INFO8006 Introduction to Artificial Intelligence

Exam of August 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. Which of the following is false?
   (a) Simple reflex agents can only work if the correct decision can be made on the basis of the current percept only.
   (b) Simple reflex agents can work if the environment is fully observable.
   (c) Model-based agents can only work if the correct decision can be made on the basis of the current percept only.
   (d) Model-based agents handle partial observability by keeping track of the part of the world they cannot perceive now.

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 the best to use?
   (a) $h_1(s) = \sqrt{(i_s - x_s)^2 + (j_s - y_s)^2}$
   (b) $h_2(s) = 0$
   (c) $h_3(s) = \max(|i_s - x_s|, |j_s - y_s|)$
   (d) $h_4(s) = \min(h_1(s), h_3(s))$

3. The unsatisfiability theorem states that...
   (a) $\alpha \vDash \beta$ iff $(\neg\alpha \wedge \beta)$ is unsatisfiable.
   (b) A sentence γ is unsatisfiable iff $M(\gamma) = \{\}$.
   (c) α entails β iff $M(\alpha \wedge \neg\beta) = \{\}$.
   (d) β follows logically from α iff $(\neg\alpha \wedge \beta)$ is unsatisfiable.

4. In two-player turn-taking zero-sum stochastic games, ...
   (a) an EXPECTIMINIMAX agent is always successful.
   (b) an EXPECTIMINIMAX agent maximizes the worst-case outcome.
   (c) an EXPECTIMINIMAX agent is optimal, even if its opponent is suboptimal and predictable.
   (d) an EXPECTIMINIMAX agent makes rational moves, assuming an optimal opponent.

5. Which of the following holds true? (for all a, b, c)
   (a) $P(a|b, c) = P(b|a, c)P(a|c) \,/\, P(b|c)$
   (b) $P(a|b, c) = P(b|a, c)P(c|a) \,/\, P(c|b)$
   (c) $P(a|b, c) = P(a|b, c)P(b|c) \,/\, P(a|c)$
   (d) $P(a|b, c) = P(c|a, b)P(b|a) \,/\, P(b|c)$

6. The Normal distribution $\mathcal{N}(\mu, \sigma)$ is described by the density function

(a) $p(x) = \frac{1}{\sigma} \exp\left(-(z + \exp(-z))\right)$, with $z = \frac{x-\mu}{\sigma}$.

(b) $p(x) = \frac{1}{2\sigma} \exp\left(-\frac{|x-\mu|}{\sigma}\right)$.

(c) $p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$

(d) $p(x) = 1 \,/\, \pi\sigma \left[1 + \left(\frac{x-\mu}{\sigma}\right)^2\right]$.

7. The time complexity of smoothing using the forward-backward algorithm applied to a sequence of size t is
   (a) $O(t)$.
   (b) $O(t \log t)$.
   (c) $O(t^2)$.
   (d) $O(\log t)$.

8. In Markov Decision Processes, ...
   (a) a small discount factor γ effectively increases the horizon of the agent.
   (b) a large discount factor γ typically results in a greedy policy.
   (c) a small discount factor γ penalizes short-term rewards.
   (d) a small discount factor γ typically results in a greedy policy.

9. Neural networks are typically trained using...
   (a) the Newton-Raphson method.
   (b) gradient descent (or a variant thereof).
   (c) the bisection method.
   (d) the simplex algorithm.

10. In speech recognition systems based on HMMs, decoding is performed using...
    (a) the Variable Elimination algorithm.
    (b) the Viterbi algorithm.
    (c) a Bayes filter.
    (d) the Policy Iteration algorithm.

Question 2 [3 points]

Consider the following two-player turn-taking game which initial configuration is shown in Figure **??**. Player A moves first. Each player must move his token to an open adjacent cell in either direction. If the opponent occupies an adjacent cell, then a player may jump over the opponent to the next open cell if any. (For example, if A is on 3 and B is on 2, then A may move back to 1.) The game ends when a player reaches the opposite end of the board. If player A reaches cell 4 first, then the value of the game to A is +1; if player B reaches cell 1 first, then the value of the game to A is -1.

1. Define the search problem associated with this game.
2. Draw the complete game tree, using the following conventions:
   - Put each terminal state in a square box and write its game value in a circle.
   - Put loop states (states that already appear on the path to the root) in double square boxes. Since their value is unclear, annotate each with a '?' in a circle.
3. Now mark each node with its backed-up minimax value (also in a circle). Explain how you handled the '?' values and why.
4. Explain why the standard minimax algorithm would fail on this game tree and briefly sketch how you might fix it, drawing on your answer to (3).
5. This 4-cell game can be generalized to n cells for any $n > 2$. Prove that A wins if n is even and loses if n is odd. (Hint: Use induction.)

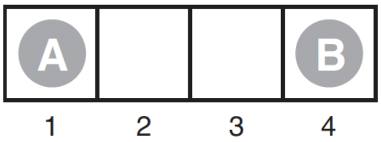

Figure 1: Initial configuration.

Question 3 [3 points]

Two astronomers in different parts of the world make measurements M_1 and M_2 of the number of stars N in some small region of the sky, using their telescopes. **Normally** (i.e. when the focus is right), there is a small possibility e of error by up to one star in each direction. Each telescope can also (with a much smaller probability f) be badly out of focus (events F_1 and F_2), in which case the scientist will undercount by three or more stars (or if N is less than 3, fail to detect any stars at all). Consider the three networks shown in Figure **??**.

1. Which of these Bayesian networks is a wrong representation of the preceding information?
2. Which is the best network? Explain.
3. Write out a conditional distribution for $P(M_1|N)$, for the case where $N \in \{1, 2, 3\}$ and $M_1 \in \{0, 1, 2, 3, 4\}$. Each entry in the conditional distribution should be expressed as a function of the parameters e and/or f.
4. Suppose $M_1 = 1$ and $M_2 = 3$. What are the possible numbers of stars if you assume no prior constraint on the values of N?
5. What is the most likely number of stars, given these observations? Explain how to compute this, or if it is not possible to compute, explain what additional information is needed and how it would affect the result. (Hint: to find the final answer, use bounds on the probability of the possible values found in previous question.)

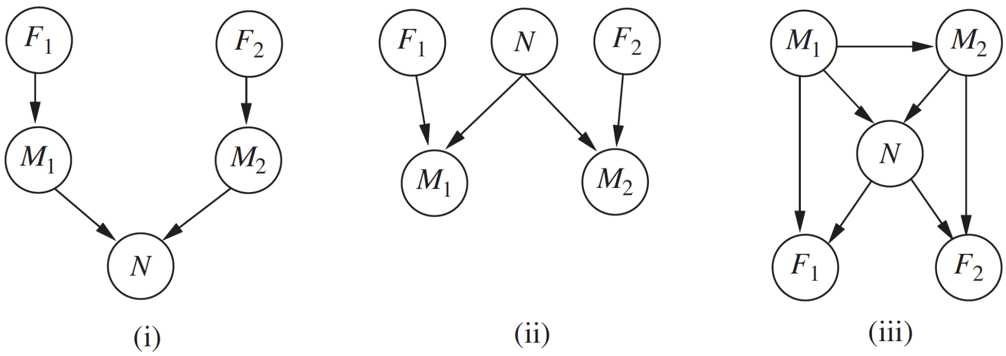

Figure 2: Three possible networks for the telescope problem.

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 **??** 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]

- Write down the Bellman equation(s) for a Markov Decision Process.
- Describe (formally) the Value Iteration algorithm for solving the Bellman equation(s). What are known issues with this algorithm?
- Let us now consider two-player MDPs that correspond to zero-sum turn-taking games with a *finite* maximum number of moves. Let the players be A and B, and let $R(s)$ be the reward for player A in state s. (The reward for B is always equal and opposite.)

  1. Let $V_A(s)$ be the utility of state s when it is A's turn to move in s, and let $V_B(s)$ be the utility of state s when it is B's turn to move in s. All rewards and utilities are calculated from A's point of view (just as in a minimax game tree where rewards are calculated from the MAX player's point of view). Write down Bellman equations defining $V_A(s)$ and $V_B(s)$ when both players optimal.
  2. Explain how to do two-player Value Iteration with these equations, and define a suitable termination criterion.

Question 6 [3 points]

A professor wants to know if students are getting enough sleep. Each day, the professor observes whether the students sleep in class, and whether they have red eyes. The professor has the following domain theory:

- The prior probability of getting enough sleep, with no observations, is 0.7.
- The probability of getting enough sleep on night t is 0.8 given that the student got enough sleep the previous night, and 0.3 if not.
- The probability of having red eyes is 0.2 if the student got enough sleep, and 0.7 if not.
- The probability of sleeping in class is 0.1 if the student got enough sleep, and 0.3 if not.

Answer the following questions:

1. Formulate this information as a dynamic Bayesian network that the professor could use to filter or predict from a sequence of observations. Provide the conditional probability tables.
2. Then reformulate the dynamic Bayesian network as a hidden Markov model that has only a single observation variable. Give the complete probability tables for the model.
3. For the evidence values e_1 = 'not red eyes, not sleeping in class', e_2 = 'red eyes, not sleeping in class' and e_3 = 'red eyes, sleeping in class' compute the following conditional probability distributions:
   (a) $P(\text{EnoughSleep}_t|e_{1:t})$ for $t = 1, 2, 3$.
   (b) $P(\text{EnoughSleep}_t|e_{1:3})$ for $t = 1, 2, 3$.

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

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