INFO8006 Introduction to Artificial Intelligence

January 2019

Instructions

- *Duration: 4 hours.*
- *Answer the questions on separate sheets, labeled with the question number, your first name, last name and student id.*
- *Answer in English or in French.*
- *Follow the same mathematical notation conventions as in the course, or properly define your conventions otherwise.*
- *Non-programmable calculators are allowed.*
- *Notes or documents of any kind are forbidden.*

Question 1 [3 points]

Multiple choice questions. Choose one of the four choices. Correct answers are graded $+\frac{3}{10}$*, wrong answers are graded* $-\frac{1}{10}$ *and the absence of answers is graded* 0*. The total of your grade for Question 1 is bounded below at* 0/3*.*

1. In a partially observable and stochastic environment, a rational agent is ...
   (a) an agent that chooses whichever action that minimizes the expected value of the performance measure, given its percept history.
   (b) an agent that chooses whichever action that maximizes the value of the performance measure, given its percept history.
   (c) an agent that chooses whichever action that maximizes the value of the performance measure, regardless of its percept history.
   (d) an agent that chooses whichever action that maximizes the expected value of the performance measure, given its percept history.

2. Consider the game of Pacman in the absence of ghosts and with a single food dot in the maze. In state s, Pacman is located at position (i_s, j_s) while the food dot is located at (x_s, y_s). At timestep t, the score of the game is $500 - t$. The game ends when the food is eaten. Which of the following heuristics is not admissible?
   (a) $h(s) = |i_s - x_s| + |j_s - y_s|$
   (b) $h(s) = \sqrt{(i_s - x_s)^2 + (j_s - y_s)^2}$
   (c) $h(s) = (i_s - x_s)^2 + (j_s - y_s)^2$
   (d) $h(s) = \min(|i_s - x_s|, |j_s - y_s|)$

3. Which of the following propositions is not equivalent to the others?
   (a) The sentence α entails the sentence β.
   (b) α is true in all models where β is true.
   (c) β follows logically from α.
   (d) $\alpha \vDash \beta$.

4. In Monte Carlo Tree Search, at a node n during the selection step, the UCB1 policy picks the child node n' of n that maximizes
   $$\frac{Q(n', p)}{N(n')} + c\sqrt{\frac{2 \log N(n)}{N(n')}}.$$
   Which of the following is true?
   (a) The first term encourages the exploitation of higher-reward nodes, while the second encourages the exploration of less-visited nodes.
   (b) The first term encourages the exploration of less-visited nodes, while the second term encourages the exploitation of higher-reward nodes.

(c) The first term encourages the exploitation of highly-visited nodes, while the second term encourages the exploration of lesser-rewarding nodes.

(d) The first term encourages the exploration of lesser-rewarding nodes, while the second term encourages the exploitation of highly-visited nodes.

5. The Bayes' rule states that ...

   (a) $P(x|y) = P(y|x)P(x) \,/\, P(y)$.
   (b) $P(x|y) = P(y|x)P(y) \,/\, P(x)$.
   (c) $P(x|y) = P(y) \,/\, P(y|x)P(x)$.
   (d) $P(y|x) = P(x|y)P(x) \,/\, P(y)$.

6. Which of the following is false?

   (a) Variable elimination is an exact inference algorithm that can be used both for discrete and continuous variables (under appropriate assumptions).
   (b) In variable elimination, the elimination ordering can greatly affect the computational complexity.
   (c) In likelihood weighting sampling, the evidence variables are all taken into account by the sampling distribution.
   (d) Gibbs sampling settles into a dynamic equilibrium in which the long-run fraction of time spent in each state is exactly proportional to its posterior probability.

7. In a first-order Markov process, as time passes ...

   (a) the state distribution always converges to a fixed point, called the stationary distribution.
   (b) the state distribution always converges to a fixed point, called the stationary distribution. This distribution is always of maximum uncertainty.
   (c) the state distribution sometimes converges to a fixed point, called the stationary distribution.
   (d) the state distribution sometimes converges to a fixed point, called the stationary distribution. This distribution is always of maximum uncertainty.

8. In a Markov decision process, the goal is ...

   (a) to find the optimal next action to take.
   (b) to find the action that maximizes the reward in the next state.
   (c) to find an optimal policy that maps states to actions.
   (d) to find an optimal plan, or sequence of actions, from start to goal.

9. Logistic regression models $P(Y = 1|\mathbf{x})$ as ...

   (a) $P(Y = 1|\mathbf{x}) = \sigma(\mathbf{w}^T\mathbf{x} + b)$, where $\sigma(x) = \frac{1}{1+\exp(-x)}$.
   (b) $P(Y = 1|\mathbf{x}) = \text{sign}(\mathbf{w}^T\mathbf{x} + b)$.
   (c) $P(Y = 1|\mathbf{x}) = \tanh(\mathbf{w}^T\mathbf{x} + b)$.
   (d) $P(Y = 1|\mathbf{x}) = \mathbf{w}^T\mathbf{x} + b$.

10. State-of-the-art approaches for speech recognition are based on ...

   (a) A*.
   (b) hidden Markov models.
   (c) neural networks.
   (d) the Kalman filter.

Question 2 [3 points]

The game "21" is played as a misère game with any number of players who take turns saying a number. The first player says "1" and each player in turn increases the number by 1, 2, or 3, but may not exceed 21; the player who says "21" or a larger number loses.

1. Define the search problem associated with the 2-player version of the "21" game.
2. For this subquestion and the following, consider the game of "5" (still in its 2-player version) which has the same rule except that you should not say 5 or more. Show the whole game tree.
3. (a) Using the minimax algorithm, mark on your tree the backed-up values, and use those values to choose the best starting move
   (b) Assume alpha–beta pruning were applied in optimal order. Draw the game tree containing only the nodes that would be evaluated.

Question 3 [3 points]

In order to quickly estimate the exam results, the teaching assistant wants to predict whether or not a student will pass the exam (Y) based on the following evidence:

- X_1: The student got a grade greater or equal than 10/20 at the first exercise (Yes/No).
- X_2: The student got a grade greater or equal than 10/20 in average for the projects (Yes/No).
- X_3: The student liked the course (Yes/No).

To simplify his model, the teaching assistant assumes that X_1, X_2 and X_3 are pairwise conditionally independent given Y. Answer the following questions:

1. Draw a Bayesian network that represents the joint distribution of Y, X_1, X_2 and X_3, and that incorporates the independence assumptions listed above. What is the name of this model?
2. Use Table 1 to compute an estimate of the conditional probability tables in the Bayesian network.
3. Give the prediction rule of Y for this model. What are the parameters of the model? Predict whether or not Pierre will pass the exam, knowing that he got 15/20 at the first question, liked the course and got 14/20 in average for the projects.

Y	Passed	Failed
Number students	78	22
Number of students who got $\geq 10/20$ at the first exercise	54	10
Number of students who got $\geq 10/20$ at the projects	70	15
Number of students who liked the course	45	4

Table 1: Data collected for 100 students.

Question 4 [4 points]

1. Define mathematically i) first-order Markov processes and ii) the inference tasks of prediction and filtering. Discuss how the latter can be useful to an agent.
2. Derive the recursive update equation of the Bayes filter, assuming discrete variables. Does the Bayes filter generalize to continuous variables? If yes, outline how? If not, why?
3. Let us consider a continuous Pacman world in which Pacman cannot directly observe ghosts. However, Pacman is equipped with a device that yields noisy estimates of the ghost positions. Assuming
   - a one-dimensional world with a single ghost,
   - a Gaussian prior of constant variance σ_0^2 for the ghost position,
   - a Gaussian transition model that nudges the ghost with random perturbations of fixed variance σ_x^2,
   - a sensor model that yields measurements with Gaussian noise of fixed variance σ_z^2 of the ghost position,

   derive the update equations from timestep t to $t+1$ of the parameters of the belief distribution of the ghost position, including its mean μ_{t+1} and its variance σ_{t+1}^2. You are free to use identities from Appendix A if needed.
4. Let us examine the behavior of the variance update.
   - As $t \to \infty$, σ_t^2 converges to a fixed point σ^2. Calculate the value of σ^2.
   - Give a qualitative explanation for what happens as i) $\sigma_x^2 \to 0$ and ii) as $\sigma_z^2 \to 0$.

Question 5 [4 points]

1. Formally define what is i) a Markov Decision Process (MDP) and ii) an optimal policy.
2. Sometimes MDPs are formulated with a reward function $R(s, a)$ that depends on the action taken or with a reward function $R(s, a, s')$ that also depends on the outcome state.
   (a) Write the Bellman equations for these formulations.
   (b) Show how an MDP with reward function $R(s, a, s')$ can be transformed into a different MDP with reward function $R(s, a)$, such that optimal policies in the new MDP correspond exactly to optimal policies in the original MDP. Then do the same to convert MDPs with $R(s, a)$ into MDPs with $R(s)$.
3. Discuss whether an agent taking actions in the real world can be modeled as a MDP. If yes, how? If not, why?

Question 6 [3 points]

The Bayesian network shown in Figure 1 models the quantity of salt in a chocolate cake. In this model, X_1, X_2, X_3 respectively denote the quantity of salt in the chocolate before cooking, the quantity of salt in the eggs and the quantity of salt in the egg white beaten in snow. To cook a chocolate cake, you first have to mix chocolate with egg yellow, together with sugar and butter. The quantity C of salt in the resulting dough is modeled as $C = a_2X_1 + a_3X_2 + \mathcal{N}(\mu_C, \sigma_C^2)$. When the egg white is beaten in snow, a bit of salt is usually added, which we model as $X_3 = a_1X_2 + \mathcal{N}(\mu_3, \sigma_3)$. Finally, the quantity of salt in the chocolate cake is modeled as $Y = a_4C + a_5X_3 + \mathcal{N}(\mu_y, \sigma_y^2)$. Answer the following questions about this model:

1. Assume the parameters of the model are known, give the prediction rule of Y given X_1 and X_2. Simplify this rule and give a minimal set of parameters that are required to predict Y given X_1, X_2. What is the name of the resulting predictive model?

2. From a data set of N points $d = \{(x_{1,1}, x_{2,1}, y_1), \ldots, (x_{1,N}, x_{2,N}, y_N)\}$ (that you assume to be iid) explain mathematically how you would compute the parameters of the simplest model.

3. Is this model always realistic? Given a large amount of data, what other model could you use instead? What if you only have a small amount of data?

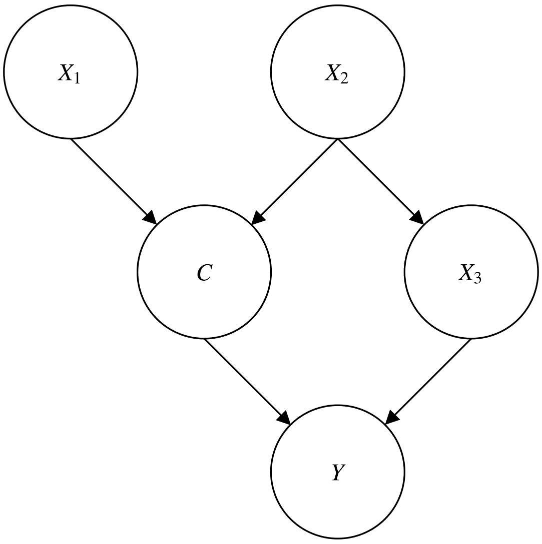

Figure 1: Bayesian network of the final grade calculation.

A Cheat sheet for Gaussian models (Bishop, 2006)

Given a marginal Gaussian distribution for $\mathbf{x}$ and a linear Gaussian distribution for $\mathbf{y}$ given $\mathbf{x}$ in the form

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mu, \mathbf{\Lambda}^{-1})$$
$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}|\mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{L}^{-1})$$

the marginal distribution of $\mathbf{y}$ and the conditional distribution of $\mathbf{x}$ given $\mathbf{y}$ are given by

$$p(\mathbf{y}) = \mathcal{N}(\mathbf{y}|\mathbf{A}\mu + \mathbf{b}, \mathbf{L}^{-1} + \mathbf{A}\mathbf{\Lambda}^{-1}\mathbf{A}^T)$$
$$p(\mathbf{x}|\mathbf{y}) = \mathcal{N}(\mathbf{x}|\mathbf{\Sigma}\left(\mathbf{A}^T\mathbf{L}(\mathbf{y} - \mathbf{b}) + \mathbf{\Lambda}\mu\right), \mathbf{\Sigma})$$

where

$$\mathbf{\Sigma} = (\mathbf{\Lambda} + \mathbf{A}^T\mathbf{L}\mathbf{A})^{-1}.$$