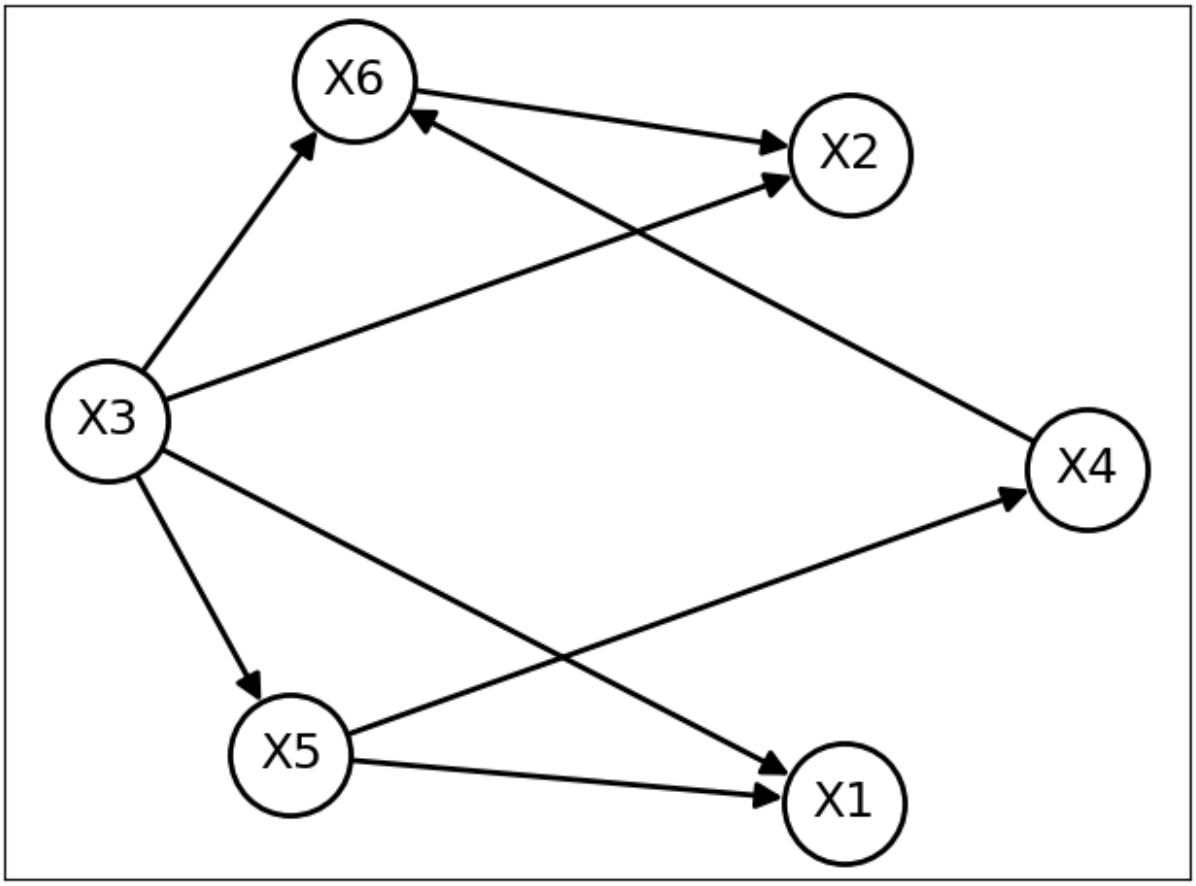

X3 is independent of X4

- True
- False

X3 is conditionally independent of X4 given X1

- True
- False

X2 is conditionally independent of X5 given {X1, X4}

- False
- True

X3 is conditionally independent of X4 given {X1, X2, X5}

- False
- True

X2 is conditionally independent of X4 given {X1, X3, X5, X6}

- False
- True

Maximum marks: 2.5

3(b) Choose one alternative for each of the following five statements. You get 0.5 points per correct answer. All statements relate to the following Bayesian network wit 6 variables (X1, X2, X3, X4, X5 and X6):

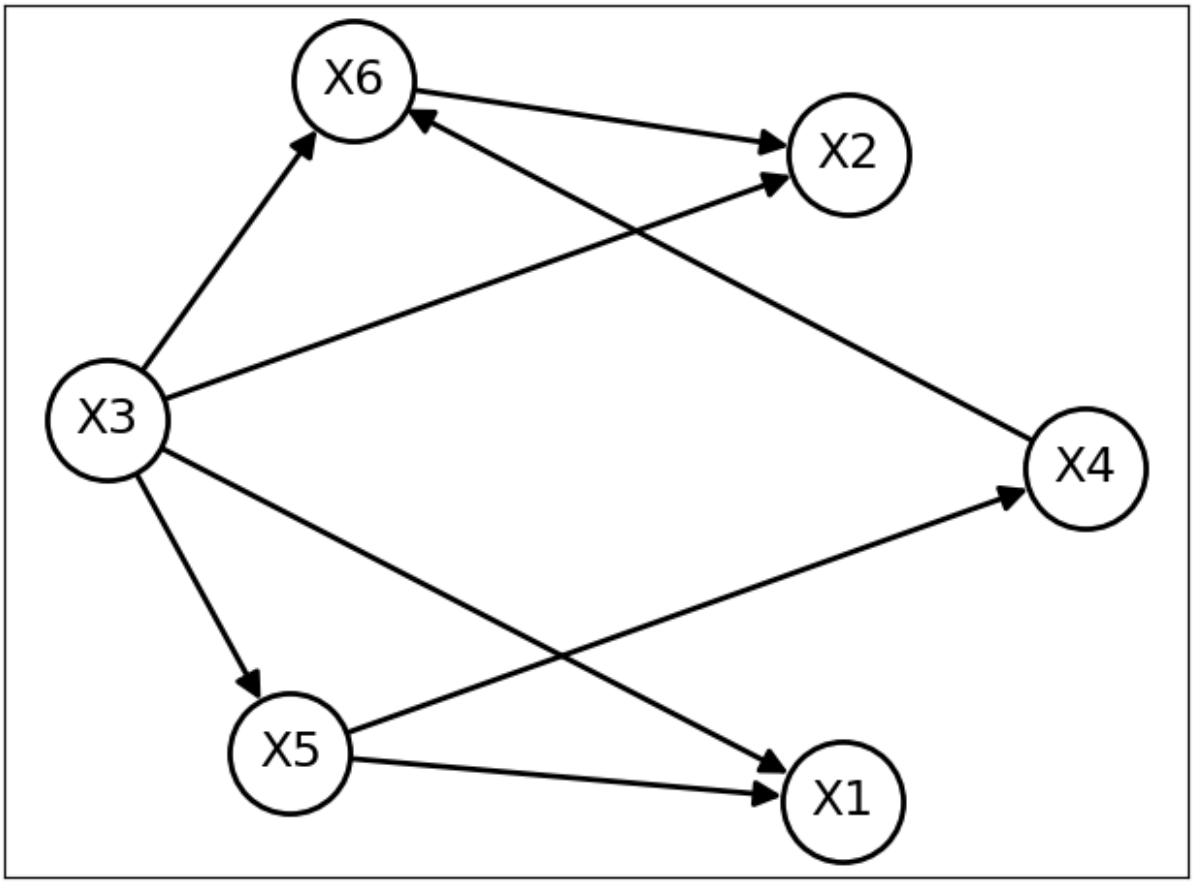

X2 is independent of X4

- False
- True

X1 is conditionally independent of X6 given X2

- False
- True

X1 is conditionally independent of X4 given {X3, X6}

- False
- True

X1 is conditionally independent of X2 given {X3, X5, X6}

- False
- True

X2 is conditionally independent of X5 given {X1, X3, X4, X6}

- True
- False

Maximum marks: 2.5

4(a)

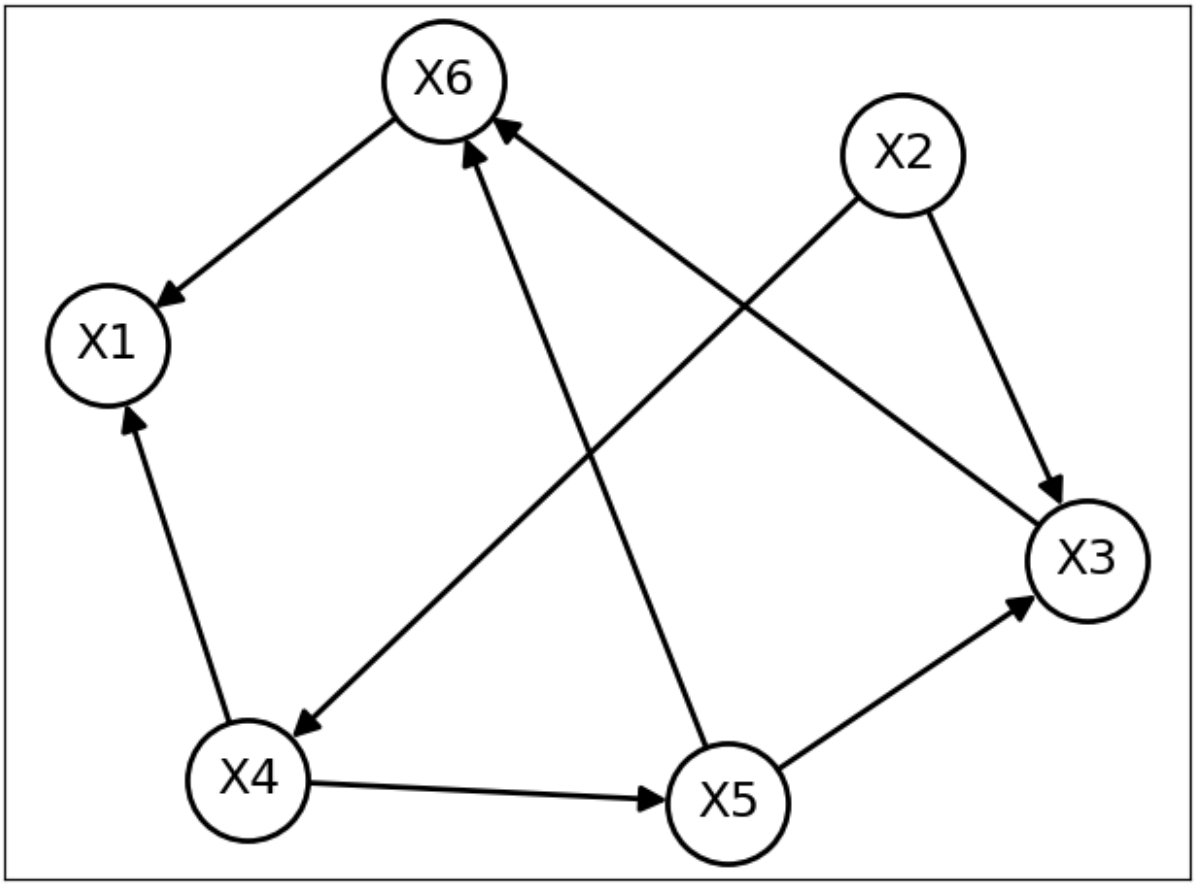

Your answer should state **how many** of the five claims below that are true:

- X4 is independent of X6
- X1 is conditionally independent of X3 given X5
- X1 is conditionally independent of X3 given {X4, X6}
- X1 is conditionally independent of X5 given {X2, X3, X4}
- X1 is conditionally independent of X5 given {X2, X3, X4, X6}

Choose a number from 0 to 5: .

Maximum marks: 2.5

4(b)

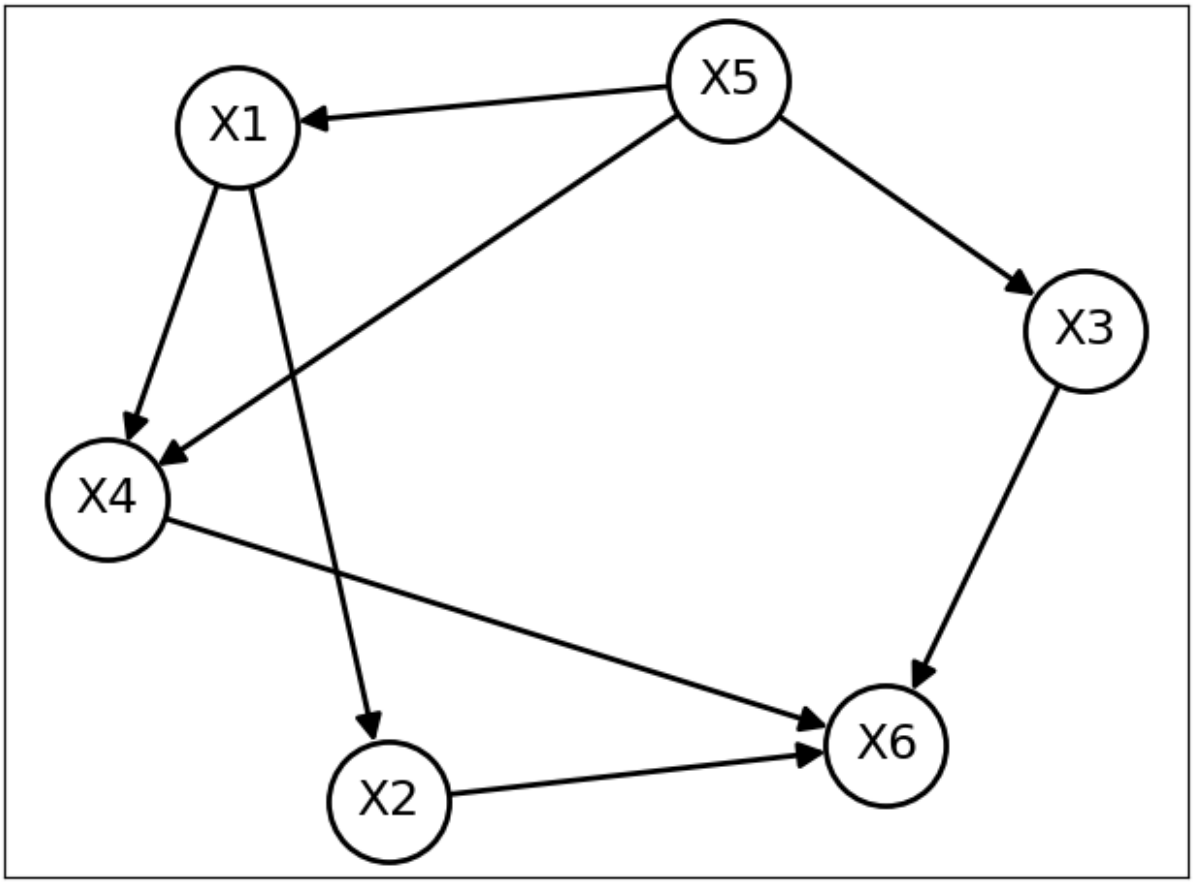

Your answer should state **how many** of the five claims below that are true:

- X1 is independent of X6
- X3 is conditionally independent of X4 given X1
- X2 is conditionally independent of X5 given {X1, X4}
- X1 is conditionally independent of X6 given {X3, X4, X5}
- X3is conditionally independent of X4 given {X1, X2, X5, X6}

Choose a number from 0 to 5: .

Maximum marks: 2.5

5(a) A deterministic boolean function *f* over the *k* boolean variables $X_1, X_2, \ldots, X_k$, i.e. the function $f(X_1, X_2, \ldots, X_k)$, can be represented by a Bayesian network if and only *f* is linearly separable.
Select one alternative:

- False
- True

Maximum marks: 1

5(b) If you want to use Bayes' rule, you first have to make the Markov assumption.
Select one alternative:

- False
- True

Maximum marks: 1

5(c) Case Based Reasoning uses a distance measure to compare a new problem (*query*) with the examples stored n the case-base
Select one alternative:

- True
- False

Maximum marks: 1

5(d) If a system passes the Turing-test, it proves that the strong AI hypothesis is true
Select one alternative:

- False
- True

Maximum marks: 1

5(e) Among the types of inference that can be done with a *Hidden Markov Model*, we find *filtering* and *smoothing*.

Select one alternative:

- True
- False

Maximum marks: 1

5(f) *Overfitting* in machine learning means that a machine learning model is trained using both expert knowledge as well as training data, while *underfitting* means that only information from the training data is used during training.

Select one alternative:

- True
- False

Maximum marks: 1

5(g) For any $k > 3$ it holds that a decision tree with k internal nodes can express any boolean function of $k-1$ binary variables.

Select one alternative:

- True
- False

Maximum marks: 1

5(h) If A and B are discrete random variables with $P(A = a, B = b) = P(A = a)$ we know that $\frac{1}{2} \leq P(B = b|A = a) \leq \frac{1}{\sqrt{2}}$

Select one alternative:

- True
- False

Maximum marks: 1

5(i) Deep learning cannot be used to analyze text data

Select one alternative:

- False
- True

Maximum marks: 1

5(j) Reinforcement learning is a kind of deep learning

Select one alternative:

- True
- False

Maximum marks: 1

5(k) When using Hidden Markov Models we have to make both the Markov-assumption and the sensor Markov assumption.

Select one alternative:

- False
- True

Maximum marks: 1

5(l) Rational behavior can be described using mathematical terms

Select one alternative:

- True
- False

Maximum marks: 1

5(m) If one succeeds in building a fully rational agent, this agent will pass the Turing test if and only if the agent has calculated that passing the Turing test maximises its expected utility.

Select one alternative:

- False
- True

Maximum marks: 1

5(n) Assume we have *n* variables in a Bayesian network, X_1, X_2 , ..., X_n and that the first *k* variables X_1, X_2 , ..., X_k for a given $k < n$ are such that they do not have any parents in the network. Then the model always asserts that
$P(x_1, x_2, \ldots x_n) = P(x_1) \cdot P(x_2) \cdot P(x_3) \cdot \ \ldots \ \cdot P(x_n)$.
Select one alternative:

- False
- True

Maximum marks: 1

5(o) A decision tree describing a concept over *n* binary attributes must have a size (measured by the number of internal nodes) of at least *n* log(*n*)
Select one alternative:

- False
- True

Maximum marks: 1

5(p) When a learning algorithm for decision-trees uses reduction in remaining entropy to find the next split-node, then that is just a heuristic and does not guarantee that the learned decision tree will be optimal.
Select one alternative:

- False
- True

Maximum marks: 1

5(q) Weight-regularization as used in deep learning is a technique to avoid overfitting.
Select one alternative:

- True
- False

Maximum marks: 1

5(r) The reason for making the assumption of stationarity when building Hidden Markov models is that stationarity induces conditional independences in the model that makes the calculations more efficient.
Select one alternative:

- False
- True

Maximum marks: 1

5(s) Case-based reasoning is used for machine learning
Select one alternative:

- True
- False

Maximum marks: 1

5(t) In this part of the CBR cycle the system stores the solved problem in the case-base:
Select one alternative:

- Retain
- Retrieve
- Revise
- Reuse

Maximum marks: 1

5(u) In this part of the CBR cycle the system finds the previous solved case (or cases) that are most similar to the problem it is confronted with now:
Select one alternative:

- Retain
- Reuse
- Revise
- Retrieve

Maximum marks: 1

5(v) In this part of the CBR cycle the system adapts previous solution(s) to the new problem-case:
Select one alternative:

- Reuse
- Revise
- Retain
- Retrieve

Maximum marks: 1

5(w) The reason why deep learning typically works better than a single perceptron in terms of training loss is that the gradient descent algorithm scales well with respect to the number of weights in the model.
Select one alternative:

- True
- False

Maximum marks: 1

5(x) The perceptron has limited expressive power as long as we insist that the transfer-function should be monotonically increasing, but if we do not put any restrictions on this transfer-function the perceptron can represent any continuous function.
Select one alternative:

- False
- True

Maximum marks: 1

5(y) When one is to make a deep learning to classify images it is common to use convolutions among other things to ensure that the same mathematical operations are done on all parts of the input-image.

Select one alternative:

- False
- True

Maximum marks: 1

5(z) If two agents both are rational, then they must have the same underlying utility function.

Select one alternative:

- True
- False

Maximum marks: 1

5(aa) The local and the global semantics of a Bayesian network are identical in the sense that you can guarantee one if the other holds.

Select one alternative:

- False
- True

Maximum marks: 1

5(ab) When a probability *P(x)* is *subjective* then two intelligent agents can hold different opinions about how the probability should be quantified even though both agents are rational.

Select one alternative:

- False
- True

Maximum marks: 1

5(ac) The computational cost of doing smoothing in Hidden Markov Models grows linearly in the number of time-steps for which you have observed data.

Select one alternative:

- True
- False

Maximum marks: 1

5(ad) There always exists a utility-function that can be used to explain the behavior of a human being precisely.

Select one alternative:

- True
- False

Maximum marks: 1

5(ae) There always exists a utility-function that can be used to explain the behavior of a rational agent.

Select one alternative:

- False
- True

Maximum marks: 1

5(af) A Bayesian network will often give a compact representation of the joint distribution modelled. Nevertheless, the worst-case space complexity of the model is exponential in the number of variables in the domain.

Select one alternative:

- True
- False

Maximum marks: 1

5(ag) To use the algorithm "Q-learning" when working with reinforcement learning, one must make the Markov-assumption

Select one alternative:

- False
- True

Maximum marks: 1

5(ah) The algorithms "Value iteration" and "Policy iteration" can both be used to solve sequential decision problems, and rest on exactly the same set of assumptions.

Select one alternative:

- False
- True

Maximum marks: 1

5(ai) When solving a Markov Decision Problem with infinite time horizon by discounting, the solution no longer describes a rational behavior.

Select one alternative:

- False
- True

Maximum marks: 1

6 In what ways do you think artificial intelligence will change the future? Be ***realistic*** and use terminology that is established in the field. Start from your own reflections regarding the weak and the strong AI hypothesis. Describe both negative and positive effects, if possible both regarding everyday life and work-life. What *measures* should be implemented to maximize positive and minimize negative effects? ***Give your reasons.***

Note: Your answer is limited to 500 words.

Fill in your answer here

Words: 0/500

Maximum marks: 5