Bayes' Theorem

September 21 & September 28, 2022

Note: This is the first part of a two-lecture series on Bayes' Theorem–this half will explain the intuition behind the formula while the second half will give more formal definitions and proofs.

Section 0: A Motivating Example

Somebody stole the handouts from next week's class, and you have been called to solve the case. Here is what you know about the case:

Whoever stole the handouts was from the Bay Area, was around 6 feet 6 inches tall, wore vintage edition Groucho Marx glasses, and had a three-foot-tall tophat, and that the perpetrator still possesses these items to this day.

The Bureau of Strange Statistics has determined that only approximately one in five million people are 6 feet 6 inches tall and own both vintage edition Groucho Marx glasses and a three-foot-tall tophat.

The police discover a suspect, James Smith, who is from the Bay Area, is 6 feet 6 inches tall, and owns both vintage edition Groucho Marx glasses and a three-foot-tall tophat.

With just this information, what are the odds that James Smith is actually the handout thief?

(If you answered "one in five million," congratulations! The Los Angeles Jury from 1968 would agree with you, and you might just end up almost sending an innocent person to prison.)

Section 1: A Brief Introduction/Review of Probability

Since some of these sections will be fairly basic, more advanced problems will be attached to the back of this handout for those who feel that they have a good grasp of the section. These problems will not be discussed in class and are only there for those people who already firmly grasp the concepts taught in the corresponding section, so please do not feel obligated to complete them.

That said, if you elect to do the problems in the back, make sure to follow along with the lecture so you do not accidentally miss anything or fall behind.

To set up the following problems, let there be 6 red marbles and 2 blue marbles in a brown bag.

Problem 1.1: If two marbles are randomly drawn from the brown bag, without replacement, what is the probability that both are red?

Problem 1.2: In the same scenario as Problem 1.1, what is the probability that both marbles are blue?

Problem 1.3: What is the probability that one marble is blue and the other is red?

Without fully diving into things yet, we will introduce some notation. Let X be an event. Then we denote the probability that X occurs as $P(X)$.

For example, if A is the event that two red marbles are drawn in the situation of Problems 1.1, 1.2, and 1.3, then Problem 1.1 is asking to solve for $P(A)$.

Problem 1.4: Two marbles are randomly drawn from the brown bag, without replacement. Let B be the event that the first marble drawn is red. What is $P(B)$?

Problem 1.5: Two marbles are randomly drawn from the brown bag, without replacement. What is the probability that both are the same color?

If two events X and Y are mutually exclusive, then the probability that one of X or Y happens is the probability that X happens plus the probability that Y happens. For example, in problem 1.4, the probability that both were the same color was the probability that both were blue, plus the probability that both were red. In our notation, we could say that if *X* and *Y* are mutually exclusive, $P(X \text{ or } Y)=P(X)+P(Y)$.

Note that if you have events $A_1, A_2, \ldots A_n$ which are mutually exclusive and where one of the events always occurs, then $P(A_1)+P(A_2)+ \ldots P(A_n)$ is equal to one. This corresponds to the fact that the answers to Problems 1.1, 1.2, and 1.3 sum to 1, since any two marbles drawn from the bag will either be both red, both blue, or one red and one blue.

Problem 1.6: Two marbles are randomly drawn from the brown bag, without replacement. What is the probability that both marbles are red, given that the first marble drawn was red?

Given two events *X* and *Y*, we write the probability that *X* occurs given that *Y* occurs as $P(X|Y)$ (read as "the probability of *X* given *Y*"). Recalling that *A* is the event that two red marbles are drawn and *B* is the event that the first marble drawn is red, Problem 1.6 is just asking for $P(A|B)$.

Problem 1.7: Compute $\frac{P(A)}{P(B)}$. Compare this to your results for Problem 1.6. Can you explain why they must be equal?

Section 2: Monty Hall and its Cousins

Problem 2.1 (Monty Hall): You are in a game show, trying to win a car. In this game show, there are three doors, one on the left, one in the middle, and one on the right. Behind one door is a sports car, and behind the other two doors or nothing. You are allowed to choose a door, and you choose the door on the right. The game show host then says “now, I will open one of the other two doors with nothing behind it,” and opens the door on the left, revealing nothing. You are allowed to switch your guess to the unopened middle door. What is the probability that the middle door has the sports car behind it?

Problem 2.2 (Monty Hall Variant): In the same scenario as before, you choose the door on the right, but the game show host then says “now, I will open the leftmost of the remaining two doors with nothing behind it,” and opens the door on the left, revealing nothing. You are allowed to switch your guess to the unopened middle door. What is the probability that the middle door has the sports car behind it?

Problem 2.3 (Another Monty Hall Variant): In the same scenario as before, you choose the door on the right, but the game show host then says “now, I will choose one of the two other doors at random and open it,” and opens the door on the left, revealing nothing. You are allowed to switch your guess to the unopened middle door. What is the probability that the middle door has the sports car behind it?

So what is going on in these Monty Hall problems? Why is the answer different between the first and the last two? One way to think about it is that in the latter two cases the host was “luckier” that the leftmost door had nothing behind it, so the host didn’t give as much information by choosing the left door to open specifically over the center one. But this is a very fuzzy concept, and ideally we could find a cleaner way to solve these sorts of problems. To do that, let’s talk about false positives

Section 3: False Positives

Disaster has struck, and Mathland has been struck by a disease, J-21315973, which is causing people to be unable to use statistics properly. This disease is highly contagious, and hard to detect in its early stages–all we have are a few disease testing kits. Now, let us introduce some terminology:

The event of a J-21315973 test coming up positive will be denoted P_T, the event of a test coming up negative will be denoted N_T, the event of actually having the disease will be denoted P_J, and the event of not having the disease will be denoted N_J.

The *sensitivity* of any given test is equal to $P(P_T|P_J)$ for that test, and the *specificity* of the test is equal to $P(N_T|N_J)$. For example, a test that always comes up positive, without regards for whether you actually have J-21315973 or not, will have a sensitivity of 100% and a specificity of 0%.

With these definitions in mind, let's launch into the problems.

Problem 3.1: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 95%. Given that you tested positive with this test, what is the probability that you actually have J-21315973 (in other words, what is $P(P_J|P_T)$?

Trick question! Just knowing the sensitivity and specificity of the test is not actually enough information to deduce how likely you are to be positive/negative for the disease given that you tested positive or negative, as the next problems will show.

Problem 3.2: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 99%. On the other hand, exactly 0% of people in your region of Mathland have the disease. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

Problem 3.3: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 99%. On the other hand, everybody in your region of Mathland has the disease. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

On the one hand, the answers to the previous two questions should be intuitively obvious, but on the other hand it feels weird that the prior distribution of the disease somehow impacts the accuracy of a positive or negative test, especially to such a large extent. A good way to get around this feeling is to think of a sample of 10,000 people from your region of mathland, and then to see what will "tend" to happen to them. This approach will be especially helpful in the next few questions.

Problem 3.4: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 100%. Exactly 10% of people in your region of Mathland have the disease. Given that you tested negative with this test, what is the probability that you nevertheless have J-21315973?

Problem 3.5: A certain test for J-21315973 has a sensitivity of 100% and a specificity of 95%. Exactly 20% of people in your region of Mathland have the disease. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

Problem 3.6: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 95%. Exactly 10% of people in your region of Mathland have the disease. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

Believe it or not, you've probably already used Bayes' formula in answering the previous questions, even if you do not know it yet! Let's introduce some variables now.

Problem 3.7: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 99%. The probability that a random person in your region of Mathland has the disease is w. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

Problem 3.7: A certain test for J-21315973 has a sensitivity of x and a specificity of y. The probability that a random person in your region of Mathland has the disease is w. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

Problem 3.8: Provide an alternate interpretation for the denominator in your answer to Problem 3.7.

Section 4: Bayes' Theorem

In problems 3.7 and 3.8, what you essentially did was find a formula for $P(P_J|P_T)$ using only $P(P_J)$, $P(P_J|P_T)$, and $P(P_T)$. But there's no reason that we had to restrict this formula just to the specific hypothesis P_J and evidence P_T. In fact, we could just have easily used any other hypothesis H in place of P_J and evidence E in place of P_T. With that in mind, write Bayes' Theorem below:

You should have a decent understanding of why the theorem works from Section 3, but a little practice never hurts. Try solving each of the following problems twice: once intuitively and once by explicitly using Bayes' Theorem, and note that both methods generally end up doing the same thing:

Problem 4.1: You have two bags of marbles, one with 7 red and 3 blues marbles, and the other with 3 red and no blue marbles. You choose a bag randomly, then you draw a marble from the bag at random. The marble is blue. If you were to draw another marble from the same bag, what is the probability that it would be blue?

Problem 4.2: You have two bags of marbles, one with 7 red and 3 blues marbles, and the other with 3 red and 5 blue marbles. You choose a bag randomly, then you draw a marble from the bag at random. The marble is blue. If you were to draw another marble from the same bag, what is the probability that it would be blue?

Problem 4.3: A certain test for J-21315973 has a sensitivity of 90% and a specificity of 95%. Exactly 10% of people in your region of Mathland *test positive* for the disease. Given that you tested positive with this test, what is the probability that you actually have J-21315973?

If there is time, go back to the Monty Hall problems and look at them from a Bayesian perspective. Hint: set your hypothesis to be that the middle door has the sports car, and set the evidence to be that the host opened the left door *and* there was nothing behind it. If there is still time, I can make up some more Monty Hall variants or we can do some Wordle-based Bayesian exercises.

Section 5: Remember Section 0?

Now, with our new Bayesian machinery, we can finally answer the question posed in Section 0. First, let us set up our hypothesis and evidence.

Our hypothesis, *H*, should be that a given person in the Bay Area is the perpetrator of the crime, and our evidence, *E*, is that the given person (from the Bay Area) is 6 feet 6 inches tall, owns vintage edition Groucho Marx glasses, and owns a three-foot-tall tophat.

Now, let us compute. What is *P(H)*? Well, *P(H)* should just be the probability a given person in the Bay Area is the perpetrator of the crime. Now there are approximately 8 million people in the Bay Area, and only one of them is the perpetrator. So *P(H)* is $\frac{1}{8{,}000{,}000}$.

What about *P(E)*? Well, the Bureau of Strange Statistics tells us that approximately one in five million people are 6 feet 6 inches tall and own both vintage edition Groucho Marx glasses and a three-foot-tall tophat, so the odds that any person in the Bay Area satisfies these conditions, in the absence of other information, should be $\frac{1}{5{,}000{,}000}$.

And what about *P(E|H)*? Well, we know that the perpetrator is 6 feet 6 inches tall and own both vintage edition Groucho Marx glasses and a three-foot-tall tophat, so *P(E|H)* should just be 1.

Putting it all together using Bayes' Theorem, we find that $P(H|E) = \frac{P(E|H)*P(H)}{P(E)} = \frac{\frac{1}{8.000.000}\cdot 1}{\frac{1}{5{,}000{,}000}} = \frac{5{,}000{,}000}{8{,}000{,}000} = \frac{5}{8}$. So since James Smith satisfies the evidence criteria and is in the Bay Area, the probability that he is the real perpetrator is $\frac{5}{8}$, which is a lot less than the $\frac{4{,}999{,}999}{5{,}000{,}000}$ that a certain Los Angeles jury might assume.

So Bayes' Theorem comes in to save the day, and we have the true probability that James Smith is actually the handout thief. Right? Wrong.

As it turns out, there is an error in the above reasoning, and even $\frac{5}{8}$ is too high an estimate. While our approach got a lot closer to the correct answer than the Los Angeles jury ever would, there's still a little issue, and we will dive a lot more into that next week.

Bonus Section: Advanced Problems for Section 1

These problems are meant only to be done if you really feel like you already know everything you need to know about the section being taught. Since these problems aren't really part of the main lecture material, no questions will be taken about them.

6.1) Bob says: "if I draw two marbles from a bag containing *any* number of red marbles and *any* number of blue marbles (and at least two marbles), I am at least as likely to draw two of the same color as I am to draw two different colors."

a) If Bob is drawing *with* replacement, is he correct? Prove or find all counterexamples.

b) If Bob is drawing *without* replacement, is he correct? Prove or find all counterexamples.

6.2) Not to be outdone, Alice says: "if I draw any even number of marbles from a bag containing *any* number of red marbles and *any* number of blue marbles (and enough marbles to sustain my drawing), I am at least as likely to draw an even number of red marbles as I am to draw an odd number of red marbles."

a) If Alice is drawing *with* replacement, is she correct? Prove or find as many counterexamples as you can.

b) If Alice is drawing *without* replacement, is she correct? Prove or find as many counterexamples as you can.

6.3) In Section 3, I will introduce the J-21315973 virus. Explain the naming of this virus.