STUDENT NAME: ______________________________

STUDENT ID: ______________________________

McGill University
Faculty of Science
School of Computer Science
Final exam

COMP-424A
Introduction to Artificial Intelligence

April 19, 2010
9:00-12:00

Examiner: Prof. Doina Precup

Associate Examiner: Prof. Prakash Panangaden

This examination is closed-book, closed-notes.

There are 16 pages, including the title page.

Start by writing down your name. Answer the questions directly on the exam booklet. Additional pages are provided if necessary. Do not forget to write your name on the additional pages too.

Read the whole exam before starting to work on it. There are 10 questions, all of which require written answers. Values for each question are shown in brackets. Partial credit will be given for incomplete or partially correct answers.

Good luck!

1. [10 points] **Search algorithms**

   (a) [5 points] Suppose you have an admissible heuristic h. For what numbers a is $a * h$ admissible? If you do not know anything about h, are there any values of a for which using $a * h$ in an A^* algorithm would be better than using h? Justify your answers

(b) [2 points] Suppose that you are implementing a game player and using alpha-beta search to go to some depth k. You bought a new machine which allows you to go up to at most $k + 1$. Are you guaranteed to get a better player with the new machine, or not? Explain your answer.

(c) [3 points] You have a problem with a very high branching factor, so you cannot run A^* on it. Instead, you decide to use a *beam search* algorithm. In this case, for every node, you will insert in the queue only the k best successors, according to $f = g + h$. Here, k is a fixed parameter that you set based on your memory requirements (e.g., $k = 2$). If the heuristic is admissible, is this algorithm guaranteed to find the optimal solution or not? Justify your answer.

2. [6 points] **Constraint satisfaction**

   Consider the game of mine sweeper, in which you have a square board. In each square you can have either a number n between 0 and 8, meaning that exactly n of the 8 neighbors of the square contain a bomb, or the square could be unknown. Suppose that you have a *given* board configuration, and you are trying to find a safe square to play. Explain how you would model this as a constraint satisfaction problem. What would be the variables, the domains of the variables and the constraints?

3. [15 points] **Logic**

   (a) [10 points] Translate the following sentences in first-order logic. Make sure that you use the predicates that you define consistently among the different statements.

      i. [2 points] All robots are smart.

      ii. [2 points] Some robots are smart.

      iii. [1 point] Robbie is a robot.

      iv. [2 point] All robots are nice to their owner

      v. [1 point] Bob is Robbie's owner

      vi. [2 points] Some people who own robots are smart.

(b) [3 points]
Prove formally that Robbie is nice to Bob

(c) [2 point] Given this knowledge base, can you prove that Bob is smart? If you answer is yes, show the proof. If your answer is no, explain why not.

4. [20 points] **Neural networks and decision trees** Suppose that you have been hired by a large on-line company to build a predictor of whether customers will buy their products, based on customer attributes. They already have 5 predictors that they built and would like you to make use of them in a smart way, in order to get better accuracy than these predictors.

   (a) [5 points] The first thing that comes to mind is that you want to take a majority vote among the existing predictors. Can you implement this function using a perceptron? If so, show the perceptron. If no, explain why not.

   (b) [5 points] Suppose that instead of 5 predictors they have some arbitrary number n. Explain what changes you need to make to your previous answer to obtain a majority vote in this case.

(c) [5 points] Suppose that you wanted a decision tree that represents the majority vote over 5 predictors. Would this be a good choice? Explain why or why not.

(d) [5 points] Suppose that one of your colleagues has puled a prank you and made 3 of the 5 predictors identical. Would the majority vote provide better accuracy than the single predictors in this case? Justify your answer.

5. [5 points] **Gradient descent for learning**

   Suppose that you want to train a hypothesis of the form:

   $$h(x) = w_0 + w_1 \cos(w_2 x)$$

   Give update rules for w_0, w_1 and w_2, assuming you do incremental gradient descent using the squared error as an error function. Explain if in this case local minima are a problem or not.

6. [5 points] **Overfitting**

   Suppose that we have a data set in which there are n binary attributes and the desired output is also binary. The examples that you see come from a random function and are perfect (the training data has the correct label). You have a "learning algorithm" that simply memorizes the training data. When asked for the output for a new instance, the algorithm will look up the instance and if it is in memory, it answers with the recorded label. Otherwise, it will answer randomly.

   (a) [1 point] What is the training error of this algorithm?

   (b) [2 points] Suppose the algorithm has seen k distinct training examples. It is tested on a new example drawn randomly from the *entire* space of possible instances. What is the expected error on this test instance? What happens to the error as k increases?

   (c) [1 points] Suppose that after seeing k training examples, the algorithm is tested on an instance which has not been part of the training data. What is the testing error as a function of k?

   (d) [1 point] Explain if the cross-validation algorithm would evaluate the error of this learner correctly or not.

7. [4 points] **Probabilities**

   Suppose you have two independent binary random variables A and B, and you know that $P(A = 1) = 0.1$ and $P(B = 1) = 0.5$.

   (a) [3 points] What is the probability that *at least* one of the variables would be 1?

   (b) [3 points] What is the probability that *exactly* one of the variables would be 1?

8. [15 points] **Bayes nets**

   Students in the AI class like to relax after their final exams by watching movies. How much they like a movie depends on how happy they are and what kind of movie it is. If they are happy, they like any movie with probability 0.8. If they are not happy, they still like sci-fi movies with probability 0.8, but they like non-sci-fi movies only with probability 0.5. They have a 0.6 chance to be happy on any given night. The movie theatre has a 0.5 chance of showing a sci-fi movie on any night.

   (a) [5 points] Draw a Bayes net describing the problem statement above. Specify both the graph structure and all the parameters.

   (b) [3 points] Compute the probability that the students will like the movie they see on any given night. Please show the formula you are using, then substitute the adequate parameters and show the resulting number.

(c) [2 points] Suppose that the students go to the movies for 10 nights. What is the expected number of nights on which they will like the movie?

(d) [5 points] The cinema manager notices that on nights on which he plays sci-fi movies, the students have a 0.7 chance of buying popcorn, while on nights on which he plays non-sci-fi movies, they have a 0.5 chance of buying popcorn. How does this change the structure of the Bayes net and parameters described above? Based on this information, and considering that the movie theatre also has non-student customers, can you determine with what probability the manager should show sci-fi movies? Explain your answer.

9. [15 points] **Markov Decision Processes and reinforcement learning**

   You have a house with two rooms (left and right) and a robot can move between the rooms. The right room contains a battery charger, so when the robot is in this room, it gets a reward of +1. In the other room, the reward is 0. The robot can choose to stay in the same room, or move to the other room. The stay action always succeeds but the move action fails 1/4 of the time.

   (a) [5 points] Describe the problem as a Markov Decision Process. What are the states, actions, transition probabilities and rewards?

   (b) [1 point] What is the optimal policy?

   (c) [4 points] Suppose that the discount factor is $\gamma = 1/2$. Compute the optimal value function by solving the linear system of equations corresponding to the optimal policy. Show all your work.

(d) [3 points]

Suppose that your robot has performed a sequence of actions in which it stayed 3 times in the left room, it moved (successfully) to the right room, it stayed in the right room once, then it moved successfully to the left room. If it has been doing TD-learning with a learning rate of 0.1, what are its estimates for the values of the two rooms at this point?

(e) [2 points] Suppose that you start with all values at -1 and use a policy that never explores. Will you be able to find the optimal policy? Justify your answer.

10. [5 points] **Problem formulation**

    You got a summer job at a sandwich shop downtown. The shop sells different kinds of sandwiches, each consisting of a type of bread, a type of meat (or substitute), a type of cheese (if the customer wants some) and a type of vegetable. Each day, you have to go to the grocery store and buy supplies for the day's sandwiches. Bread has to be discarded at the end of the day. Meat and cheese have to be discarded after a week. Vegetables are discarded after 3 days. You know that customers come in randomly. If a customer cannot have the desired sandwich, she will leave and there is a 0.5 probability that she will never come to the shop again. Of course, customers pay different amounts for different sandwiches, and you have to pay for the groceries. Having taken AI, you want to optimize, in a principled way, the amount of purchases you are making. Describe this problem using any AI technique of your choice. Explain *all* the components of your model. What algorithms could you use to solve this problem?