Probability, Statistics, and Bayes' Theorem
Session 3

1 Introduction

Now that we know what Bayes' Theorem is, we want to explore some of the ways that it can be used in real-life situations. Often the results are surprising and seem to contradict common sense. Before we turn to these, we'll have a quick review of what Bayes' Theorem says.

1.1 Bayes' Theorem: Review

Recall the definition of the conditional probability of A given B:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}.$$

On a formal level, Bayes' Theorem relates two conditional probabilities:

$$P(A|B) = P(B|A)\frac{P(A)}{P(B)}.$$

Let the event A be interpreted as a proposition and let the event B be interpreted as evidence. Then $P(A)$ is our initial belief in the truth of proposition A. This is called the *prior* probability for A, or just the prior. $P(A|B)$ represents our belief that A is true taking into account the evidence B. It is called the *posterior* probability of A given B, or just the posterior. The ratio

$$\frac{P(B|A)}{P(B)}$$

is interpreted as the support that the evidence B provides for A. With these interpretations, Bayes' Theorem gives us a way to update of belief in the truth of proposition A based upon the available evidence B; we start with the prior belief in A, then we gather evidence B, and then we form the posterior belief $P(A|B)$.

2 Medical Testing

Bayes' Theorem can be effectively used in assessing the results of any kind of test where there is not complete accuracy in the outcome of the test. These kinds of tests are often given in a medical context. Before we turn to the problem, we have to formally define what we mean when we say that the outcome of a test isn't completely accurate.

2.1 False Positives and False Negatives

We are going to assume that the test has only two outcomes, which will be called "Positive" and "Negative," which we will sometimes just shorten to P and N. We are also going to assume that there are only two possible states that we can be in; we either have what the test is checking for, or we don't. We will call these states "Sick," if we have it, and "Healthy" if we don't. Again, we will often shorten these to S and H. There are now four possible combination of the two states and the two outcomes:

$$
\begin{aligned}
&(N|H) : \text{The test comes back negative and we do not have the disease}\\
&(P|S) : \text{The test comes back positive and we do have the disease}\\
&(P|H) : \text{The test says that we have the disease but we don't}\\
&(N|S) : \text{The test says that we don't have the disease but we do}
\end{aligned}
$$

In the first two cases, the test is doing what it is supposed to do: if we're healthy, the test says that we are healthy, while if we are sick then the test confirms this. In the second two cases, however, the test is not working, but it's not working in different ways. Given (H,P), we are healthy but the test is coming back positive. This means that we think that we have the disease even though we don't. This has widespread ramifications, emotionally, physically, and financially. When this happens, we say that the test has produced a *false positive*, since the outcome of the test is positive, but that is not the state that we are in. Similarly, given (S,N), we are sick but the test says that we are healthy. This means that we don't think that we have the disease even though we do. This has even more dangerous ramifications because we might fail to get the treatment that we need to regain our health because we don't think that we are sick. When this happens, we say that the test has produced a *false negative*, because the outcome of the test is negative but that id not the state that we are in. To keep terminology consistent, we refer to (H,N) as a *true negative* and (S,P) as a *true positive*. This can all be summarized as

$$
\begin{aligned}
&(N|H) : \text{True Negative}\\
&(P|S) : \text{True Positive}\\
&(P|H) : \text{False Positive}\\
&(N|S) : \text{False Negative}
\end{aligned}
$$

These values reflect the accuracy of the test, and they need to be taken into account before any decisions are made based on the outcome of the test.

2.2 The Situation

Suppose that you go to the doctor to have a check-up. You feel perfectly healthy and look perfectly healthy. When you get to the doctor's office, he tells you that there is a pretty serious disease going around, and it seems like 1 in 5000 people in the Waterloo area have contracted the disease. A reasonab;e description of your knowledge would be the prior probabilities

$$P(\text{Healthy}) = .9998, \quad P(\text{Sick}) = .0002$$

He then tells you that he has a test for the disease, and that the test has about 5% false positives and about 2% false negatives, so that

$$
\begin{aligned}
&P(\text{Positive}|\text{Healthy}) = .05\\
&P(\text{Negative}|\text{Healthy}) = .95\\
&P(\text{Positive}|\text{Sick}) = .98\\
&P(\text{Negative}|\text{Sick}) = .02
\end{aligned}
$$

You decide to take the test. A few days later, the doctor calls you with your results.

2.2.1 Exercises

1. How would you feel if the doctor told you the test came back negative?

2. How would you feel if the test came back positive?

3 The Monty Hall Problem

This is a famous application of Bayes' Theorem to a probabilistic problem which explains the counterintuitive results.

3.1 The Problem

Suppose that you are a contestant on a game show. There are three doors, which are closed, and which are number 1 to 3. Behind one of doors is a brand new sports car. Behind the other two are goats. You have one chance to pick a door, and if you pick the door with the sports car behind it, you win the sports car. If you pick a door with a goat, you win the goat. Say that you pick door 1, and announce it to the host. Before he opens door 1, he opens door 3 and reveals a goat. The host turns to you and asks you if you want to change your choice to door 2.

3.1.1 Exercises

1. Is it to your advantage to switch?

3.2 Discussion and another Problem

Before the host opens door 3, you figure that you have a 1/3 chance of winning the sports car. When he opens door 3 and shows you a goat, you now know that the sports car is either behind door 1 or door 2. You might think to yourself that since you now know that there are only two possibilities (door 1 and door 2), that the probability of it being behind door 1 is now 1/2 and the probability that it is behind door 2 is 1/2, so there is no advantage in switching. In essence, once the host opens door 3, you think to yourself that the sample space has shrunk from three options to just two, and since the information provided by opening door 3 doesn't change the location of the car, each option is equally likely.

You're standing there thinking and thinking. The host can tell that you're working through the various possibilities, and he asks you to share your thoughts with him. You explain that you're not going to switch because you figure that the probabilities of being behind door 1 and door 2 are both 1/2 and you'll stick with your gut instinct which told you to pick door 1. The host listens to you, nodding as you explain your reasoning. When you're finished, he says "What if I told you that the probability of the car being behind door 2 is 2/3?"

3.2.1 Exercises

1. Do you believe him?

2. Would you switch now?

3.3 The Mathematical Model

Let's turn this into a probabilistic model that we can investigate rigorously. Our sample space is $\Omega = \{1, 2, 3\}$, these numbers representing the respective doors. There are three random variables involved:

C, the number of the door with the car behind it
S, the number of the door selected by you
H, the number of the door opened by the host

At the start of the game, $P(C = c) = 1/3$ for all $c \in \Omega$ since we are assuming that the car is equally likely to be behind any door. We also know that $P(C = c|S = s) = P(C = c)$ since our choice doesn't affect the location of the car in any way. Finally, we have the following

$$P(H = h|C = c, S = s) = \begin{cases} 0 & h = s \\ 0 & h = c \\ 1/2 & h \neq s, s = c \\ 1 & h \neq c, h \neq s, s \neq c \end{cases}$$

These probabilities are all that we have to work with. It turns out that they are all the we need.

3.3.1 Exercises

1. Explain the values of $P(H = h|C = c, S = s)$.
2. Use Bayes' Theorem to get an expression for $P(C = c|H = h, S = s)$ in terms of the probabilities that we have above.
3. Compute $P(C = 2|H = 3, S = 1)$.

4 The Prosecutor's Fallacy

This is an example from a legal setting. It is a fictional example based upon real trials that have affected real lives.

4.1 On Trial

A person is accused of committing a pretty serious crime. They maintain that they are innocent, the victim of bad luck. There is some evidence that seems to link the defendant to the crime, but as with most evidence it is not absolutely conclusive. During the trial, the prosecutor calls to the witness stand an expert in interpreting exactly this kind of evidence. At one point, the prosecutor asks the expert "What is the probability of finding this evidence, if the defendant was innocent?" The expert replies "1 in 73 million." Later in the trial, during the final argument, the prosecutor says to the jury "You've heard an expert witness testify that the probability that the defendant is innocent, given this evidence, is 1 in 73 million."

4.1.1 Exercises

1. If you were the defense attorney, would you object to the prosecutor's last statement?

4.2 Mathematical Formulation

The prosecutor's fallacy arises from the confusion of two different, though related, conditional probabilities. Let E represent the evidence, and I represent the innocence of the defendant. The expert witness testified to the probability of finding the evidence, assuming innocence of the defendant. This can be represented in terms of conditional probabilities as $P(E|I)$. What the prosecutor said to the jury involved the probability of the defendant being innocent given the evidence. This can be represented in terms of conditional probabilities as $P(I|E)$. The fallacy in the prosecutor's statement is that $P(E|I) \approx P(I|E)$. But we know from Bayes' Theorem that we can't just switch the order of terms in conditional probabilities. The correct relationship between the two probabilities is instead given by

$$P(I|E) = P(E|I)\frac{P(I)}{P(E|I)P(I) + P(E|I^c)(1 - P(I))}$$

Here $P(I)$ is the prior probability of innocence independent of all evidence, and I^c represents the defendant being guilty.

4.2.1 Exercises

1. What would happen if the evidence were truly a "smoking gun," meaning that $P(E|I) = 0$?

2. What would happen in the opposite case that the evidence was absolutely untrustworthy, meaning that $P(E|I) = 1$?

3. What would happen if the evidence didn't really provide any additional insight, meaning that $P(E|I) = 1/2$?