

INFO8006 Introduction to Artificial Intelligence

Exam of August 2020

Instructions

- The exam starts at 8:30 AM.
- Question 1 must be answered on eCampus. The deadline is set to 9:30 AM.
- Questions 2 to 5 must be answered on paper. Your sheets must be scanned and submitted on eCampus before 12:30 AM.
- 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.

Question 1 [4 points]


Multiple choice questions. Choose one of the four choices. 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. Let us consider the task environment for the card game of Belotte. Which of the following is true?
 - (a) Belotte is deterministic, partially observable, continuous and single agent.
 - (b) Belotte is stochastic, partially observable, discrete and multiagent.
 - (c) Belotte is stochastic, fully observable, discrete and multiagent.
 - (d) Belotte is stochastic, partially observable, dynamic and multi-agents.
2. Let h_A and h_B be two admissible heuristics, which of the following propositions is correct:
 - (a) $\min(h_A, h_B)$ dominates h_A .
 - (b) $h_A + h_B$ is admissible.
 - (c) $\min(\max(h_A, h_B), \min(h_A, h_B))$ is admissible.
 - (d) $\max(h_A + h_B, h_A - h_B)$ is admissible.
3. Let α and β denote two sentences from propositional logic. Which of the following propositions is not equivalent to the others:
 - (a) $\alpha \models \beta$.
 - (b) $M(\alpha) \subset M(\beta)$.
 - (c) $(\alpha \wedge \neg\beta)$ is unsatisfiable.
 - (d) $M(\alpha \wedge \neg\beta) = \{\}$.
4. Which of the following propositions is correct:
 - (a) A transposition table can only be used for deterministic 2-player games.
 - (b) A dictionary is an efficient way of implementing transposition tables.
 - (c) Minimax is optimal if the adversary is not optimal.
 - (d) In MCTS, the backpropagation phase optimizes the parameters of the algorithm.
5. Which of the following propositions is not equivalent to the others:
 - (a) $\forall x, y, P(x, y) = P(x)P(y)$.
 - (b) $\forall x, y, z, P(x)P(z|x, y)P(y) = P(x, y, z)$.
 - (c) $X \perp Y$.
 - (d) $X = f(Y)$ (read the random variable X is a function of the random variable Y).
6. You are tasked to implement a procedure for sampling from a 100-dimensional multivariate Normal distribution. Which of the following would be the most appropriate algorithm?
 - (a) Inference by enumeration.

- (b) Inference by variable elimination.
 - (c) The Kalman Filter.
 - (d) Gibbs sampling.
7. Let us consider the RAGI robot wandering around at the Montefiore Institute. Which of the following is wrong?
- (a) Filtering can make use of its sensor data to maintain a belief state about its position.
 - (b) A neural network can be used for detecting obstacles.
 - (c) The H-Minimax algorithm can be helpful for decoding the speech of its visitors.
 - (d) Policy iteration can be used to make rational decisions about its next move.
8. Let us consider the following version of the Bellman equation: $V(s) = R(s) + \gamma \max_a \sum_{s'} P(s'|s, a)V(s')$. Which of the following is false?
- (a) The equation holds even if the environment is continuous.
 - (b) The equation holds even if the environment is deterministic.
 - (c) The equation holds even if the environment is partially observable.
 - (d) The equation holds even if the environment is fully observable.
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. The search algorithm in AlphaGo is based on...
- (a) the Minimax algorithm.
 - (b) Monte Carlo tree search.
 - (c) A*.
 - (d) a Markov Chain Monte Carlo method.

Question 2 [4 points]

1. Yann, Yoshua and Geoffrey are playing poker together. The objective of the game is to become the wealthiest player and so to avoid being liable. Fill in the following multiplayer minimax tree, in which the players utilities is given in dollars and as (Yann's utility, Yoshua's utility, Geoffrey's utility). Yann is the first player, Yoshua is second and Geoffrey the third.





2. As you may have noticed, poker is a zero-sum game. Below, apply an algorithm similar to $\alpha - \beta$ pruning. Cross off any branches that can be safely pruned. If no branches can be removed, explain why not. Make the assumption that the nodes are explored from left to right.
3. In addition to the zero-sum game assumption, we now further assume that all utilities are bounded between $[-10, 10]$. Below, cross off any branches that can be safely unexplored. If no branches can be removed, explain why not. Make the assumption that the nodes are explored from left to right.
4. Again, consider the zero-sum, bounded-utility, 3-player game described above. Here, we assume the utilities are bounded by $[-C, C]$ for some constant C (the previous problem considered the case $C = 10$). Assume there are only 3 actions in the game, and the order in which players play is Yann, Yoshua, then Geoffrey. Below is pseudo-code that computes the utilities at each node, but the pruning conditions are left out. Fill in the conditions below that prune wherever safe.

Algorithm 1 Compute palyers' values.

```
1: procedure COMPUTE-GEOFFREY-VALUE(state,  $C$ ,  $\beta$ )
2:    $v_G = -\infty$ 
3:    $\gamma = -\infty$ 
4:   for each successor of state do
5:      $(w_{Y_a}, w_{Y_o}, w_G) = \text{EVALUATE-UTILITY}(\text{successor})$ 
6:     if  $w_G > v_G$  then
7:        $v_{Y_a} = W_{Y_a}$ 
8:        $v_{Y_o} = W_{Y_o}$ 
9:        $v_G = W_G$ 
10:    end if
11:    if ???? then
12:      return  $(v_{Y_a}, v_{Y_o}, v_G)$ 
13:    end if
14:     $\gamma = \max(\gamma, v_G)$ 
15:  end for
16:  return  $(v_{Y_a}, v_{Y_o}, v_G)$ 
17: end procedure

1: procedure COMPUTE-YOSHUA-VALUE(state,  $C$ ,  $\alpha$ )
2:    $v_{Y_o} = -\infty$ 
3:    $\beta = -\infty$ 
4:   for each successor of state do
5:      $(w_{Y_a}, w_{Y_o}, w_G) = \text{COMPUTE-GEOFFREY-VALUE}(\text{successor}, C, \beta)$ 
6:     if  $w_{Y_o} > v_{Y_o}$  then
7:        $v_{Y_a} = W_{Y_a}$ 
8:        $v_{Y_o} = W_{Y_o}$ 
9:        $v_G = W_G$ 
10:    end if
11:    if ???? then
12:      return  $(v_{Y_a}, v_{Y_o}, v_G)$ 
13:    end if
14:     $\beta = \max(\beta, v_{Y_o})$ 
15:  end for
16:  return  $(v_{Y_a}, v_{Y_o}, v_G)$ 
17: end procedure

1: procedure COMPUTE-YANN-VALUE(state,  $C$ )
2:    $v_{Y_a} = -\infty$ 
3:    $\alpha = -\infty$ 
4:   for each successor of state do
5:      $(w_{Y_a}, w_{Y_o}, w_G) = \text{COMPUTE-YOSHUA-VALUE}(\text{successor}, C, \alpha)$ 
6:     if  $w_{Y_a} > v_{Y_a}$  then
7:        $v_{Y_a} = W_{Y_a}$ 
8:        $v_{Y_o} = W_{Y_o}$ 
9:        $v_G = W_G$ 
10:    end if
11:    if ???? then
12:      return  $(v_{Y_a}, v_{Y_o}, v_G)$ 
13:    end if
14:     $\alpha = \max(\alpha, v_{Y_a})$ 
15:  end for
16:  return  $(v_{Y_a}, v_{Y_o}, v_G)$ 
17: end procedure
```



Question 3 [4 points]

A small robot going under the name of WALL-E wanders around the University to find and collect waste on the campus. When it detects an interesting object on the ground, WALL-E computes and executes a sequence of instructions specific to the type of the object it has recognized. Unfortunately, random bugs introduced by the students who wrote its program may appear, which results in the sequence of instructions to be improperly executed and therefore in more frequent failures when trying to grasp an object. Consider the Boolean variables S (the grasp succeeds) and B (it has randomly reached a state in which its program causes the gripper to be buggy), as well as the multivalued variables P (the sequence of pick up instructions), I (the image recorded by its camera) and O (the actual object to pick up).

- (a) Given the following factorization of the joint distribution,

$$P(B, I, O, P, S) = P(B)P(I|O)P(O)P(P|I)P(S|B, P, O),$$

draw the most sparse Bayesian network that correctly captures the independence assumptions.

- (b) From the topology of your network, determine which of the following is true, false, or cannot be determined.

- i. $I \perp B$
- ii. $O \perp S$
- iii. $P \perp S|I$
- iv. $O \perp S|\{B, P\}$
- v. $I \perp B|S$
- vi. $S \perp I|B$

- (c) Assume there are only two types of objects, bottles and boxes, each with their own and appropriate pickup sequence. Under normal conditions, if the appropriate pickup sequence is executed, then the probabilities of success are x_1 and x_2 respectively for bottles and boxes, and y_1 and y_2 when WALL-E attempts to execute the wrong pickup sequences. When WALL-E is in a state that causes bugs, all probabilities of success are halved. Give the conditional probability table associated with S .
- (d) Assume the expected accuracy of the convolutional network used to recognize an object from images of the camera is 90% for bottles and 80% for boxes. Explain how and why you can make use of this assumption to simplify your Bayesian network.
- (e) Using this simplified Bayesian network, calculate an expression for the probability that the object is actually a box, given a successful grasp and a recognition of the object as a box.

Question 4 [4 points]

You decide to start your own game company after graduating. Your first game starts with two ghosts, Greeny and Blinky. On each turn, the player clicks on a ghost. Clicking on any Greeny (action a_G) produces a new Greeny with probability $(2g + b)/(2(g + b))$ and a new Blinky otherwise, whereas clicking on any Blinky (action a_B) produces a new Blinky with probability $(g + 2b)/(2(g + b))$ and a new Greeny otherwise, where g and b are the numbers of Greenies and Blinkies in the current state. There is a reward +1 for producing an offspring of the same kind. The game ends after T steps.

- (a) Sometimes MDPs are formulated with a reward function $R(s, a, s')$ defined as the reward obtained by performing the action a from state s and resulting in state s' . Write the Bellman equations for this formulation.
- (b) Assuming the formalism above, define the game as an MDP, with a minimal state space.
- (c) Count the total number of states in the MDP.
- (d) For $\gamma = 0.9$, $T = 3$ and $V_0(s) = 0$ for all s , run the Value Iteration algorithm for 3 iterations.
- (e) Explain why Value Iteration applied to this problem converges in $O(T)$ iterations.

Question 5 [4 points]

You observe a Grandmaster agent playing Pacman. How can you use the moves you observe to train your own agent?

- (a) Describe formally the data you would collect, the inference problem you would consider, and how you would solve it.
- (b) How would you design a neural network to control your agent? Define mathematically the neural network architecture, its inputs, its outputs, its parameters, as well as the loss you would use to train it.
- (c) Discuss the expected performance of the resulting agent when (i) the Grandmaster agent is optimal, and (ii) the Grandmaster agent is suboptimal.