You Know I’m All About that Bayes: Crash Course Statistics #24

Crash Course: Statistics
https://youtube.com/watch?v=9TDjifpGj-k
https://nerdfighteria.info/v/9TDjifpGj-k

Hi, I’m Adriene Hill, and Welcome back to Crash Course, Statistics.

We all have ideas about how the world works. And even if we haven’t ever used numbers to describe them, we see different beliefs everywhere we go.

From whether it’s healthier to be a vegetarian... to whether school uniforms are a good idea; we all have slightly different “models” of how the world works. And yet we all agree on a lot. For example, we all believe the sun will come up tomorrow morning--or at least that it’ll come up eventually if you live at the poles.

And we all believe that the Oxygen atoms in the room won’t all suddenly move to one corner of the room, leaving us to suffocate. And almost all our beliefs change based on our experience. That’s why your friend with a snake that he cuddles with all the time isn’t as afraid of snakes as you are...since your only exposure to snakes is that one time when you were hiking and a rattlesnake almost bit your dog.

Our beliefs are numerous, sometimes complex, and consistently changing. So it can be useful to have a way of doing statistical inference that reflects that. INTRO Bayes’ Theorem--or Bayes Rule--tells us the that probability of A given B, is the probability of B given A times the probability of A, all divided by the probability of B.

And remember that the numerator in this equation is just another--way of writing the probability of A and B. For example, When you’re out to lunch, your sister mentions that she have a friend who has breast cancer, but doesn’t say much else. You recently saw a documentary about males with breast cancer.

Because it’s so fresh in your mind, you wonder if your sister’s friend is a male. Your gut feeling is that it’s not that likely they’re male... but let’s quantify that. You want to know the probability that your sister’s friend is male, given that you know that friend has breast cancer.

Using Bayes’ theorem, we can calculate this probability. The probability of being male, given that you have breast cancer is equal to the probability of having breast cancer given that you are male times the probability of being male, divided by the probability of having breast cancer. Thanks to government health agencies, we know many of these statistics.

The probability of getting breast cancer given that you’re male is 0.001, and we will assume the probability of being male is 0.5. The overall probability of getting breast cancer is 0.063. Armed with your facts you calculate that the probability that your sister’s friend is male is only about 0.79%...so...not very likely.

But maybe more likely than you would have anticipated. If we rearrange Bayes’ Theorem slightly you can see that it allows you to update your beliefs based on new information: When we used Bayes’ Theorem what we were really doing was updating our belief that a person was male (probably about 50/50 odds if you know nothing else about them) with the new information that they had breast cancer. This new information changed our belief; we went from a 50 percent chance to about a 0.79% chance just by taking into account this new information.

This idea about updating beliefs is core to Bayesian statistics and can be used to test hypotheses. We start with some idea or belief about how something works. For example, you set your friend Maria up on a blind date.

Maria’s excited, but nervous, and on her way to the coffee shop to meet her blind date Jordan, she wonders whether he shares her love of Star Wars. From her experience meeting people in the city, she believes that in general, there are slightly more Star Wars fans than non fans. She guesses there’s a 60% chance that a given person is a Star Wars fan, and a 40% chance they are not, which means that she thinks it’s 1.5 times more likely that someone is a fan.

When Maria arrives at the coffee shop she and Jordan do the normal first date small talk. He asks her what she did this last weekend, and she told him that she saw the new Star Wars movie. Jordan says he did too!

After hearing this, Maria feels like it’s more likely that she might have met her Porg loving soulmate. She knows that not everyone who’s seen Star Wars is a fan, but she can use the fact that Jordan has seen it to update her belief about whether or not he is one… or she could ask! Maria knows that the probability of having seen the last Star Wars movie given that you’re a fan is 0.99, since pretty much all the fans rushed to see the movie.

But not everyone who went to see the movie were fans. Some were just curious, and others were dragged by family or friends to see it. She thinks that the approximate probability of having seen the movie given that you’re not a fan is 0.5, since some but not all non-fans went to see it.

Maria can use the ratio of these two probabilities: To see which hypothesis is more probable given that we know Jordan saw the movie. Based on Maria’s quick calculations, this new information means that it’s now 1.98 times more likely that Jordan is a Star Wars fan than not. Her heart starts beating a little faster!

This ratio of the probability of our information under one hypothesis--that he’s a fan--compared to another--that he’s not a fan--is called a Bayes’ Factor. It represents the amount of information that we’ve learned about our hypotheses from the data. Maria can use it to update her previous belief--or prior odds--that it’s 1.5x more likely that Jordan is a fellow Star Wars fan.

All she has to do is multiply her prior beliefs--the one’s she held before she had any new information--by the Bayes Factor which tells her how much to change her belief, now that she has gotten some evidence. The resulting belief is called her posterior belief in this case 2.97. And she can continue to incorporate new information.

When Jordan says that his dog is named Anakin, she can again update her beliefs. Or just ask! Mathematically, we took Maria’s prior belief: And updated it with our Bayes Factor, which told us how much our data--Jordan seeing the new Star Wars movie--should change her beliefs about his fan-hood.

This is a very simple example of how we can use Bayesian Hypothesis Testing to compare the probabilities of different hypotheses based on data that we observe. But this doesn’t look exactly like the Bayes’ Theorem that we saw at the beginning... That's because instead of looking at the probability of one hypothesis given the data, we’re looking at the ratio of two hypotheses.

Instead of just calculating the probability that Jordan was a Star wars fan, given that he’d seen the latest film, we compared the probabilities of the two hypotheses, given that he’d seen the movie. So we’re really looking at the ratio of two calculations of Bayes’ Theorem because we’re comparing two posterior probabilities. Luckily, the probability of having seen the latest Star Wars movie is the same in both equations, so it cancels out and we end up with this: In Bayesian Statistics, these things are called: the Prior--what you believed before you saw any evidence the likelihood--a measure of how much your evidence should change your prior

beliefs And the Posterior--what you believe after you've seen the evidence In a more general form we can say that after we see the data, how likely one hypothesis is compared to the other is equal to the ratio of how likely we thought these hypotheses were before we got any evidence, adjusted by the evidence with which the data provided us.

This reflects the core idea of Bayesian Hypothesis testing: Updating what you currently believe, with new information. But notice that I said that you update your belief. Inherently, what we believe is subjective.

It depends on who we are, and what we've experienced. While Maria initially believed that a star wars fan is 1.5 times more likely than an non fan, you may believe something else. Like that it's just as likely that someone is a fan and not a fan.

And since this is just a personal belief, it's okay that you and Maria believe something different to begin with. But we used Maria's prior beliefs in our calculations, if you were to do the same calculations, you'd come up with a different number. And this is one criticism that Bayesian statistical inference faces.

One of the main uses of statistics is science which is supposed to be relatively "objective" and not influenced by opinion, and yet, here's a method that includes beliefs in its calculation. For example, say a scientist bases her conclusion that Extra Sensory Perception (ESP) exists on the posterior odds of her Bayesian calculation. She concludes from her study that it is 5 times more likely that ESP exists than doesn't exist.

But upon reading her paper, you find that her prior beliefs about the probability of ESP were way higher were way higher than yours. She assumed that it was just as likely that ESP exists as it is that ESP doesn't exist. And that just doesn't seem right to you.

You could find another scientist who has about the same prior beliefs about ESP as you do... but that seems difficult and a bit inefficient… There is a better solution. Often studies that use Bayesian calculations will not just report their posterior odds, but also the Bayes Factor that they calculated. If you disagreed with a researcher's prior odds, you could use the reported Bayes Factor to adjust your own, different beliefs about these two specific hypotheses.

For example, if you believed that it was 1,000 times more likely that ESP doesn't exist, you could use the researcher's reported Bayes Factor, 5, and adjust your own beliefs: Even though the evidence in the study makes you believe that ESP is more likely than before, you still think it's relatively unlikely that ESP exists. Even though Bayesian Hypothesis Testing includes subjective beliefs, the Bayes factor allows you, and anyone else, to use the evidence from a study or analysis to update whatever your prior beliefs about the two specified hypotheses! Just like you and the ESP favoring researcher, sometimes evidence can lead two people to very different conclusions.

But often, unless someone has already decided something has a 0% probability, when there's sufficient evidence, two people with different prior odds will come to the same conclusion. For example, you initially believe that sushi is pretty dangerous and has a high risk of infecting you with parasites, and your co-worker thinks that the risk is low to moderate. If you both see your boss and entire team go out to sushi every week for 2 years and not have any issues with parasites, both of you, despite your initial differences, would probably have updated your beliefs with this new information and concluded that sushi is pretty safe after all.

If you had only seen your boss go out to sushi 4 times without getting a parasite, you may have each come to different conclusions since that's not as much evidence. You may still think sushi is pretty risky, but that may have been enough evidence to convince your co-worker it's safe. Your current beliefs would rely more on your prior beliefs than the new evidence.

But the huge amount of evidence provided by a group of healthy co-workers over 2 years was enough to overwhelm your and your coworker's prior beliefs. So your new, posterior beliefs are more affected by the evidence than your prior beliefs. Bayesian hypothesis testing provides a structured way to quantify a logical process that we do every day, incorporating new events into the way that we see the world.

It provides an explanation...or at least a hypothesis--about why two people can see the same evidence and reach different conclusions. In some situations, the logic of Bayesian methods similar to how we think naturally. Like a doctor who uses patient symptoms like fever and fatigue to update the prior odds that a patient has the flu compared to a cold so that they can prescribe the correct treatment.

Or the way that you updated your belief that your best friend is a kind, caring person by continuously incorporating evidence of their kindness--like covering the cost of your Starbucks when you lost your wallet, or helping you move. In real life you don't ignore all previous pieces of evidence you saw as soon as you get a new one, and Bayesian Inference allows for you to take your new updated beliefs and update them again. As some Bayesians say, "yesterday's posterior (your updated belief), is today's prior (the beliefs to be updated)" Thanks for Watching, I'll see you next time.