CS 337 (Spring 2019): Class Test 3

Instructor: Shivaram Kalyanakrishnan

11.05 a.m. – 12.20 p.m., March 27, 2019, LA 201

Total marks: 20

Note. Provide brief justifications and/or calculations along with each answer to illustrate how you arrived at the answer.

Question 1. Consider the linearly-separable set of data points shown below. The points are in 2-dimensional (x_1-x_2) space. A, B, and C are labelled positive, while D, E, F, and G are negative.

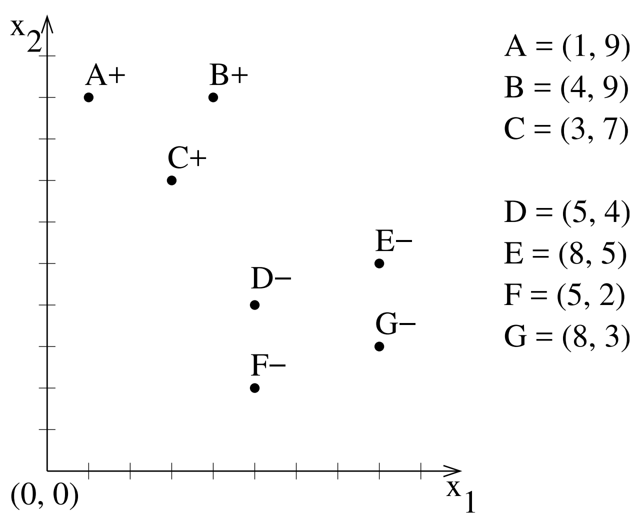

1a. What is the equation of the line separating the positives from the negatives with *maximum margin*? You do not have to furnish a formal proof; an informal argument based on your visual judgement will suffice. [3 marks]

1b. Which among the 7 points are support vectors for your answer to 1a? [1 mark]

1c. What is the value of the maximum margin for this data set? [1 mark]

Question 2. A fly F is inside a bottle with three chambers: C_1, C_2, and C_3, as shown below. F is initially in chamber C_1. At each time step, F can go either left or right. From C_1, going left releases F from the bottle; from C_2 and C_3 going left results in arriving at the chamber immediately to the left. Going right from C_1 and C_2 takes F to the chamber immediately to the right. Going right from C_3 retains F in C_3.

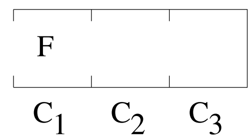

At each time step, suppose F decides to go left with probability $p \in (0, 1)$ and right with probability $q = 1 - p$. What is the expected number of time steps that F remains in the bottle? [5 marks]

Question 3. This question pertains to heuristic functions for search. Assume that every state in the problem instance has a finite-length path to a goal.

3a. What is an *admissible* heuristic? What is a *consistent* heuristic? [2 marks]

3b. Consider statements S1 and S2.

(S1) Every admissible heuristic is consistent.

(S2) Every consistent heuristic is admissible.

Only one of these statements is true; which one? [1 mark]

3c. For whichever statement in 3b you identified as *true*, provide a proof. You do not have to prove that the other statement is false. [2 marks]

Question 4. Consider the game tree shown below: first Max plays, then Min, and then the game ends. Each player has two actions at each stage. As usual, each leaf shows the reward for Max upon termination; Min obtains the negative of this reward.

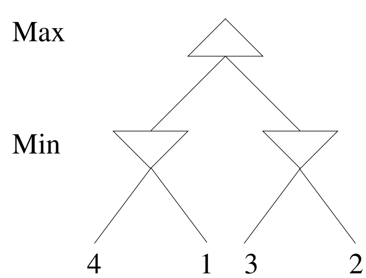

Suppose Max uses α-β pruning to compute its minimax value, and suppose that at each node, it expands a yet-unexpanded child picked *uniformly at random*. What is the expected number of leaves that it expands? [5 marks]

Solutions

1a. From the visualisation provided, it is clear that any separating line must intersect DC. Of all such lines, the one that maximises the minimum distance to C and D is the perpendicular bisector of DC, given by

$$\begin{aligned}
\left((x_1, x_2) - \frac{C+D}{2}\right) \cdot (C - D) &= 0, \text{ that is} \\
((x_1, x_2) - (4, 11/2)) \cdot (-2, 3) &= 0, \text{ that is} \\
-2x_1 + 3x_2 - (-8 + 33/2) &= 0, \text{ that is} \\
4x_1 - 6x_2 + 17 &= 0.
\end{aligned}$$

For any other line, either the distance from C will be larger or the distance from D will be larger. To show that this line is a maximum margin separator for the entire data set, it suffices to verify that its distances to A, B, E, F, and G are at least as much as $\frac{CD}{2}$. This can be done either visually or analytically.

The distance of a point $\mathbf{z}$ to a line $\mathbf{w} \cdot \mathbf{x} + b = 0$ is given by $\frac{|\mathbf{w} \cdot \mathbf{x} + b|}{\|\mathbf{w}\|}$. The distances from each of the data points to the maximum margin line are given below.

$$\begin{aligned}
\text{From } A: \frac{|4(1) - 6(9) + 17|}{\sqrt{52}} &= \frac{33}{\sqrt{52}}. \\
\text{From } B: \frac{|4(4) - 6(9) + 17|}{\sqrt{52}} &= \frac{21}{\sqrt{52}}. \\
\text{From } C: \frac{|4(3) - 6(7) + 17|}{\sqrt{52}} &= \frac{13}{\sqrt{52}}. \\
\text{From } D: \frac{|4(5) - 6(4) + 17|}{\sqrt{52}} &= \frac{13}{\sqrt{52}}. \\
\text{From } E: \frac{|4(8) - 6(5) + 17|}{\sqrt{52}} &= \frac{19}{\sqrt{52}}. \\
\text{From } F: \frac{|4(5) - 6(2) + 17|}{\sqrt{52}} &= \frac{25}{\sqrt{52}}. \\
\text{From } G: \frac{|4(8) - 6(3) + 17|}{\sqrt{52}} &= \frac{31}{\sqrt{52}}.
\end{aligned}$$

1b. C and D. Every other point is farther away from the maximum-margin separating line.

1c. $\frac{CD}{2} = \frac{\sqrt{13}}{2}$.

2. For $i \in \{1, 2, 3\}$, let x_i denote the expected number of time steps to termination starting at Chamber C_i. The description of the problem provides the following recurrence.

$$x_1 = 1 + p(0) + qx_2,$$
$$x_2 = 1 + px_1 + qx_3,$$
$$x_3 = 1 + px_2 + qx_3,$$

which yields

$$x_1 = \frac{p^2 - p + 1}{p^3},$$
$$x_2 = \frac{p^2 + 1}{p^3},$$
$$x_3 = \frac{2p^2 + 1}{p^3}.$$

The required answer is x_1.

3a.

- A heuristic h is admissible if for every node n,

$$h(n) \leq \text{optimal-cost-to-goal}(n).$$

- A heuristic h is consistent if for every pair of nodes n and n' and action a such that $\text{transition}(n, a) = n'$,

$$h(n) \leq \text{cost}(n, a, n') + h(n').$$

3b. S2 is true.

3c. Let h be consistent; we have to show that it is also admissible. We can show the result by induction on the length of the shortest optimal path from a node to the goal.

Since, by convention, h is 0 for goal states, we trivially have $h(n) = \text{optimal-cost-to-goal}(n) = 0$ for goal states n. For some $d \geq 0$, assume that all nodes n that have a d-length optimal path to goal satisfy $h(n) \leq \text{optimal-cost-to-goal}(n)$. Consider an arbitrary node n whose shortest optimal-cost path is of length $d + 1$; let this path reach a goal state n_{goal}, and be of the form

$$n, a, n', a', n'', a'', \ldots, n_{\text{goal}}.$$

Clearly n' has a d-length shortest optimal-cost path, and by our induction hypothesis, $h(n') \leq \text{optimal-cost-to-goal}(n')$. Since $\text{optimal-cost-to-goal}(n') = \text{cost}(n, a, n') + \text{optimal-cost-to-goal}(n')$ and h is consistent, we get

$$\text{optimal-cost-to-goal}(n') = \text{cost}(n, a, n') + \text{optimal-cost-to-goal}(n') \geq \text{cost}(n, a, n') + h(n') \geq h(n),$$

which completes the proof.

4. Minimax search with α-β pruning expands nodes depth-first. The current question involves picking children to expand in a sequence picked uniformly at random. For the game tree in question, there are 8 distinct depth-first traversals of the form $(ab)(cd)$, where a and b are, in sequence, children of the first Min node expanded by Max, and c and d are, in sequence, children of the second Min node expanded by Max. These traversals, which are equally likely, are listed in the table below.

Depth-first traversal	Number of nodes expanded
(14)(23)	4
(14)(32)	4
(41)(23)	4
(41)(32)	4
(23)(14)	3
(23)(41)	4
(32)(14)	3
(32)(41)	4

If the traversal is $(ab)(cd)$, the a, b, and c will necessarily be expanded. However, d will be expanded if and only if $c > \min\{a, b\}$. By going through each case, we discover that d is expanded in 6 of the 8 traversals. Hence, the expected number of leaves expanded is $\frac{1}{8}(6 \times 4 + 2 \times 3) = \frac{15}{4}$.