INFO8006 Introduction to Artificial Intelligence

Exam of August 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 goal-based agent ...
   - ◯ maximizes the expected utility of the outcomes of its actions.
   - ◯ selects actions on the basis of the current percept, ignoring the past and the future.
   - ◯ generates possible sequences of actions and their outcomes and selects actions that achieve its goal.
   - ◯ can run without maintaining a model of the environment.
2. The 2048 game is an example of a ...
   - ◯ fully observable, deterministic game.
   - ◯ partially observable, deterministic game.
   - ◯ fully observable, stochastic game.
   - ◯ partially observable, stochastic game.
3. Which is false? In the tree-search implementation of the A* algorithm, ...
   - ◯ repeated states can hurt the time complexity.
   - ◯ is optimal if the heuristic is admissible.
   - ◯ nodes n are expanded in increasing order of $f(n) = g(n) + h(n)$.
   - ◯ the search terminates when a goal state is added to the fringe.
4. Which is false? The Kolmogorov's axioms state or imply that ...
   - ◯ the probability $P(\omega) \geq 0$ for all sample points $\omega \in \Omega$.
   - ◯ the probability $P(\omega) \leq 1$ for all sample points $\omega \in \Omega$.
   - ◯ the probability $P(\Omega) = 1$.
   - ◯ the probability $P(\{\omega_1, \ldots, \omega_n\}) = \prod_{i=1}^{n} P(\omega_i)$ for any set of mutually exclusive sample points $\{\omega_1, \ldots, \omega_n\} \subseteq \Omega$.
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. Fitting a regression model $P(y|\mathbf{x}) = \mathcal{N}(y; \mu = \mathbf{w}^T\mathbf{x} + b, \sigma^2)$ is equivalent to ...
   - ◯ minimizing the zero-one loss.
   - ◯ minimizing the mean squared error.
   - ◯ minimizing the cross-entropy loss.
   - ◯ minimizing the likelihood of the parameters.
7. In a multi-layer perceptron, ...
   - ◯ the parameters are the weights $\mathbf{W}_k$ and biases $\mathbf{b}_k$ (for $k = 1, \ldots, L$) of each layer.
   - ◯ non-linear activation functions accelerate training by making the model linear.
   - ◯ the output layer is usually set to the hardmax function for classification tasks.
   - ◯ gradient ascent is used to minimize the loss function.

8. Let $\mathbf{x} = (0, 0, 3, 0, 0, -3, 0, -3, 0)$ and $\mathbf{u} = (1, 0, 1)$. The convolution $\mathbf{x} \circledast \mathbf{u}$ (as usually defined in convolutional networks) is equal to ...
   - ◯ $(0, 0, 0, 0, 0, -3, 0)$.
   - ◯ $(1, 0, 1, -1, 0, -2, 0)$.
   - ◯ $(3, 0, 3, -3, 0, -6, 0)$.
   - ◯ $(0, 0, 0, 0, 0, 0, 0)$.
9. In Markov Decision Processes, infinite sequences ...
   - ◯ are not allowed.
   - ◯ can be discounted to bound the expected return.
   - ◯ cannot be avoided.
   - ◯ make the value function $V(s)$ undefined.
10. Temporal-difference learning ...
    - ◯ estimates the value function V^π of a policy.
    - ◯ is a model-based reinforcement learning algorithm.
    - ◯ does not work with stochastic environments.
    - ◯ requires a known transition model $P(s'|s, a)$.

Question 2 [4 points] Let us assume a grid-world in which there are three kinds of food pellets, each with a different color (blue, green and red). Mrs Pacman is only interested in tasting two different kinds of food: the game ends when she has eaten (at least) two pellets of distinct colors (red and blue, green and blue or red and green). Mrs Pacman has eight actions: moving in any of the eight directions (up, down, left, right, and the four diagonals), but she cannot move into any of the B walls. There are L pellets of each kind and the dimensions of the grid-world are $W \times H$.

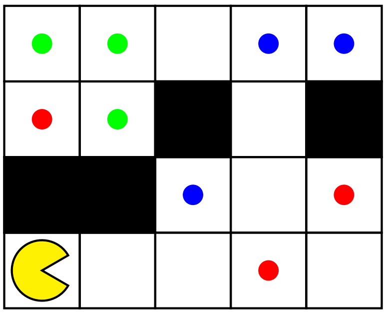

(a) Formally define an *efficient* search problem for Mrs Pacman, for which an optimal solution would minimize the number of moves while achieving the goal. Assume she starts in location (i, j).

(b) Provide a tight upper bound on the size of the state space.

(c) Provide a tight upper bound on the branching factor of the search problem.

(d) For each of the following heuristics, indicate (yes/no) and motivate briefly (1-2 sentences) whether or not it is admissible.

- h_1: the number of remaining red pellets.
- h_2: the Manhattan distance to the closest remaining pellet.
- h_3: the Euclidean distance to the closest remaining pellet.
- h_4: the maximum Euclidean distance between any two remaining pellets of different colors.

(e) Propose an admissible heuristic h that performs better than h_1, h_2, h_3 and h_4 in terms of the number of nodes expanded by the A* algorithm. Explain why it is admissible and why it performs better.

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 positive scalar $\hat{y} = f(x;\theta) \in \mathbb{R}$. The neural network is defined as

$$f(x;\theta) = \text{ReLU}(w_4\text{ReLU}(w_0 x + w_1) + w_5\text{ReLU}(w_2 x + w_3)),$$

where $\theta = (w_0, w_1, w_2, w_3, w_4, w_5, w_6)$ 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 $\theta = (1, -2, -1, -2, 1, 2, 0)$, draw the function $\hat{y} = f(x;\theta)$ for $x \in [-10, 10]$.

(c) Using the data point $(x, y) = (5, 5)$ and the value of θ given above, we want to fine-tune the parameter w_0 such that $f(x; \theta)$ produces a more accurate prediction of y.

(i) Evaluate the squared error loss at the data point $(x, y) = (5, 5)$ and the current value of θ.

(ii) Derive an expression for the derivative of the squared error loss with respect to w_0.

(iii) Update the parameter w_0 using one step of gradient descent with a learning rate $\gamma = 0.01$.

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

Question 5 [4 points] Pacman is in an unknown MDP where there are three states (A, B, C) and two actions (Stop, Go). We are given the following samples collected by taking actions in the environment. For the following, assume a discount factor $\gamma = 1$ and a learning rate $\alpha = 0.5$.

s	a	s'	r
A	Go	B	2
C	Stop	A	0
B	Stop	A	-2
B	Go	C	-6
C	Go	A	2
A	Go	B	-2

(a) Outline the pseudo-code of the Q-learning algorithm.

(b) What are the estimates of the Q-values $Q(C, \text{Stop})$ and $Q(C, \text{Go})$ as obtained by running Q-learning updates on the samples above (in the same order)? All Q-values are initialized to 0.

(c) We now consider a feature-based representation

$$Q(s,a) = w_1 f_1(s,a) + w_2 f_2(s,a),$$

using two features

$$f_1(s,a) = 1,$$
$$f_2(s,a) = \begin{cases} 1 & \text{if } a = \text{Go}, \\ -1 & \text{otherwise.} \end{cases}$$

Starting from initial weights of $w_1 = 0$ and $w_2 = 0$,

(i) Compute the updated weights w_1 and w_2 after the first sample $(A, \text{Go}, B, 4)$.

(ii) Compute the updated weights w_1 and w_2 after the second sample $(B, \text{Stop}, A, 0)$.

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 / 2.

Extra page 2 / 2.