DSC 40A - Homework 7

Due: Thursday, June 1 at 11:59pm

Write your solutions to the following problems by either typing them up or handwriting them on another piece of paper. Homeworks are due to Gradescope by 11:59pm on the due date. You can use a slip day to extend the deadline by 24 hours.

Homework will be evaluated not only on the correctness of your answers, but on your ability to present your ideas clearly and logically. You should **always explain and justify** your conclusions, using sound reasoning. Your goal should be to convince the reader of your assertions. If a question does not require explanation, it will be explicitly stated.

Homeworks should be written up and turned in by each student individually. You may talk to other students in the class about the problems and discuss solution strategies, but you should not share any written communication and you should not check answers with classmates. You can tell someone how to do a homework problem, but you cannot show them how to do it.

For each problem you submit, you should **cite your sources** by including a list of names of other students with whom you discussed the problem. Instructors do not need to be cited.

This homework will be graded out of 50 points. The point value of each problem or sub-problem is indicated by the number of avocados shown.

Notes:

- This assignment has six problems instead of the usual five, because some are short. The total number of points for all problems is 50 as usual.
- For full credit, make sure to **assign pages to questions** when you upload your submission to Gradescope.
- For Problem 6 part (b), code your answer in the supplementary Jupyter notebook (linked). You'll need to turn in your completed Python file to Gradescope separately from the rest of this homework, in a file called `hw7code.py`. This part will be autograded, so no explanation is needed.

Problem 0. Reflection and Feedback Form

Make sure to fill out this Reflection and Feedback Form, linked here for three points on this homework! This form is primarily for your benefit; research shows that reflecting and summarizing knowledge helps you understand and remember it.

Problem 1. Independence and Conditional Independence

Consider the sample space $S = \{a, b, c, d, e, f, g\}$ with associated probabilities given in the table below.

outcome	a	b	c	d	e	f	g
probability	$\frac{5}{21}$	$\frac{2}{21}$	$\frac{1}{21}$	$\frac{4}{21}$	$\frac{2}{21}$	$\frac{4}{21}$	$\frac{3}{21}$

Let $X = \{d, e\}$ and $Y = \{e, f\}$. Remember to show your work for all calculations.

a) Are X and Y independent?

b) Determine if X and Y are conditionally independent given each of the following events Z.

1. $Z = \{a, b, d, e, f, g\}$
2. $Z = \{a, d, e, f, g\}$
3. $Z = \{d, e, f, g\}$

Problem 2. Independence and Complements

Let E and F be two events in a sample space S, with $0 < P(F) < 1$.

a) If $P(E|F) = P(E|\overline{F})$, must it be true that E and F are independent? Provide a proof of independence, or give a counterexample by specifying a sample space S and two dependent events E and F that satisfy the given conditions.

b) If $P(E|F) = P(\overline{E}|F)$, must it be true that E and F are independent? Provide a proof of independence, or give a counterexample by specifying a sample space S and two dependent events E and F that satisfy the given conditions.

Problem 3. Probability Theory

Let S be a sample space, and let A, E_1, E_2, E_3 be events in that sample space. Suppose that $E_1 \cap E_2$, $E_1 \cap E_3$, and $E_2 \cap E_3$ are all empty. Given the following probabilities, find $P(E_2|A)$:

$$\begin{aligned} P(E_1) &= 1/6 \quad & P(A|E_1) &= 1/9 \\ P(E_2) &= 1/3 \quad & P(A|E_2) &= 1/7 \\ P(E_3) &= 1/2 \quad & P(A|E_3) &= 1/5 \end{aligned}$$

Problem 4. Double Deck of Cards

A standard deck of cards contains 52 cards. There are 13 cards in each of 4 suits (hearts ♡, spades ♠, diamonds ♢, and clubs ♣.) Within a suit, the 13 cards each have a different rank. In ascending order, these ranks are 2, 3, 4, 5, 6, 7, 8, 9, 10, Jack, Queen, King, Ace.

You are playing a four-player card game using two regular decks of cards. Each player will be dealt 26 cards as follows:

- The first deck of cards will be randomly shuffled and dealt out, 13 cards to each of the four players.
- Then the second deck of cards will be randomly shuffled and dealt out, 13 cards to each of the four players.

a) Let Event A = 'you are dealt two Kings of Hearts' and Event B = 'some other player has a pair of Kings of Hearts'. Are these two events independent of each other?

b) Let Event A = 'you are dealt two Kings of Hearts' and Event C = 'some other player has a pair of Aces of Hearts'. Are these two events independent of each other?

c) Now suppose that you know you do not have any Aces of Hearts. Let Event A = 'you are dealt two Kings of Hearts' and Event C = 'some other player has a pair of Aces of Hearts'. Now that you know you don't have any Aces of Hearts, are these two events independent of each other?

d) Let Event A = 'you are dealt two Kings of Hearts' , Event C = 'some other player has a pair of Aces of Hearts', and Event D = 'you do not have any Aces of Hearts'. State the result of parts (b) and (c) in words, in terms of independence and conditional independence.

Problem 5. Baby Avi

When Avi was a baby, he was a picky eater. When he was offered a food, he'd eat it only sometimes, depending on the type of food. Baby Avi ate

- bananas 95% of the time,
- crackers 60% of the time,
- meat 45% of the time, and
- zucchini 30% of the time.

a) Baby Avi's grandpa gave him one of the above four foods and Baby Avi ate it all up. You have no idea which of the four foods it was, so assume it was equally likely to be any of them.

Given that Baby Avi ate the food, what's the probability that it was a banana? A cracker? Meat? Zucchini? Show your work.

b) As before, Baby Avi's grandpa gave him one of the above four foods and Baby Avi ate it all up. This time, suppose you know that Baby Avi's grandpa offers

- bananas 15% of the time,
- crackers 40% of the time,
- meat 10% of the time, and
- zucchini 35% of the time.

Given that Baby Avi ate the food, what's the probability that it was a banana? A cracker? Meat? Zucchini? Show your work.

c) Compare your answers to part (a) and part (b) above. Identify which of the four probabilities you computed increased and which decreased, and explain why this makes sense intuitively.

Problem 6. Comic Characters

In this problem, we'll work with a dataset of characters from comic books. At first, we look at a limited set of 25 of the more popular comic characters. You can find this dataset on the last page of this assignment as well as on Datahub in the same directory as the supplementary Jupyter notebook (linked).

For each character, we have the following information.

Variable	Definition
`ID`	The identity status of the character ("Public Identity" or "Secret Identity")
`SEX`	Sex of the character ("Male Characters" or "Female Characters")
`ALIVE`	Whether the character is alive ("Living Characters" or "Deceased Characters")
`COMPANY`	The comic company that created the character ("DC" or "Marvel")
`ALIGN`	The alignment of the character ("Bad Characters", "Good Characters", or "Neutral Characters")

a) Use the set of 25 popular characters to make a prediction using Naive Bayes (without smoothing) for the following character's alignment:

- "Secret Identity"
- "Male Characters"

- "Living Characters"
- "Marvel"

Your prediction should be "Bad Characters", "Good Characters", or "Neutral Characters", whichever is most likely according to Naive Bayes. Do this part by hand, and show your work. You can check your answer using the code you'll write in part (b).

b) In the supplementary Jupyter notebook (linked), we've provided not only the dataset of 25 popular characters, but also a much larger dataset of over ten thousand comic characters. For this part, you will be implementing a Naive Bayes classifier (without smoothing) to predict the alignment of any character based on their features (ID, SEX, ALIVE, COMPANY), using the larger dataset.

Complete the function `predict_align`, which takes in a DataFrame of comic characters and the features of one particular character, and returns the predicted alignment for that character according to Naive Bayes without smoothing, using the input DataFrame of characters.

Your function should return a string, either "Bad Characters", "Good Characters", or "Neutral Characters".

c) Use your `predict_align` function to predict the alignment for the same character as you did in part (a), this time using the larger dataset of comic characters. As a reminder, the features were

- "Secret Identity"
- "Male Characters"
- "Living Characters"
- "Marvel"

You should get a different prediction than you got in part (a), when you used only the 25 popular characters. Why do you get different results? What specific difference between the two datasets explains why your predictions are different?

Hint: If you're stuck, try printing out each term in the products that you calculate in `predict_align`. Compare these terms when you input the DataFrame of 25 popular characters and when you input the larger DataFrame.

Hint: Another way to get started on this is to try other combinations of features as input to `predict_align`. When using the DataFrame of 25 popular characters, try to find a combination of features that leads to a different result than you got in part (a).

name	ID	SEX	ALIVE	COMPANY	ALIGN
Wilson Fisk (Earth-616)	Public Identity	Male Characters	Living Characters	Marvel	Bad Characters
Joker (New Earth)	Secret Identity	Male Characters	Living Characters	DC	Bad Characters
Kent Nelson (New Earth)	Secret Identity	Male Characters	Deceased Characters	DC	Good Characters
Timothy Dugan (Earth-616)	Public Identity	Male Characters	Deceased Characters	Marvel	Good Characters
Samuel Wilson (Earth-616)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Hulk (Robert Bruce Banner)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Reed Richards (Earth-616)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Benjamin Grimm (Earth-616)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Iron Man (Anthony \"Tony\" Stark)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Jessica Drew (Earth-616)	Secret Identity	Female Characters	Living Characters	Marvel	Good Characters
Johnathon Blaze (Earth-616)	Secret Identity	Male Characters	Living Characters	Marvel	Good Characters
Brian Braddock (Earth-616)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Dane Whitman (Earth-616)	Secret Identity	Male Characters	Living Characters	Marvel	Good Characters
Captain America (Steven Rogers)	Public Identity	Male Characters	Living Characters	Marvel	Good Characters
Crystalia Amaquelin (Earth-616)	Public Identity	Female Characters	Living Characters	Marvel	Good Characters
Superman (Clark Kent)	Secret Identity	Male Characters	Living Characters	DC	Good Characters
Batman (Bruce Wayne)	Secret Identity	Male Characters	Living Characters	DC	Good Characters
Franklin Rock (New Earth)	Public Identity	Male Characters	Living Characters	DC	Good Characters
Cassandra Sandsmark (New Earth)	Public Identity	Female Characters	Living Characters	DC	Good Characters
Garth (New Earth)	Public Identity	Male Characters	Deceased Characters	DC	Good Characters
James Rhodes (Earth-616)	Secret Identity	Male Characters	Living Characters	Marvel	Good Characters
Deadpool (Wade Wilson)	Secret Identity	Male Characters	Living Characters	Marvel	Neutral Characters
Odin Borson (Earth-616)	Public Identity	Male Characters	Living Characters	Marvel	Neutral Characters
Otto Octavius (Earth-616)	Secret Identity	Male Characters	Deceased Characters	Marvel	Neutral Characters
Wolverine (James \"Logan\" Howlett)	Public Identity	Male Characters	Living Characters	Marvel	Neutral Characters