

B551 Final Practice Problems: Fall 2011

The final exam will be similar in breadth, length, and difficulty to these problems. All topics covered in class will be fair game for the final exam.

I. Traveling on a roadmap

You are given a roadmap of some country in the form of a connected non-directed graph in which nodes represent cities and edges represent roads between cities. (A connected graph is one in which every two nodes are connected by a path made of one or several consecutive edges.) Each edge (i,j) is labeled by the length $l(i,j)$ of the road between cities i and j .

Two friends live in two different cities, a and b , of the map. They want to meet in a city of the map (any one). To do this, they move in successive turns. On every turn, the two friends start moving at the same time. Each friend moves to a neighboring city on the map; he/she cannot stay in the same city. The amount of time needed to move from city i to neighboring city j is equal to the length $l(i,j)$ of the road between cities i and j . So, the two friends may not reach their respective new cities at the same time. The friend that arrives first to his/her new city must wait until the other arrives to his/her new city (each one calls the other on his/her cell phone when he/she arrives to a new city) before the next turn can begin. The two friends want to meet as quickly as possible. Note that the goal for the two friends is to meet in a city, not anywhere on a road.

1. Formulate this problem as a state-space search problem:
 - a) What is the state space?
 - b) What are the initial and the goal states?
 - c) What is the successor function?
 - d) What is the step cost function?
2. Let $D(i,j)$ be the straight-line distance between any two cities i and j in the map. Which, if any, of the following heuristic functions are admissible? Why?
 - a) $D(i,j)$
 - b) $D(i,j)-2$
 - c) $D(i,j)/2$
3. Is the following statement true or false: "There are connected maps for which no solution exists"? If you answer 'true', give an example of such a map. If you answer is 'false', prove it.

II. Approximately optimal search

The two objectives of finding a solution as quickly as possible and finding an optimal solution are often conflicting. In some problems, one may design two heuristic functions h_A and h_N , such that h_A is admissible and h_N is not admissible, with h_N resulting in much faster search most of the time. Then, one may try to take advantage of both functions.

1. A best-first search algorithm called A_ϵ^* uses the evaluation function $f(N) = g(N) + h_A(N)$. At each iteration, A_ϵ^* expands a node N' such that $f(N') \leq (1+\epsilon) \times \min_{N \in \text{FRINGE}} \{f(N)\}$, where ϵ is any strictly positive number. What can you say about the cost of the solution returned by A_ϵ^* ?
2. Explain briefly how A_ϵ^* can use the second heuristic function h_N to reduce the time of the search. What tradeoff is being made in choosing ϵ ?


III. Modified Tic-Tac-Toe

Consider the game of 2×2 tic-tac-toe where each player has the additional option of passing, i.e., of marking no square on the 2×2 board. The two players, MAX and MIN, take turns, with MAX going first.

1. Draw the full game tree down to depth 2 (recall that the root of the tree is at depth 0). Do not show the nodes that are rotations or reflections of siblings already shown (your tree should have five leaves). [Draw your tree nicely since you will have to use it again for your answers to questions 2 and 3.]
2. Let the evaluation function of MAX be the number of MAX's marks on the board minus the number of MIN's marks. Give the values of the evaluation function for all leaves of the tree constructed in Question 1 and the values backed-up by the Minimax algorithm for all internal nodes. [Show these values on the tree drawn in Question 1.]
3. Circle all nodes [in the tree that you drew in Question 1] that would not be evaluated by the Alpha-Beta algorithm during a left-to-right depth-first exploration of your tree.
4. Suppose we wanted to solve the game to find the optimal move of MAX (i.e., by constructing a game tree with no depth limit). Explain why Alpha-Beta pruning with an appropriate node ordering can do it, while Minimax can't.

IV. Automotive Diagnosis

Consider the following simple network for car diagnosis:




Each variable is Boolean, and a value of True indicates that the aspect of the car is working properly.

1. How many independent probability values would be listed in the joint probability table for these six variables, if no independence assumptions were made?
2. How many independent probability values are listed in the conditional probability tables of this BN?
3. Given what other variables can you say that Battery is independent of Moves? Given what other variables can you say that Battery is NOT independent of Moves?
4. Given what other variables can you say that Gas is independent of Radio? Given what other variables can you say that Gas is NOT independent of Radio?
5. List at least two algorithms that can perform inference on this Bayesian network.

V. Support Vector Machines

1. Construct by hand a 2-dimensional linear classifier that is consistent with the positive examples (0,3), (1,4), (2,5) (2,3) and the negative examples (3,2), (1,1),(2,2), (4,3). Give the equation defining this classifier.
2. For the above example, consider the linear classifier $-2*x + y + 2$. What is the geometric margin of each of the data points? (If an example is misclassified, then its margin is negative)
3. Consider the following 2D dataset (filled circles indicate positive examples, empty circles indicate negative ones).



How might you transform the data into a higher-dimensional feature space in order to get a good linear classifier in that space? Hint: consider that the positive examples seem to be above a parabola.

VI. Iterated Rock-Paper-Scissors

Consider playing the game of rock-paper-scissors (RPS) against an opponent over many rounds. One round of RPS consists of both players simultaneously choosing either rock (R), paper (P), or scissors (S). Each player does not know which move the other will make. The winner of the round is determined as follows: rock beats scissors, scissors beat paper, and paper beats rock. You receive a reward of +1 for each win, and -1 for each loss. When both players choose the same move, the round is a tie (resulting in a reward of 0). This problem will consider modeling the opponent using probabilistic models.

1. Assume you are playing an opponent that chooses R, P, or S independently at random at each round. Describe a probabilistic model for this opponent (Hint: it contains two parameters). In terms of those parameters, what is the probability that the opponent plays R? What is the probability that the opponent plays R three times in a row? Given you won the past round by playing P against the opponent's R, what is the probability that the opponent plays R in the next three rounds?
2. Suppose you have observed that the opponent has played rock N_R times, paper N_P times, and scissors N_S times. What are the maximum likelihood estimates of the parameters of your model? What is your expected reward for choosing R? What are two potential disadvantages of maximum likelihood estimation?
3. Now you switch to playing a smarter opponent, where the opponent tries to modify its behavior depending on how well it has done in the past. Assume now that this opponent always chooses the move that has given it the highest average reward in the past (if multiple moves have equal average reward, the move is chosen at random). Can you do better, the same, or worse against this opponent than one that picks moves independently at random?