


Section 12: Final Exam Review

*Lecturer: Ariel Procaccia**Authors: Janani Sekar and Catherine Cui*

AI Game Playing

Problem from 2021 Midterm


Problem 1 The following game tree describes a zero-sum game between a max-player (up triangles) and min-player (down triangles). Assign payoffs (for the max player) to the leaves in such a way that alpha-beta pruning would have to examine all 8 leaves, i.e. no pruning will take place. Assume leaves are examined from left to right.



1. Refer to the 2021 midterm solutions for the solution to the first part of the problem. There are many possible solutions that work. Notice that the general order that works has the poorer payoffs at leaves that are further left.

More alpha beta pruning

Problem 2 In the following game tree representing a 2-player zero-sum game, up-triangles represent the max player, and down-triangles represent the min player. The different final outcomes or leaves have been provided in the image below. Fill out the game tree all the way to the top node. Then, use alpha-beta pruning to reduce the number of nodes you would have had to explore.



2. We prune the leaf with payoff 2 and the entire rightmost subtree with leaves 7 and 8. If you filled out the tree before pruning, the values at the nodes, going top to bottom and left to right are: 5, 5, 4, 5, 6, 4, 8

Social Choice

Problem 3 Create a voter preference profile and alternatives to demonstrate that Borda count is not a Condorcet consistent voting rule.

3. There are many different preference profiles that could work for this question, but the example from the Social Choice lecture is one such case depending on how you break ties.

Problem 4 Consider a voting rule where voters award 1 point to their most preferred alternative and -1 points to their least preferred alternative. Is this voting rule Condorcet consistent? If yes, explain why. If no, give a counterexample.


4. This voting rule is not Condorcet consistent. Consider the following preference profiles:

- 3 voters prefer: A, B, D, C
- 2 voters prefer: C, A, B, D
- 2 voters prefer: C, B, D, A

The Condorcet winner here is C. However, with our strange voting rule, the scores for each of the alternatives are $A=1$, $B=0$, $C=1$, $D=-2$. If we break ties alphabetically, for example, A wins with this voting rule instead of C.

Bayesian Networks

Problem 5 Consider the following Bayes net and answer the following questions:



- (a) What is $\mathbb{P}(\text{Grass wet} = \text{True} \mid \text{Cloudy} = \text{True})$?
- (b) Suppose we did likelihood weighting on this network, where we observe evidence $C = t$, $W = f$ and we sample $S = t, R = f$. What is the weight of this sample?

14. Using the diagram above, we have

(a) Using LOTP we have

$$\begin{aligned}
 P[W = T|C = T] &= P[W = T|C = T, S = T, R = T] \cdot P[S = T, R = T|C = T] \\
 &\quad + P[W = T|C = T, S = T, R = F] \cdot P[S = T, R = F|C = T] \\
 &\quad + P[W = T|C = T, S = F, R = T] \cdot P[S = F, R = T|C = T] \\
 &\quad + P[W = T|C = T, S = F, R = F] \cdot P[S = F, R = F|C = T]
 \end{aligned}$$

Because in a Bayes net, each random variable is conditionally independent of its predecessors given its parents, this simplifies to

$$\begin{aligned}
 P[W = T|C = T] &= P[W = T|S = T, R = T] \cdot P[S = T, R = T|C = T] \\
 &\quad + P[W = T|S = T, R = F] \cdot P[S = T, R = F|C = T] \\
 &\quad + P[W = T|S = F, R = T] \cdot P[S = F, R = T|C = T] \\
 &\quad + P[W = T|S = F, R = F] \cdot P[S = F, R = F|C = T]
 \end{aligned}$$


Plugging in the probabilities given gives

$$0.99 \cdot 0.1 \cdot 0.8 + 0.9 \cdot 0.1 \cdot 0.2 + 0.9 \cdot 0.9 \cdot 0.8 + 0$$

(b) At first $w = 1$. Since C is evidence, $w = 1 \cdot 0.5 = 0.5$. Then we sample $S = t, R = f$ and so W is false with probability 0.1. Therefore $w = 0.5 \cdot 0.1 = 0.05$.

HMM

Problem 6 *Mr. Red is happy some days and angry others. We can observe when he smiles, frowns, laughs, and yells, but not his actual emotional state. There can only be one state transition per day. It can be to the happy state or angry state. The HMM is shown below and the probability Mr. Red is happy on day 0 is 0.5.*



Let s_t be the state on day t and let e_t be the observation on day t . Suppose $e_1 = Frown$. Then what is the probability $s_1 = Happy$?

15. First we look for the probability $s_1 = Happy$ given $e_1 = Frown$.

$$\begin{aligned} P[s_1 = Happy | e_1 = Frown] &\propto P[e_1 = Frown | s_1 = Happy] \cdot (P[s_0 = Happy] \cdot P[s_1 = Happy | s_0 = Happy] \\ &\quad + P[s_0 = Angry] \cdot P[s_1 = Happy | s_0 = Angry]) \\ &= 0.1 \cdot (0.5 \cdot 0.8 + 0.5 \cdot 0.2) \\ &= 0.05 \end{aligned}$$

Now we look for the probability $s_1 = Angry$ given $e_1 = Frown$.

$$\begin{aligned} P[s_1 = Angry | e_1 = Frown] &\propto P[e_1 = Frown | s_1 = Angry] \cdot (P[s_0 = Happy] \cdot P[s_1 = Angry | s_0 = Happy] \\ &\quad + P[s_0 = Angry] \cdot P[s_1 = Angry | s_0 = Angry]) \\ &= 0.5 \cdot (0.5 \cdot 0.2 + 0.5 \cdot 0.2) \\ &= 0.1 \end{aligned}$$

In conclusion, after normalizing, the probability $s_1 = Happy$ given $e_1 = Frown$ is $\frac{0.05}{0.15} = \frac{1}{3}$

MDPs/Reinforcement Learning

Problem 7 Which of the following statements are true for an MDP? Briefly explain why.

- (a) If one is using value iteration and the values have converged, the optimal policy based on the current values must have converged as well.
- (b) Policies found by value iteration are superior to policies found by policy iteration, assuming that both algorithms have converged.

16.

- (a) Yes! Since the values have converged, the values don't change anymore. Therefore, the policy between iterations also stops changing.
- (b) False. Assuming both algorithms have converged, policies found by value iteration and policy iteration should be equal.

This is because both value iteration and policy iteration improve the policy until no improvements can be made. In other words, they find an optimal policy. If value iteration found a superior policy, then policy iteration wouldn't have found the optimal policy.

Linear Classification

Problem 8

(5 points) Higher Dimensions. Let's figure out how to linearly separate the interior of an arbitrary circle or ellipse from its surroundings. The general equation for a circle (and hence the decision boundary) is $(x_1 - a)^2 + (x_2 - b)^2 - r^2 = 0$, and the general equation for an ellipse is $c(x_1 - a)^2 + d(x_2 - b)^2 - 1 = 0$.

- (a) Expand out the equation for the circle and show what the weights w_i would be for the decision boundary in the four-dimensional feature space x_1, x_2, x_1^2, x_2^2 . Explain why this means any circle is linearly separable in this space.
- (b) Do the same for ellipses in the same feature space $(x_1, x_2, x_1^2, x_2^2, x_1x_2)$.

17.

- (a) Expanding out the equation for the circle gives

$$\begin{aligned}(x_1 - a)^2 + (x_2 - b)^2 - r^2 &= 0 \\ x_1^2 - 2ax_1 + a^2 + x_2^2 - 2bx_2 + b^2 - r^2 &= 0 \\ x_1^2 + x_2^2 - 2ax_1 - 2bx_2 + a^2 + b^2 - r^2 &= 0\end{aligned}$$

For the decision boundary in the four-dimensional feature space x_1, x_2, x_1^2, x_2^2 , the weights are $-2a, -2b, 1$, and 1 respectively.

Any circle is linearly separable in this four-dimensional feature space because while the decision boundary isn't linear in terms of x_1, x_2 , it is clearly linear in terms of x_1, x_2, x_1^2, x_2^2 .

- (b) Expanding out the equation for the ellipse gives

$$\begin{aligned}c(x_1 - a)^2 + d(x_2 - b)^2 - 1 &= 0 \\ c(x_1^2 - 2ax_1 + a^2) + d(x_2^2 - 2bx_2 + b^2) - 1 &= 0 \\ cx_1^2 - 2acx_1 + ca^2 + dx_2^2 - 2bdx_2 + db^2 - 1 &= 0 \\ cx_1^2 - 2acx_1 + dx_2^2 - 2bdx_2 + ca^2 + db^2 - 1 &= 0\end{aligned}$$

For the decision boundary in the five-dimensional feature space $x_1, x_2, x_1^2, x_2^2, x_1x_2$, the weights are $-2ac, -2bd, c, d$, and 0 respectively.

Any ellipse is linearly separable in this five-dimensional feature space because while the decision boundary isn't linear in terms of x_1, x_2 , it is clearly linear in terms of $x_1, x_2, x_1^2, x_2^2, x_1x_2$.

Decision Trees

Problem 9 Using the table above, answer the following question:

Day	Weather	Temperature	Humidity	Wind	Play?
1	Sunny	Hot	High	Weak	No
2	Cloudy	Hot	High	Weak	Yes
3	Sunny	Mild	Normal	Strong	Yes
4	Cloudy	Mild	High	Strong	Yes
5	Rainy	Mild	High	Strong	No
6	Rainy	Cool	Normal	Strong	No
7	Rainy	Mild	High	Weak	Yes
8	Sunny	Hot	High	Strong	No
9	Cloudy	Hot	Normal	Weak	Yes
10	Rainy	Mild	High	Strong	No

What is the information gain from splitting on **humidity**? What is the information gain from splitting on **wind**?

18. We find (1) the information gain from splitting on **humidity** and (2) the information gain from splitting on **wind**.

- (a) Splitting on **humidity** gives two different subsets of examples: *High* with probability $\frac{7}{10}$, *Normal* with probability $\frac{3}{7}$. Note that if humidity's *High*, we play $\frac{3}{7}$ of the time; if humidity's *Normal* we play $\frac{2}{3}$ of the time. Therefore

$$Gain(weather) = 1 - \left[\frac{7}{10} \cdot B\left(\frac{3}{7}\right) + \frac{3}{10} \cdot B\left(\frac{2}{3}\right) \right]$$

- (b) Splitting on **wind** gives two different subsets of examples: *Weak* with probability $\frac{4}{10}$, *Strong* with probability $\frac{6}{10}$. Note that if wind's *Weak*, we play $\frac{3}{4}$ of the time; if wind's *Strong* we play $\frac{2}{6}$ of the time. Therefore

$$Gain(weather) = 1 - \left[\frac{4}{10} \cdot B\left(\frac{3}{4}\right) + \frac{6}{10} \cdot B\left(\frac{2}{6}\right) \right]$$

Problem 10 Master Yoda is concerned about the number of Jedi apprentices that have turned to the Dark Side, so he's decided to train a decision tree on some historical data to help identify problem cases in the future. The following table summarizes whether or not each of the 12 initiates turned to the Dark Side based on their age when their Jedi training began, whether or not they completed their training, their general disposition, and their species.

Dark Side	Age Started Training	Completed Training	Disposition	Species
0	5	1	Happy	Human
0	9	1	Happy	Gungan
0	6	0	Happy	Wookiee
0	6	1	Sad	Mon Calamari
0	7	0	Sad	Human
0	8	1	Angry	Human
0	5	1	Angry	Ewok
1	9	0	Happy	Ewok
1	8	0	Sad	Human
1	8	0	Sad	Human
1	6	0	Angry	Wookiee
1	7	0	Angry	Mon Calamari

- (a) What is the initial entropy of Dark Side?
- (b) What attribute would the decision-tree building algorithm choose to use for the root of the tree? (You can use intuition here or actually calculate out the information gain of splitting on each feature).

19. Using the chart we get

- (a) $-\frac{5}{12} \log \frac{5}{12} - \frac{7}{12} \log \frac{7}{12} = 0.9798$
- (b) Using intuition and eyeballing, we see that splitting on **Completed Training**, everyone who hasn't completed their training goes to the Dark Side. This doesn't happen with any of the other features. Therefore, it seems like a good idea to first split on **Completed Training**.

Fairness

Problem 11 (Fill in the Blank Problem). **[BLANK]** requires the outcome to be independent of the attribute $G \in \{0, 1\}$.

20. Looking at the definition of Demographic Parity, it's clear that demographic parity requires the probability of getting an outcome to be independent of attribute G (probability of some outcome is the same whether you have $G = 0$ or $G = 1$).