Bayes in science and everyday life: Crash Course Statistics #25

Crash Course: Statistics
https://youtube.com/watch?v=51bLRF02b4w
https://nerdfighteria.info/v/51bLRF02b4w

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

Bayesian Hypothesis Testing--or Bayesian Inference--is a great way to model the way we reason about things in everyday life. We collect evidence and experience and we use it to build our beliefs about the world.

We collect information on whether certain facial expressions mean that someone is upset. Whether clouds outside mean it's going to be cold today. Or whether people who smoke are more likely to have lung cancer.

But Bayesian methods are useful above and beyond updating our personal beliefs. Bayes has helped companies make marketing decisions--like which color to use on their website. And it has helped researchers to quantify their results in scientific studies.

And today, we're going to talk about it. INTRO First, you may have noticed that so far when we talk about the math of Bayes' Theorem, we've been using discrete variables, like whether or not you're a star wars fan, or whether or not you have a disease. But Bayes' Theorem can help us update beliefs that involve continuous variables too.

The math of Bayes' Theorem with a continuous variable is a bit more complicated than in the discrete case. Science writer Sharon Bertsch McGrayne called it, "a Theorem in want of a computer...". In fact for much of the 20th century, Scientists and Statisticians who wanted to use Bayes were limited in their ability to do so by a lack of computational power.

But we still want answers to those more complicated problems. Sometimes we want to know whether dogs who are walked regularly are less likely to damage furniture. Or whether House Elves have lower intelligence than Wizards...which is an example of Bayesian Hypothesis testing from a Harry Potter themed article by Alexander Etz and Joachim Vandekerckhove.

Guess you'd better update your prior belief about how cool statisticians areðŸ˜Ž. The ideas behind continuous and discrete Bayesian Inference are exactly the same. We take our prior beliefs--what we believe before we've seen new evidence--and update it with the likelihood of our evidence.

This is called the Bayes Factor when comparing two models. Once we've updated, our new beliefs are called our "Posterior" beliefs. If we're comparing two models, they are called our posterior odds.

But instead of simple, discrete probabilities, we have probability distributions. For example, let's look at the ever present problem of whether or not a coin is biased. Before you start your experiment to test the fairness of your coin, you decide that you know almost nothing about whether or not it's biased.

So your prior probability of getting tails is a uniform distribution between 0 (never tails) and 1(always tails). You consider all probabilities of getting tails--we'll call that theta--equally likely. You have a friend flip the coin in question 5 times, and they get 1 tail.

Which seems unlikely, though not impossible for a fair coin. Using the Binomial Probability Formula we know that the probability of this happening with a fair coin is about 16%. Note this new notation for 5 choose 1, you're most likely to run into this in the stats world.

So how does this evidence update your belief about what the real probability of getting a tail is for this coin? Before we show the Bayesian calculation, let's take a moment to figure out what we think without the math. Since we saw at least one head and one tails, we can rule out both the probabilities 0 and 1.

And we think that probabilities very close to 0 and 1 are unlikely too. Because it'd be REALLY rare to see only one tail if the probability of tails were 0.99. And similarly rare to see a tail at all if the probability were 0.001.

Now we can do the Bayesian calculation and see if it matches our intuition. Here's Bayes' Theorem, but for this continuous problem: We won't get too stuck on the math, but we can see that this is the same old Bayes' Theorem that we've seen before...just continuous. When we plug in this formula to a graphing program to show our posterior, it looks like this: The Y axis tells us the relative probability of a theta--in this case theta is the probability of getting tails-- and the x axis shows us all the possible values of theta between 0 and 1.

We can see that we took our prior distribution (the dotted line)... and updated it using the likelihood of the data, which told us the probability of getting 1 out of 5 tails for EVERY potential probability of getting tails that a coin could have. Once we updated our prior beliefs, about which probabilities are the most likely, our posterior beliefs are represented like this (the solid line). Anything on the curve that is above the dotted prior line represents a theta that became more likely after we saw the data.

And anything on the curve that is below the dotted line is a theta that became less likely. And this matches our intuition; Thetas that are close to 1 and 0 became less likely, while thetas around 0.1-0.5 became more likely. So maybe we have a fair coin here...but it seems more likely that it's biased.

Businesses like Bayes because it allows them to take into account previous knowledge and expert opinion when they make their calculations. Let's look at an example of how a business might use Bayesian inference. We'll keep the math to a minimum, but if you're interested in learning more, you can check out this awesome blog post by Will Kurt on countbayesie.com which we based this next example on.

And the link is in the description. Say you're a beauty blogger, and you send out weekly emails encouraging your followers to read your latest blog post. The more people who click, the more money you make, and so you want the most clickable emails ever.

Your friend, who's also in the blogging business, told you that adding a picture at the top of your email gets more people to click, but you want to test that idea out with your own readers. Normally, your click rate is around 30%, so you decide to represent your prior beliefs about your true click rate using this function: Values around 30% are most likely, but it's possible your true click rates are higher or lower than that. You randomly select 300 of your followers to be a part of your experiment--often called an A/B test in the business world--and send half the email with a cute picture of you with your poodle, Ginger as well as the normal content.

The other half gets your standard picture-less email. You anxiously await the results, and three days later you have them: You use the new information you have about your two emails to update your original beliefs about your click rate. Since the two groups were the same before you assigned them to get either email No Dog Pictures or with Dog Pictures, you use the same prior for both groups.

Once you've incorporated this new evidence, your Posterior distributions look like this: And they tell you how likely each click rate is under your new, posterior beliefs about each group. It looks like the group with pictures is likely to have a higher click rate... but you can't know for sure. One way to get more information to make your decision is to randomly simulate a bunch of samples - one at a time.

The samples come from each of your two posterior distributions and then you count how often the group with pictures' click rate is higher than the group that didn't get a picture in their email. That percentage will tell you roughly how likely it is that the group that got pictures will have a higher click rate than the group who did not. You decide that if in 70% of your simulation samples the group with pictures has a higher click rate, you'll include glamour shots of Ginger in all your new emails.

Using Bayesian methods to analyze this question allowed you to "inject" your own prior beliefs into the analysis, which is important when making business decisions. Businesses often want to make the best decision in the most cost efficient way, which means taking advantage of all the information you have; not only data, but prior knowledge of the field and expert opinion. Your prior knowledge about the click rate of your emails made it possible for you to start your analysis knowing it's pretty unlikely your click rate was very near 0, or very near 1.

Bayesian analyses can be incredibly useful in science, as well. A study on Dissociative Identity Disorder (or DID)--formerly called Multiple Personality Disorder--looked at whether people with D-I-D had different "memory" between personalities. If one person had two separate personalities, Bob and Alice, researchers were interested in whether something that person learned as "Bob" could be remembered by that person when they were "Alice".

In order to test this idea, participants were shown a few pictures and told a story. They then waited a little while, and answered 15 multiple choice questions about the material. There were 3 different groups of participants: A group of DID patients - who were asked to learn the materials in one personality and switch to another personality before the test.

A pretend amnesiacs group - without DID who did not see the materials. And a malingers group without DID who saw the materials but were told to pretend they hadn't and answer as if they had never heard the story or seen the pictures. Researchers wanted to know whether the patients with DID, the people who had never seen the materials, and the people who were pretending not to have seen the materials had the same mean accuracy on the test.

This would help researchers and cognitive scientists understand more about how memory works in DID patients. Using Null Hypothesis Significance testing, researchers could try to address whether all three groups had the same mean score on the test, but even if they rejected the null hypothesis that all three groups are the same, they wouldn't be able to say how much more likely it was that all three groups were different. Bayesian methods can tell you that.

And a Group of researchers did analyze the data this way, and found out that the Bayes Factor for these models was about 4,000! That means that the data that the researchers saw should update our beliefs by a lot. No matter what you believed before hand, your updated beliefs will most likely reflect the fact that it's more likely that these three groups--DID patients, people who didn't see the materials, and people who pretended not to see the materials--are three distinct groups.

And it's interesting, because it provides evidence that people with DID may not just be pretending to not remember things that were learned while they were in a different personality... but they may not quite be behaving the same as people who really had never seen the materials, which is what you might expect if two personalities were completely separate. And while Bayesian inference is increasingly popular in many scientific fields like Psychology, it's also being used right now in many places near you. Bayesian methods are used to help translate one language to another, and to suggest which items you might buy next based on the fact that you just bought four silicone sponges, a Sandalwood Candle, and whiteboard markers.

Bayes can help figure out which allergy medicine you'll react best to based on your genetic profile. And Bayes plays a role in creating artificial intelligence that can do pretty amazing things, like understanding that it's more likely that you said "Siri, Turn on the lights" and not "Siri, Learn all the Sites !" Thanks for watching, I'll see you next time.