INFO8006 Introduction to Artificial Intelligence

Exam of January 2024

Instructions

- The exam lasts for 4 hours.
- You are allowed to use a calculator during the exam, but documents of any kind are forbidden.
- The last two pages can be used for scratch work or for extra space. If you want work done there to be graded, mention where to look **in big letters with a box around them**, on the page with the question.
- Write your last name, first name, and ULiège ID on the first page. Write only your ULiège ID on all the other pages.
- Before handing in your exam, **sort all the pages according to the page numbers** (even if you used additional pages to answer a question).

Good luck!

Last name, First name (in uppercase):

ULiège ID (s201234):

Question 1 [4 points] Multiple choice questions. Choose one of the four choices by filling in its circle. Correct answers are graded $+\frac{4}{10}$, wrong answers are graded $-\frac{2}{15}$ and the absence of answers is graded 0. The total of your grade for Question 1 is bounded below at 0/4.

1. A rational reflex agent ...
   - ◯ maximizes the performance measure.
   - ◯ maximizes the expected performance measure.
   - ◯ maximizes the expected performance measure, given the current percept.
   - ◯ maximizes the expected performance measure, given the current percept and all the previous percepts.
2. The board game Monopoly is ...
   - ◯ static and deterministic.
   - ◯ static and stochastic.
   - ◯ dynamic and deterministic.
   - ◯ dynamic and stochastic.
3. In the H-Minimax algorithm, ...
   - ◯ the evaluation function can be implemented with a neural network.
   - ◯ the heuristic needs to be admissible.
   - ◯ the cutoff test impacts the time complexity of the algorithm, but not its optimality.
   - ◯ it is assumed that the opponent uses the same heuristic as the agent.
4. Let A and B represent the results of two six-sided dice rolls. Let $C = A+B$ and $D = A-B$.
   - ◯ $A \perp B | C$
   - ◯ $A \perp B | D$
   - ◯ $C \perp D | A$
   - ◯ $A \perp B | C, D$
5. In filtering, the belief state is updated at each timestep using ...
   - ◯ the current belief and the transition model.
   - ◯ the current belief, the observation model and the new evidence.
   - ◯ the transition model, the observation model and the new evidence.
   - ◯ the current belief, the transition model, the observation model and the new evidence.
6. A logistic regression model $P(Y = 1|\mathbf{x}) = \sigma(\mathbf{w}^T\mathbf{x} + b)$ is fit by ...
   - ◯ minimizing the zero-one loss between the predictions and the labels.
   - ◯ minimizing the mean squared error between the predictions and the labels.
   - ◯ minimizing the cross-entropy between the predictions and the labels.
   - ◯ minimizing the likelihood of the parameters.
7. In a multi-layer perceptron, removing all the activation functions would reduce the expressiveness of the neural network to ...
   - ◯ a linear model.
   - ◯ a logistic regression model.
   - ◯ a constant function.
   - ◯ It would not reduce its expressiveness.
8. Let $\mathbf{x} = (4, 4, 3, 2, 1, 0, 0)$ and $\mathbf{u} = (-1, 1)$. The convolution $\mathbf{x} \circledast \mathbf{u}$ (as usually defined in convolutional networks) is equal to ...
   - ◯ $(0, 0, 0, 0, 0, 0, 0)$
   - ◯ $(0, -1, -1, -1, -1, 0)$

- ◯ $(0, 1, 1, 1, 1, 1, 0)$
- ◯ $(1, 0, 0, 0, 0, 0, 1)$

9. Which of the following is true? In Markov Decision Processes, ...
   - ◯ the closer the discount factor γ to 0, the higher the utility of future rewards.
   - ◯ the closer the discount factor γ to 0, the longer Value Iteration may take to converge.
   - ◯ the closer the discount factor γ to 1, the greedier the optimal agent.
   - ◯ the closer the discount factor γ to 1, the longer Value Iteration may take to converge.
10. Q-Learning ...
    - ◯ is a model-based reinforcement learning algorithm.
    - ◯ is an on-policy reinforcement learning algorithm.
    - ◯ converges to an optimal policy, but only when acting optimally.
    - ◯ requires a random exploration strategy to converge to an optimal policy.

Question 2 [4 points] The game "21" is played with any number of players who take turns increasing a counter. The counter starts at 1 and each player in turn increases the counter by 1, 2, or 3, but may not exceed 21; the player who says "21" or larger loses.

(a) Define the 2-player version of "21" formally as a game. (Define all the components of the game.)

(b) For the following, consider the game of "5" (still in its 2-player version), which has the same rules has "21" except that you should not say 5 or more.

  (i) Draw the whole game tree of "5".

  (ii) Annotate the nodes of the tree with the backed-up values computed using the Minimax algorithm.

  (iii) Using the Minimax algorithm, choose the optimal starting move.

(c) Explain how the Minimax algorithm should be adapted to deal with the generic case of a game with $n \geq 2$ players. Sketch the pseudo-code to support your answer.

Question 3 [4 points]

(a) Define first-order Markov processes formally. (Define all its components and all the assumptions it makes.)

(b) Derive the recursive update equation of the Bayes filter, assuming continuous state and observation spaces.

(c) Explain whether the Bayes filter is generally tractable for continuous state and observation spaces. If not, explain under which conditions it can be made tractable.

(d) Let us consider a continuous first-order Markov process describing the motion of a robot in a one-dimensional world. On average, the robot moves forward by a distance b at each timestep. The robot's position x_t at time t is not directly observable, but can be estimated using a sensor that measures the distance between the robot and the sensor. Assuming,

- the sensor is fixed at position 0,
- a Gaussian prior $\mathcal{N}(x_0|\mu_0, \sigma_0^2)$ on the initial position of the robot,
- a Gaussian transition model $\mathcal{N}(x_{t+1}|x_t + b, \sigma_x^2)$,
- a Gaussian observation model $\mathcal{N}(e_t|x_t, \sigma_e^2)$,

derive the update equations of the Bayes filter. More specifically, since the belief distribution is Gaussian, derive the update equations of its mean μ_{t+1} and variance σ_{t+1}^2. Gaussian identities from Appendix A can be used if needed.

Question 4 [4 points] Let us consider a neural network f with one hidden layer taking as input a scalar $x \in \mathbb{R}$ and producing as output a scalar $y = f(x;\theta) \in \mathbb{R}$. The neural network is defined as

$$f(x;\theta) = w_1 \text{ReLU}(w_2 x + w_3) + w_4 \text{ReLU}(w_5 x + w_6) + w_7 \text{ReLU}(w_8 x + w_9),$$

where $\theta = (w_1, w_2, w_3, w_4, w_5, w_6, w_7, w_8, w_9)$ is the set of parameters and $\text{ReLU}(x) = \max(x, 0)$ is the rectified linear unit function.

(a) Draw the computation graph representing the neural network and the flow of information from inputs to outputs. Your diagram should be a directed graph that follows the following conventions:

- circled nodes correspond to variables (input, output, parameters or intermediate variables),
- squared nodes correspond to primitive operations (addition, multiplication, ReLU) and produce an intermediate variable as output,
- directed edges correspond to the flow of information, from inputs to outputs.

(b) For $\theta = (-1, \frac{1}{2}, 0, -4, 1, -2, 4, 1, -5)$, draw the function $y = f(x;\theta)$ for $x \in [-10, 10]$.

(c) Using the data point $(x, y) = (10, -15)$ and the value of θ given above, we want to fine-tune the parameter w_8 of the neural network such that $f(x; \theta)$ produces a more accurate prediction of y.

(i) Evaluate the squared error loss at the data point $(x, y) = (10, -15)$ and the current value of θ.

(ii) Derive an expression for the derivative of the squared error loss with respect to w_8.

(iii) Update the parameter w_8 using one step of gradient descent with a learning rate $\gamma = 0.0005$.

(iv) Verify that the value of the loss function has decreased after the update.

Question 5 [4 points] Let us consider the following search problem formulated as a Markov Decision Process.

Pacman is looking for food pellets in a grid world with no ghosts. Pacman can move in the four cardinal directions $\{N, S, E, W\}$, but cannot move through walls. Whenever Pacman moves into a cell containing a food pellet, the food pellet is eaten and a reward of +1 is received. No penalty is received for moving into a cell without a food pellet. The game ends when all the food pellets have been eaten. Throughout this question, assume a discount factor $\gamma = 0.9$.

Let us consider the particular states shown below.

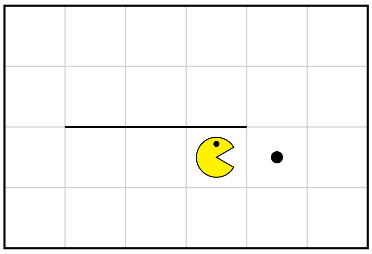

State A.

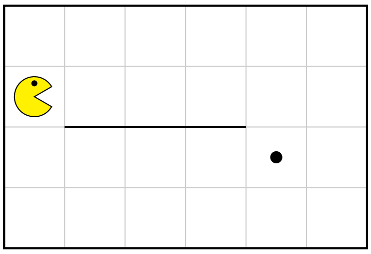

State B.

(a) (i) What is the value $V(A)$ of state A?

(ii) What is the value $V(B)$ of state B?

(b) To find the optimal policy, we decide to use Value Iteration. Assuming $V_0(s) = 0$ for all states s, at which iteration k will $V_k(B)$ first be non-zero? Motivate your answer.

(c) How do the Q-values $Q(A, W)$ and $Q(A, E)$ of moving west and east from state A compare? Is $Q(A, W) > Q(A, E)$, $Q(A, W) < Q(A, E)$ or $Q(A, W) = Q(A, E)$? Motivate your answer.

(d) If we use this MDP formulation of the search problem, is the policy found guaranteed to produce the shortest path from the start state to the goal state? If so, explain why. If not, how could we modify the MDP formulation to guarantee that the policy found produces the shortest path?

Appendix A. Gaussian identities (Särkkä, 2013).

(a) If $\mathbf{x}$ and $\mathbf{y}$ have the joint Gaussian distribution

$$p\begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix} = \mathcal{N}\left(\begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix} \middle| \begin{pmatrix}\mathbf{a}\\ \mathbf{b}\end{pmatrix}, \begin{pmatrix}\mathbf{A} & \mathbf{C}\\ \mathbf{C}^T & \mathbf{B}\end{pmatrix}\right),$$

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

$$\begin{aligned}
p(\mathbf{x}) &= \mathcal{N}(\mathbf{x}|\mathbf{a}, \mathbf{A})\\
p(\mathbf{y}) &= \mathcal{N}(\mathbf{y}|\mathbf{b}, \mathbf{B})\\
p(\mathbf{x}|\mathbf{y}) &= \mathcal{N}(\mathbf{x}|\mathbf{a} + \mathbf{C}\mathbf{B}^{-1}(\mathbf{y} - \mathbf{b}), \mathbf{A} - \mathbf{C}\mathbf{B}^{-1}\mathbf{C}^T)\\
p(\mathbf{y}|\mathbf{x}) &= \mathcal{N}(\mathbf{y}|\mathbf{b} + \mathbf{C}^T\mathbf{A}^{-1}(\mathbf{x} - \mathbf{a}), \mathbf{B} - \mathbf{C}^T\mathbf{A}^{-1}\mathbf{C}).
\end{aligned}$$

(b) If the random variables $\mathbf{x}$ and $\mathbf{y}$ have Gaussian probability distributions

$$\begin{aligned}
p(\mathbf{x}) &= \mathcal{N}(\mathbf{x}|\mathbf{m}, \mathbf{P})\\
p(\mathbf{y}|\mathbf{x}) &= \mathcal{N}(\mathbf{y}|\mathbf{H}\mathbf{x} + \mathbf{u}, \mathbf{R}),
\end{aligned}$$

then the joint distribution of $\mathbf{x}$ and $\mathbf{y}$ is Gaussian with

$$p\begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix} = \mathcal{N}\left(\begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix} \middle| \begin{pmatrix}\mathbf{m}\\ \mathbf{H}\mathbf{m} + \mathbf{u}\end{pmatrix}, \begin{pmatrix}\mathbf{P} & \mathbf{P}\mathbf{H}^T\\ \mathbf{H}\mathbf{P} & \mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R}\end{pmatrix}\right).$$

Extra page 1 / 4.

Extra page 2 / 4.

Extra page 3 / 4.

Extra page 4 / 4.