
Midterm for 6.034 Spring 2010
Name:

20	20	20	20

Good luck!

Question #1

20 points

Indicate whether the following statement is true. Provide a short explanation. You will not be awarded any points if your answer is not justified

- (i) (2 points) Alpha-beta pruning can alter the computed minimax value of the root of a game search tree.

- (ii) (2 points) When doing alpha-beta pruning on a game tree which is traversed from left to right, the leftmost branch will never be pruned.

- (iii) (2 points) Iterative Deepening (ID) is preferred to DFS in CSP as it is always guaranteed to stop.

- (iv) (2 points) If $g(s)$ and $h(s)$ are two admissible A^* heuristics, then the following heuristic $\frac{1}{2}g(s) + \frac{1}{4}h(s)$ is also admissible.

- (v) (2 points) For a search problem, the path returned by uniform cost search may change if we add a positive constant C to every edge cost.

- (vi) (2 points) With every round of perceptron training, the number of incorrectly classified points always decreases.

-
- (vii) (2 points) If the data is linearly separable, the perceptron algorithm will always find the same separator, no matter how the weights are initialized.
- (viii) (2 points) Assuming significant noise in the training data, we are likely to increase prediction accuracy on the test set by increasing k in a k -NN classifier.
- (ix) (2 points) Assuming significant noise in the training data, we are likely to increase prediction accuracy on the test set by increasing the height of the tree (by doing less pruning).
- (x) (2 points) Assuming significant noise in the training data, a learning algorithm with a low variance always has a lower cross-validation error than an algorithm with a high variance.

Question #2

20 points

1. (6 points) The following table shows a training set of several 2-dimensional training data points and their true labels. Let (x_1, x_2) denote each data point, y the corresponding true label, and i the corresponding position in the training set.

i	y	x_1	x_2
1	+1	0	-1
2	+1	1	-3
3	+1	1	-2
4	+1	2	-3
5	+1	3	-4
6	-1	3	-2
7	-1	3	1
8	-1	4	-1


In the table below, trace the perceptron parameter updates for the primal weight vector $\bar{\mathbf{w}}$ and the dual variables $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_8)$ until convergence, where α_i is the dual variable for the i -th data vector.

Remarks:


- Let learning rate $\eta = 1.0$
- The first coordinate of $\bar{\mathbf{w}}$ is the bias b of the hyperplane
- A total of six (6) updates will be performed
- **Fill in a new row only when parameters are updated.**

Iteration	i	$\bar{\mathbf{w}}$			$\boldsymbol{\alpha}$								
0	-	0	0	0	0	0	0	0	0	0	0	0	0
1													

-
2. (2 points) Draw the final decision boundary and shade the region which labels points as +1.



3. (4 points) Consider that case where only one training data point is removed. In some cases, removing a single data point will change the decision boundary of the learned classifier. Circle data points that will cause such a change.



4. (2 points) If the data points were re-ordered, would your answer in the previous question change? If yes, give an example. If no, explain why (in two or fewer sentences).


5. (2 points) Here, we consider using the decision tree learning algorithm on the same data set. Suppose each split can only be binary, i.e. it takes one of the following forms:

$$x_1 \leq t,$$


$$x_2 \leq t$$

where t is a real number. Also, let the split point be the mid-point of the closest point on either side, i.e. t is a multiple of 0.5.

Draw the decision boundary after each split. (There might be fewer than four splits.) In the final split, shade the region that labels points as +.



-
6. (4 points) Consider that case where only one training data point is removed. In some cases, removing a single data point will change the decision boundary of the learned classifier. Circle data points that will cause such a change.



Question #3

20 points

In this question, we consider several variants of the A^* **tree** search algorithm as follows:

- (a) Standard A^*
- (b) A^* , but we apply the goal test before enqueueing nodes rather than by selecting nodes from the queue Q
- (c) A^* , but $PickNode(Q)$ based only on lowest $g(n)$ (ignoring $h(n)$)
- (d) A^* , but $PickNode(Q)$ based only on lowest $h(n)$ (ignoring $g(n)$)
- (e) A^* , but $PickNode(Q)$ based on lowest $g(n) + h(n')$
- (f) A^* , but $PickNode(Q)$ based on lowest $g(n') + h(n)$

where

Q		a queue from which a node is picked
n		a node
n'		a parent of n
$g(n)$		the cumulative path cost from the start node to node n
$h(n)$		a lower bound on the shortest path from node n to a goal state
$PickNode(Q)$		a function which picks from Q according to some criterion

We also assume all costs are *non-negative*.

The search algorithm is reproduced below for your convenience:

```
Initialize Q with start node S
Repeat
  If Q is empty
    Fail
  else
    n = PickNode(Q)
  If state(n) is goal
    return n
  Dequeue n from Q
  Enqueue children of n to Q
```

1. (3 points) Which of the above variants are complete (i.e., guaranteed to find a solution when there is one), assuming all heuristics are admissible?

2. (3 points) Which of the above variants are optimal (i.e., guaranteed to find the best cost solution), assuming all heuristics are admissible?

Upper Bounds A^* exploits lower bounds h on the true completion cost h^* . Suppose now that we also have an upper bound $k(n)$ on the best completion cost (i.e., $\forall n, k(n) \geq h^*(n)$). Consider the point at which you are inserting a node n into the queue.

3. (3 points) Assume you are required to preserve optimality. In response to n 's insertion, can you ever delete any nodes m currently on the queue? If yes, state a general condition under which nodes m can be discarded, if not, state why not. Your answer should involve various path quantities (g, h, k) for both the newly inserted node n and other nodes m on the queue.

-
4. (3 points) Assume that your goal now is to find some solution of cost less than some threshold t (if one exists). Your solution does not have to be optimal. In response to n 's insertion, can you ever delete any nodes m currently on the queue? If yes, state a general condition, if not, state why not. Your answer should involve various path quantities (g, h, k) for both the newly inserted node n and other nodes m on the queue.

ϵ -Admissible Heuristics: Suppose that we have a heuristic function which is not admissible, but ϵ -admissible. That is, for some known $\epsilon > 0$ and for all nodes n :

$$h(n) \leq h^*(n) + \epsilon$$

where $h^*(n)$ is the optimal completion cost. In other words, $h(n)$ is never more than ϵ from being optimal.

5. (1 point) Is using A^* with an ϵ -admissible complete? Briefly justify.
6. (2 points) Assuming we utilize an ϵ -admissible heuristic in standard A^* search, how much worse than the optimal solution c^* might our solution be?
7. (5 points) Suggest a modification to the A^* algorithm which is guaranteed to yield an optimal solutions using an ϵ -admissible heuristic with fixed, known ϵ . Justify your answer.

Question #4

20 points

1. (6 points) In this question you will perform alpha-beta pruning on a standard game tree from left to right. The game tree is shown as a table as in the problem set — Each cell in the table corresponds to a node in the tree and the cells directly below a cell represent the children of that node. In each box, you have to fill in two numbers as in the problem set. The numbers are different depending if the box is the top or the bottom row of the MAX or MIN function:

- In the top box: the values of alpha and beta passed into each call to MAX-VALUE and MIN-VALUE (not values which are returned from MAX-VALUE and MIN-VALUE)
- In the bottom box: the value returned by MAX-VALUE or MIN-VALUE, and followed by the number of static evaluations that were done in the subtree below the box (not the actual MIN-MAX value of each subtree)

Remarks:


- If a node is pruned, so that no alpha or beta values are computed for it, enter “X” in the boxes.
- Do not leave any box empty.

MAX														
	10 11													
MIN														
MAX														
STATIC	5	-1	-1	10	2	7	-8	1	10	4	7	10	11	-4

-
2. The standard Minimax algorithm calculates worst-case values in a *zero-sum* two player game, i.e., a game in which for all terminal states s , the utilities for players A (MAX) and B (MIN) obey $U_A(s) + U_B(s) = 0$. In the zero sum case, we know that $U_A(s) = -U_B(s)$ and so we can think of player B maximizing $U_B(s)$ as simply minimizing $U_A(s)$.

In this question, we will consider the *nearly zero sum* case, in which $|U_A(s) + U_B(s)| < \epsilon$ at all terminal nodes s for some ϵ which is known in advance. For example, the game tree in the following figure is nearly zero sum for $\epsilon = 2$. In this game, each player still maximizes his or her own utility, but this no longer always minimizes the other player's utility.

- (a) (8 points) Draw a "X" in each node in this game tree which could be pruned with the appropriate generalization of alpha-beta pruning. Assume that the exploration is being done in the standard left to right depth-first order and the value of ϵ is known to be 2. Make sure you make use of ϵ in your reasoning.



- (b) (6 points) Give a general condition under which a child n of a B node (MIN node) b can be pruned. Your condition should generalize α -pruning and should be stated in terms of quantities such as the utilities $U_A(s)$ and/or $U_B(s)$ of relevant nodes s in the game tree, the bound ϵ , and so on. Do not worry about ties.