Logic, Reasoning, and Persuasion, Week 10-1 Handout (Version 2)

1 | Bayes Theorem

Inductive Reasoning is about *updating your opinions on evidence*:

$$\text{have opinion } Pr(H) \longrightarrow \text{get evidence } E \longrightarrow \text{update to opinion } Pr_E(H)$$

The standard theory of updating opinions on evidence says if you learn E, then your new opinion should be your old opinion *given that* E.

Conditionalization

Suppose that you have an estimate $Pr(H \mid E)$ of the probability of H, given E. Then if you get E as evidence, you should **conditionalize**: your new opinion Pr_E should be given by

$$Pr_E(H) = Pr(H \mid E).$$

For example, if $Pr(\text{you are sick} \mid \text{you test positive}) = 0.18$, then if you learn that you tested positive, you should now have the opinion $Pr_E(\text{you are sick}) = 0.18$.

To conditionalize, we have to know how to calculate conditional probabilities. We've already learned how to do this when we have probability truth tables. But in most cases, we do not have probability truth tables. Here, the following formula often comes in handy:

Bayes' Theorem

To find the probability of a hypothesis H *given* some evidence E, take the previous probability of H, $Pr(H)$, and multiply it by $\frac{Pr(E|H)}{Pr(E)}$.

$$Pr(H \mid E) = Pr(H) \times \frac{Pr(E \mid H)}{Pr(E)}.$$

Example 1: Recall the COVID-19 test case. Suppose:

1. You know that 5% of the people in your population are sick, so you think it's 0.05 likely that you are sick: $Pr(\text{sick}) = 0.05$.
2. Overall, there is a 25% chance you test positive: $Pr(+) = 0.25$.
3. You know that if you are sick, there is an 80% chance that you test positive: $Pr(+ \mid \text{sick}) = 0.80$

Problem 1: You take the test, and it comes up positive. Conditionalize with Bayes' Theorem to calculate $Pr_E(\text{sick})$, your new estimate of the probability you are sick.

2 | Bayes' Theorem with Total Probability

Using Bayes' Theorem requires knowing the probability that you get some evidence E, $Pr(E)$, which is sometimes not readily available. However, we'll often have estimates of other hypotheses, which together can be used to calculate $Pr(E)$:

The Law of Total Probability

For any evidence statement E and possible hypotheses $H_1, H_2, \ldots, H_n$, where exactly one hypothesis is true, the probability of E is given by:

$$\begin{aligned} Pr(E) = \quad & Pr(E \mid H_1)Pr(H_1) \\ & + Pr(E \mid H_2)Pr(H_2) \\ & \ldots \\ & + Pr(E \mid H_n)Pr(H_n). \end{aligned}$$

where $Pr(E \mid H)Pr(H)$ is shorthand for multiplication: $Pr(E \mid H) \times Pr(H)$.

Substituting the this formula into Bayes' Theorem, we get the following:

Bayes' Theorem with Total Probability

For any evidence statement E and possible hypotheses $H_1, H_2, \ldots, H_n$, where exactly one hypothesis is true, the probability of E conditional on any H_i is given by the following formula:

$$Pr(H_i \mid E) = Pr(H_i) \times \frac{Pr(E \mid H_i)}{Pr(E \mid H_1)Pr(H_1) + \ldots + Pr(E \mid H_n)Pr(H_n)}.$$

In particular, suppose that we are only considering one hypothesis H and its negation $\neg H$. Then, since exactly one of H and $\neg H$ is true, the probability of E is given by

$$Pr(E) = Pr(E \mid H)Pr(H) + Pr(E \mid \neg H)Pr(\neg H),$$

which allows us to calculate Bayes' Theorem with the following:

$$Pr(H \mid E) = Pr(H) \times \frac{Pr(E \mid H)}{Pr(E \mid H)Pr(H) + Pr(E \mid \neg H)Pr(\neg H)}.$$

Example 2: Suppose in the COVID-19 test case from above, we don't know the overall probability of testing positive, but we know a conditional probability:

1. You know that 5% of the people in your population are sick, so you think it's 0.05 likely that you are sick: $Pr(\text{sick}) = 0.05$.
2. ~~Overall, there is a 25% chance you test positive.~~
3. You know that if you are sick, there is an 80% chance that you test positive: $Pr(+ \mid \text{sick}) = 0.80$, and if you are not sick, there is an 15% chance that you test positive: $Pr(+ \mid \neg\text{sick}) = 0.15$.

Problem 2: You take the test, and it comes up positive. Use conditionalization with Bayes' Theorem to calculate $Pr_E(\text{sick})$.

Main Exercise: The Kahneman and Tversky Problem

Here's a problem from a study that Kahneman and Tversky ran (I copied this text from Weisberg, Ch 8, changing some numbers):

A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data:

1. 70% of the cabs in the city are Green and 30% are Blue.
2. A witness identified the cab as Blue.
3. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 80% of the time and failed 20% of the time.

What is the probability that the cab involved in the accident was blue?

Let $Pr(\text{car was blue})$ be your estimate of probability that the car in the accident is blue, before learning what the witness said. And let $Pr_{\text{witness said blue}}(\text{car was blue})$ be your estimate of probability that the car in the accident is blue, *after* learning that the witness said the car was blue.

1. The question asks for $Pr_{\text{witness said blue}}(\text{car was blue})$, the probability that the cab involved in the accident was blue, after you learn that the witness identified the cab as blue. What's your intuitive estimate of this probability?
2. Use Bayes' Theorem to calculate $Pr_{\text{witness said blue}}(\text{car was blue})$, using these steps:
   (a) Calculate
   $$Pr(\text{witness said blue} \mid \text{car was blue})$$
   and
   $$Pr(\text{witness said blue} \mid \text{car was green}).$$
   (b) Use those values and the law of total probability to find
   $$Pr(\text{witness said blue}).$$
   (c) Use the probabilities
   $$Pr(\text{witness said blue}), \quad Pr(\text{car was blue}), \text{ and}$$
   $$Pr(\text{witness said blue} \mid \text{car was blue})$$
   to calculate $Pr(\text{car was blue} \mid \text{witness said blue})$.

   By conditionalization, after learning that the witness said blue, you should have the new opinion
   $$Pr_{\text{witness said blue}}(\text{car was blue}) = Pr(\text{car was blue} \mid \text{witness said blue}).$$
3. Suppose that the facts were like in the original case in the box, except that only 50% of the cabs in the city are green. Then what would you estimate for the probability that the cab involved in the accident was blue, after learning what the witness said?

4. Suppose that the facts were like in the original case in the box, but the witness is asymmetrically reliable. If the car is Green, then the witness will identify it correctly 80% of the time but misidentify it as blue 20% of the time. But if the car is blue, the witness always correctly identifies it. Then what would you estimate for the probability that the cab involved in the accident was blue, after learning what the witness said?
5. Suppose that, as originally, the witness correctly identifies each one of the two colors 80% of the time and failed 20% of the time. However, you don't know what percent of the cabs in the city are Green and what percent are blue. What percentage of cabs have to be green so that, after you learn that the witness said the car was blue, you think there is a 50% probability that the car was blue?

3 | Challenge Exercises

1. Suppose the test for some disease is perfect for people who have the disease, $Pr(\text{positive} \mid \text{sick}) = 1$. And it's almost perfect for people who don't have the disease: $Pr(\text{positive} \mid \neg\text{sick}) = 1/99$. How high does the base rate $Pr(\text{sick})$ have to be for the test to be 99% reliable, i.e. to have $Pr(\text{sick} \mid \text{positive}) = 99/100$?[1]
2. I have an even mix of black and white marbles. I pick one at random and put it in a bag. Then I put a white marble in the bag, shake the bag, and take out a marble at random. If the marble I took out is white, what is the probability that a second marble I take out is white? [2]
3. Without putting the white marble I took out back in the first bag, I pull out another bag and put the white marble in along with two black marbles. Now, what gives you a better chance of drawing a white marble?[3]
   (a) Randomly choose one of the bags, and then draw a marble from that bag.
   (b) Pour the marbles into the same bag, and draw a marble from that bag.
4. Bayes Theorem says that to find the probability of a hypothesis H *given* some evidence E, take the previous probability of H, $Pr(H)$, and multiply it by $\frac{Pr(E|H)}{Pr(E)}$. Suppose that the only thing you know is the value of $\frac{Pr(E|H)}{Pr(E)}$. How would you figure out if learning E should make you more confident or less confident in your hypothesis H?
5. Bayes' Theorem can be proved in just a few steps from the ratio formula for conditional probability, applied to $Pr(H \mid E)$ and $Pr(E \mid H)$:

$$Pr(H \mid E) = \frac{Pr(H\&E)}{Pr(E)}, \qquad Pr(E \mid H) = \frac{Pr(H\&E)}{Pr(H)}.$$

   Complete the proof.
6. Using the ratio formula (see the exercise right above), the law of total probability can be rewritten to not have any conditional probabilities. How?

1. Weisberg, *Odds and Ends*, Ch 8
2. This problem is from Lewis Carroll, "Pillow Problems" (1958). The answer is *not* 1/2.
3. Also from Carroll (1958).