國立陽明交通大學試題紙

112 學年度第一學期
博士班資格考

科目：人工智慧 A　　日期：113 年 1 月 9 日　第 1 頁 共 2 頁

請 "✓" 明　✓不可看書　可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. (15 pts total) Consider the following relation ρ on variables A, B, C, and D, ρ(A,B,C,D)={(a,a,a,a), (a,b,b,b), (b,b,a,c)}. A search space is backtrack-free for a particular order of variable instantiation if all the paths in the search space are full lengths, i.e. all variables can be successfully instantiated.
   (a) (5 pts) Show *proj*(ρ).
   (b) (5 pts) Give a variable ordering such that *proj*(ρ) is backtrack-free, and show the search tree.
   (c) (5 pts) Give a variable ordering such that *proj*(ρ) is not backtrack-free, and show the search tree.

2. (15 pts total) Consider the following 2-player game, X vs. Y. At each stage, an integer utility value N is assigned to a legal move. Each of X and Y has two legal moves, respectively, as shown below.

Player X:
Move A: new N= old N + (old N mod 23) -11
Move B: new N= old N + 20 – (old N mod 41)

Player Y:
Move C: new N= old N + 2*(old N mod 17) -16
Move D: new N= old N + 25 – (old N mod 51)

At the initial state, N = 100, and player X makes the 1st move, followed by player Y. This game is played for 4 ply, i.e. X plays first, Y plays second, then X plays again, and at last Y plays. Then the game is over. X will try to maximize N. On the contrary, Y will try to minimize N.
Now assume the game tree is evaluated using Minimax that follows the alphabetical order, i.e. Move A before Move B, also Move C before Move D, and using alpha-beta to prune the tree from left to right.

(a) (12 pts) Draw the game tree with alpha-beta pruning, and mark the pruned nodes (or branches), showing the utility value of each node according to the definitions above and all the utility values backpropagated from the bottom level.
(b) (3 pts) Which move will X make?

A partial sample game tree with the top two levels is shown below for your reference. The numbers in the boxes are the values of N. You need to finish the rest of the game tree, and backpropagate the appropriate N values. Show the backpropagated N values next to the boxes, e.g. N_0, N_1, N_2.

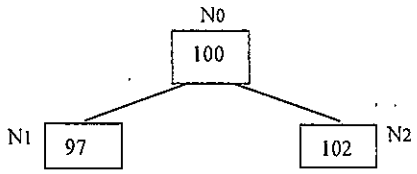

3. (20 pts total) Let a, b, c, and d be four propositions.
   A knowledge base, KB, contains: $a \rightarrow c$, $b \rightarrow (c \vee d)$, $a \vee b$.
   (a) (5pts) Does KB entail $(c \vee d)$? (Yes/No) Explain why and how by a truth table
   (b) (5 pts) Can modus ponens derive (i.e. prove) $(c \vee d)$? (Yes/No) Explain why and how.
   (c) (4 pts) Given your answers to (a) and (b), is modus ponens sound for this KB? (Yes/No), and is modus ponens complete for this KB? (Yes/No)
   (d) (6 pts) Can resolution derive (i.e. prove) $(c \vee d)$? (Yes/No) Explain why and how.

國立陽明交通大學試題紙

112 學年度第一學期
博士班資格考

科目：人工智慧 B　　日期：113 年 1 月 9 日　第 1 頁 共 1 頁

請 "✓" 明　✓不可看書　可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

Answer the following questions concisely (5 points each)

1. What is the difference between feedforward and recurrent artificial neural networks?
2. Consider the regression problem of fitting a function, $y = \sum_{k=1}^{p} a_k f_k(x)$, to a set of points, $\{(x_i, y_i) \mid 1 \le i \le n\}$. There's no adjustable parameter in the functions $f_k(x)$. Derive the min-squared-error solution of the coefficients a_k.
3. Following question #2, Derive the update equation of the coefficients if we want to obtain the coefficients via gradient descent. Assume that only one sample, (x_i, y_i), is used per step.
4. L2 regularization uses the sum of squared coefficients as the regularization term. Add L2 regularization to the cost function in question #2 and re-derive the solution.
5. We want to do k-means clustering to a set of three 1-D samples: 1, 3, and 6. Let k=2. Start with prototypes at 1 and 6, give the steps of k-means until convergence.
6. What are the meanings of "bias" and "variance in the bias-variance dilemma regarding supervised learning? You can use a regression problem, such as the one in question #2, as an example.
7. Explain what a policy is in reinforcement learning.
8. The following is the standard update equation in Q-learning. Explain the meanings of as many symbols in this equation as possible.

$$Q(s,a) \leftarrow Q(s,a) + \alpha\left[r + \gamma \max_{a'} Q(s',a') - Q(s,a)\right]$$

9. Sketch a logistic function, and explain how it is used in binary classification (i.e., logistic regression).
10. For decision-tree learning, describe a criterion used for selecting the attribute to split a tree node.