

QUESTION 1

Question 1 can in total give 20 points. All ten sub-questions are to be answered by either "True" or "False". No explanation is required. If you answer a sub-question correctly you get +2 points. A wrong answer gives -3 points. Not answering a sub-question results in 0 points. Nevertheless, please note that your total score from Question 1 will never be negative. A point total of zero will be carried forward if the sum of the scores from the ten sub-questions is negative.

- a) Every Boolean function can be represented by a Bayesian network.
- b) Suppose that none of the variables X_1, X_2, \dots, X_k have parents in a Bayesian network that in total consists of $n > k$ variables. Then the Bayesian net asserts that

$$\mathbf{P}(X_1, X_2, \dots, X_k) = \mathbf{P}(X_1) * \mathbf{P}(X_2) * \dots * \mathbf{P}(X_k).$$

- c) Bayesian networks are particularly well suited for detecting objects in a greyscale image if the pixels of the image are represented as an array of intensity values.
- d) A perfectly rational agent will never loose in a game of Poker.
- e) An agent that senses only partial information about the state of its environment can still be perfectly rational.
- f) A system must think like a human to pass the Turing Test reliably.
- g) The Chinese Room argument explains why the Weak AI hypothesis is not achievable.
- h) Deep neural networks are called deep because they require deep engineering knowledge to be learned efficiently.
- i) A key ingredient of convolutional neural networks is weight sharing.
- j) The gradient of the loss function wrt. the network weights is not sufficient to learn the weights of a deep learning network.

QUESTION 2

Question 2 can give 18 points in total. Each of the three sub-questions has a maximum score of 6 points.

An agent inhabits a world with two states, named "**S**" and "**¬S**" (NOT S). At each point in time the agent does exactly one of the following two actions:

A1: This action has no effect. The agent remains in the same state as at the previous time-step.

A2: The state of the world is flipped. If the agent was in **S**, he will be in **¬S** after **A2** is performed; if he was in **¬S** he will be in **S** after **A2** is performed.

At each point in time the agent receives a reward, R , depending on which state the agent was in: If the agent is in state **S** the reward is 3, otherwise the reward is 2.

We want to find the optimal strategy π , that for each possible state of the world determines the rational action to perform.

a) Explain why the problem can be represented as a *Markov Decision Problem*. Make any additional assumption(s) you make explicit.

b) The Bellman-equations state that

$$U^i(s) = R(s) + \gamma \max_a \sum_{s'} P(s'|s, a) \cdot U^i(s')$$

Explain how this equation can help the agent determine the optimal strategy.

c) Calculate the U-values for the first three iterations of the value iteration scheme. After each iteration, you should also produce the associated strategy.

Set $\gamma = 0.5$, and start from initial U-values of zero.

QUESTION 3

Question 3 can give 22 points in total. The scoring rule for each sub-question is given below.

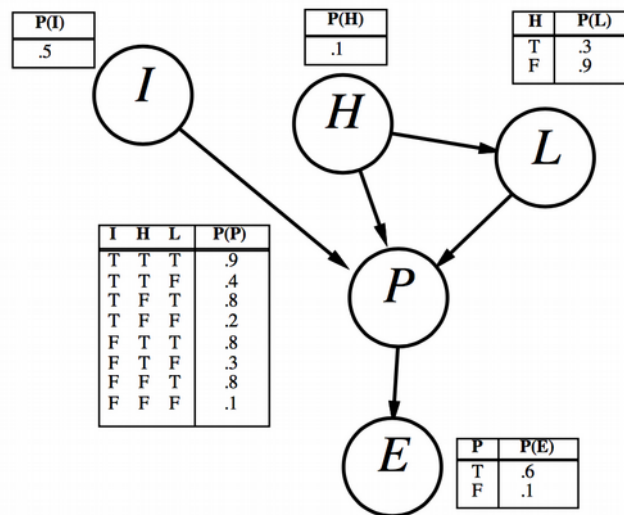


Figure 1: A simple Bayesian network with Boolean variables

In Question 3 we consider the simple Bayesian network model shown in Figure 1. The model describes the process of getting elected as the American president. All the variables are binary, and have the following meaning:

- I : Intelligent
- H : Honest
- P : Popular
- L : LotsOfCampaignFunds
- E : Elected

For each node we have given the associated conditional distribution. Note that only the required numbers are given, with the syntactic definition that, for instance, “ $P(H)$ ” is a shorthand for $P(H=\text{True})$. With this number being given as 0.1, we know that (since all variables are binary) $P(H=\text{False}) = 1 - P(H=\text{True}) = 0.9$.

- a) Considering only the *structure* of the model (and not the allocated conditional probability tables), which - if any - of the following statements are True:
- $\mathbf{P}(I, L) = \mathbf{P}(I) \mathbf{P}(L)$
 - $\mathbf{P}(E|P, L) = \mathbf{P}(E|P, L, H)$
 - $\mathbf{P}(P|I, H) \neq \mathbf{P}(P|I, H, L)$

This sub-question is worth 3 points.

b) Calculate $P(I=True, H=True, L=False, P=True, E=False)$.

This sub-question is worth 3 points.

c) Calculate the probability that someone is intelligent given that they are honest, have few campaign funds, and are elected.

This sub-question is worth 4 points.

d) Suppose that we want to add the variable $R=RigsElection$ to the network. Draw the new network and provide new or modified conditional probability tables as needed.

This sub-question is worth 6 points.

e) True or False: If there are two candidates in the race, then making two copies of the network will correctly represent the joint distribution over the two sets of variables.

This sub-question is to be answered by "True" or "False" only. A correct answer gives 2 points, a wrong answer results in -3 points. No answer gives zero points.

f) Suppose that you are considering making a large *Donation* to help elect your candidate. Draw the influence diagram/decision network (only the structure, conditional probability tables not required) obtained by adding the appropriate node(s) to Figure 1.

This sub-question is worth 4 points.

QUESTION 4:

Question 4 can give 20 points in total, with max-score for each sub-question given below.

a) Consider the four different training sets given in Figure 2 below.

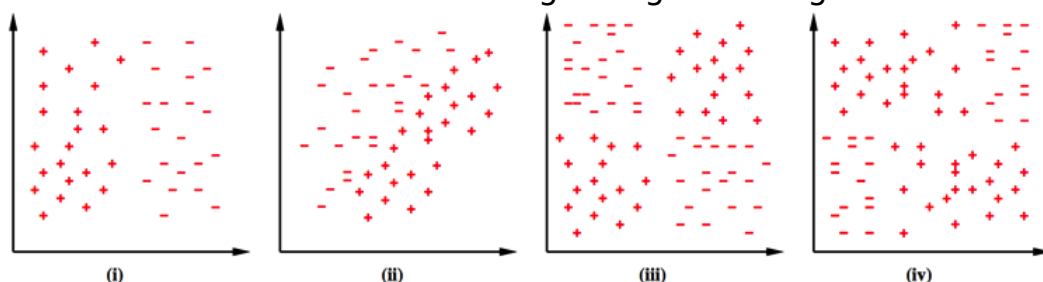


Figure 2 Four different training sets with two numerical input-attributes

Each training set has 2 continuous input attributes, one measured along the x-axis, the other along the y-axis. Each training example is either a member of the *positive class* (shown by a "+") or the *negative class* (shown by a "-").

For each of the four training sets, you should choose one machine

learning method that can classify that training set correctly. Choose among the following:

- Decision trees
- Perceptron
- Case-based reasoning with Euclidean distance measure
- Deep learning with convolutional filters

If more than one approach will be successful, you chose the simplest one. You do not have to select the same classifier for all four data sets.

This sub-question is worth 8 points: 2 points per dataset in Figure 2.

b) **Briefly** explain the four steps of the CBR cycle.

This sub-question is worth 6 points.

c) You wish to compare your new machine learning algorithm to decision trees. Your algorithm, called YOU, has an important parameter that users can set. Call this parameter θ and assume it only has three possible settings (1, 2, and 3). You have three standard datasets used in machine-learning research, call them A, B, and C. You split each data-set into a training set and a test set, train using the training set, and calculate the following test set accuracies:

Dataset	Decision Tree	YOU with $\theta=1$	YOU with $\theta=2$	YOU with $\theta=3$
A	79%	75%	81%	91%
B	62%	67%	56%	61%
C	89%	84%	93%	87%

Which interpretation of these results is best?

1. YOU is better because for EACH dataset SOME setting of θ produces results better than the Decision Tree.
2. The Decision Tree is better because for SOME setting of θ it beats YOU on each dataset.
3. An additional dataset (a *validation set*) should have been used to independently choose a good setting for θ for each Dataset A, B and C. Then YOU should report only the score with that value of θ .
4. θ should *not* be a settable parameter in YOU and instead should be always set to 3 because on average it produces the best accuracy. Therefore YOU is better than Decision Trees for dataset A, but poorer for datasets B and C.

State your selection (1, 2, 3 or 4) and **briefly** explain your answer.

This sub-question is worth 6 points.