國立交通大學試題紙

一百零九學年度第一學期
博士班資格考

科目：人工智慧 A　　　　日期： 110 年1月 19 日　第 1 頁 共 3 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. (12 pts total) Consider the graph coloring problem below.

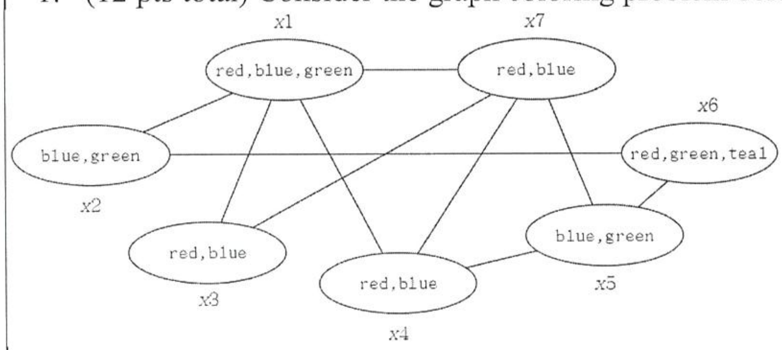

Assume we start with variable x_7, and instantiate it by color red. After performing Forward Checking, show the domains of the remaining variables.

D_1=__________________ D_2=__________________ D_3=__________________

D_4=__________________ D_5=__________________ D_6=__________________

2. (12 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.
   X player:
   Move A: new N= old N + (old N mod 23) -11
   Move B: new N= old N + 20 – (old N mod 41)
   Y player:
   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 X player makes the 1st move, followed by Y player. 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.
Now assume the game tree is evaluated using DFS 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) (6 pts) Draw the game tree with alpha-beta pruning, and mark the pruned nodes (or branches).
(b) (3 pts) Show the utility value of each node according to the definitions above,
(c) (3 pts) Show all the utility values backpropagated from the bottom level,

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

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.

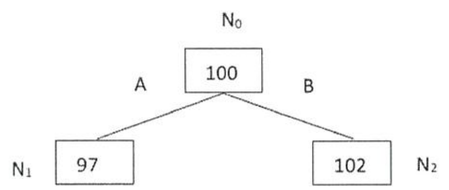

3. (17 pts total) You are asked to solve a murder case by applying resolution. Below is what we know about this case after some investigation.
According to biologists, whoever was born in City AA is a blonde. Meanwhile, psychologists suggest that if X and Y are friends, then X like Y and Y likes X. In addition, if X and Y are friends, and Y and Z are friends, then X and Z are friends. Furthermore, criminologists conclude that whoever was born in City BB, and has low income is desperate. Also, anyone that is desperate, and likes a blonde will murder that blonde. The body of Mary was found in her bedroom. She was brutally murdered.
A detective has done some digging into Mary's social background, and found the following. Helen is a friend of Mary's, and John is Helen's friend. David likes Mary. Both Mary and David were born in City AA. John and Helen were both born in City BB. John and David have low income, but Helen and Mary have high income.
Disclaimer: This is a make-up murder case for an exam. No offense to blondes.

Now apply resolution to find out who murdered Mary.
You can only use the predicates below.
Born_AA(X): X was born in City AA; **Born_BB(X)**: X was born in City BB;
Like(X,Y): X likes Y; **Friend(X,Y)**: X and Y are friends; **Blonde(X)**: X is a blonde;
Desperate(X): X is desperate; **Lowin(X)**: X has low income; **Hiin(X)**: X has high income;
Murder(X,Y): X murders Y
(a) (6 pts) First translate the necessary info above into first-order logic (FFL) statements.
(b) (5 pts) Convert FFL into clausal forms.
(c) (6 pts) Apply resolution to identify the one that is responsible for the murder of Mary.

4. (9 pts total) Given the problem space below, where S is the starting state, G is the goal state, the number next to a node is an estimated cost (i.e. *h*) to goal, and the number by a link is real cost (i.e. *g*), for a **greedy stochastic** search (i.e. **no** backtrack is allowed) strategy, SA-like (simulated annealing-like), we define the probability of being the next state out of the current child nodes as follows.
Let *s* denote a state, and *C* as the set of its child nodes, i.e. $C=\{c_i \mid \text{any } c_i \text{ is a child node of } s\}$.

$$Prob(c_i\ being\ next\ state) = \frac{e^{-h_{c_i}}}{\sum_{c_j \in C} e^{-h_{c_j}}},$$ where hc_i is the estimated cost of c_i.

Note. Let $e = 3.0$.

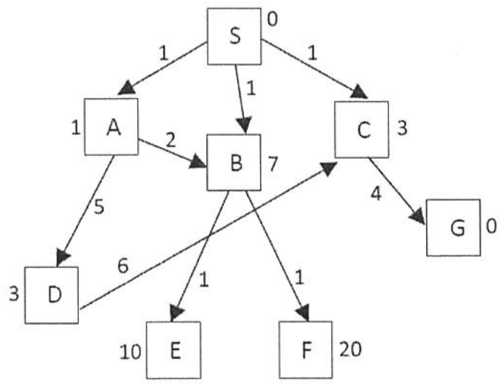

(a) (3 pts) Show the solution path with the minimum cost that can be found by SA-like, and calculate its probability.
(b) (3 pts) What is the most probable solution path found by SA-like, and what is the probability?
(c) (3 pts) Show the search path that gets stuck with the maximum probability, and calculate its probability.

國立交通大學試題紙

一百零九學年度第一學期
博士班資格考

科目：人工智慧 B　　　　日期： 110 年 1 月 19 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.

1. [18%] Consider the problem of estimating the probability distribution of a scalar random variable x given a set of samples: $x_1, x_2, \ldots, x_n$.

(a) [4%] Let the desired probability distribution be $p(x; \alpha)$, where α is a parameter of the functional form of p. Describe the procedure of using the maximum-likelihood method to estimate α.

(b) [4%] Let the desired probability distribution be a 1-D normal distribution with mean μ and standard deviation σ. Derive the maximum-likelihood estimation of μ. Note: Use log-likelihood.

(c) [2%] A more complex probability distribution can be estimated as a mixture of parameterized distributions. Write down the 1-D probability distribution as a mixture of Gaussians.

(d) [2%] Explain the concept of "hidden" or "latent" variables in this problem.

(e) [6%] The standard algorithm to estimate the parameters in a mixture model is the EM algorithm. What do the terms E and M stand for? You have to describe their meanings. Show the steps of the EM algorithm with pseudo-code; you don't need to give the actual equations.

2. [16%] Assume a system with four states S_1, S_2, S_3, and S_4 with rewards of R_1, R_2, R_3, and R_4, respectively. There are three possible actions a_1, a_2, and a_3 from each state. Use the system to answer the following question about reinforcement learning:

(a) [3%] What is a policy?

(b) [4%] What is a Q-function (in Q learning), and how is it related to the policy?

(c) [5%] Assume that the episode below is executed:

$S_1 \rightarrow$ (action a_2) $\rightarrow S_4 \rightarrow$ (action a_1) $\rightarrow S_3$

Which Q values are updated after this episode? What are their new values? You can assume the original Q values are all zero. Use α and γ to represent the learning rate and discount factor, respectively.

(d) [2%] What is the effect of the discount factor in general?

3. [16%] Answer the following questions about supervised learning:

(a) [4%] What is the difference between parametric and non-parametric models? Also give an example each.

(b) [5%] What is the purpose of k-fold cross-validation? How is it done?

(c) [3%] Explain the concept of boosting when learning an ensemble based classifier.

(d) [4%] How is a naïve Bayesian model different from a regular one? What is the motivation (or advantage) of using a naïve model?

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：