Bayes' theorem

Vincent Thorne

December 5, 2019

Intuition

One main intuition for the Bayes' theorem is: "How much can I trust a particular test?". In other words, how much can I rely on the evidence that is given to me by a test to assess the probability of something happening. Bayes' theorem tries to makes sense of the error that a test contains.

Elements of a test

event the thing that we actually want to measure

test the test that *imprecisely* (i.e. with some error) indicates if the event occurred.

true positives the cases where the test correctly indicated that the event happened

false positive the cases where the test showed that the event occurred when it actually didn't

true negatives the test didn't indicate that the event was happening when it was indeed not happening

false negatives the test didn't indicate that the event was happening when it was actually happening

Examples of test

- A fire alarm (test) reports if a fire is raging (event) in a building. If the fire is indeed raging *and* the alarm sounds, that's a true positive.
- A lie detector (test) tells you if someone is lying (event). If the detector bips but the person is *not* lying, that's a false positive.
- A bomb detector (test) rings if it detected a bomb in someones bag (event). If the alarm doesn't ring but there actually was a bomb in the bag, that's a false negative

- A road radar (test) flashes flashes if it detects a car driving faster than the speed limit (event). If it didn't flash when the car was going below the speed limit, then that's a true negative

What does Bayes' theorem do

Each of these examples illustrates one outcome of the test: true/false positive and true/false negatives. Bayes' theorem tell us "What is the probability of being a true positive given all the positives" (i.e. false + negative positives). In other words, out of all the times an alarm rings, what's the probability that it rings because the event is actually occurring. Here lies the essence of Bayes.

We talk about *prior* and *posterior* probabilities (or odds when not expressed in percentages) to indicate that we filter the event through a test, and we want to know what is the chance that the event occurred given the result of the test. Namely, we *update* our probabilities after making the event go through a test.

The cancer test example given in class is a nice illustration. Before the test, we know that cancer hits 1% of the population. Moreover, we are aware of the performance of our test: we know that the test correctly identifies cancer with a 90% success rate, and does not detect cancer *when there is no cancer* with 90% accuracy (i.e. it will "bip" 10% of the time if the person doesn't have cancer).

Bayes theorem tells us the chances we have of actually having cancer if we got a positive test result. In this case, it's 8.3% (try to work out the math on that one), which means there is a 8.3% chance that we have cancer if the test was positive. That's pretty low "efficiency score" for a test, but still, we know with greater certainty whether we have cancer or not after the test: we've updated our probabilities thanks to the test.

Math

If $P(B)$ is the probability of a fire (thus $P(B^C)$ is the probability of no fire), $P(A)$ the probability that the alarm rings, $P(A \mid B)$ the probability that the alarms rings when there is a fire (true positive), and $P(A \mid B^C)$ the probability that the alarm rings when there is no fire.

Bayes' theorem tells us what is the probability that there is a fire when the alarm rings, that is $P(B \mid A)$. The formula describing this relation is

$$P(B \mid A) = \frac{P(A \cap B)}{P(B)} \tag{1}$$

which in plain English is

$$\frac{TruePositives}{AllPositives} = \frac{AlarmsRingsWhenFire}{AllTimesAlarmRings}. \tag{2}$$

Using some probability rules, we can change equation 1 to

$$P(B \mid A) = \frac{P(A \mid B) \times P(B)}{P(A \mid B) \times P(B) + P(A \mid B^C) \times P(B^C)} \tag{3}$$

which in words translates to

$$\frac{ProbOfAlarmIfFire \times ProbOfFire}{ProbOfAlarmIfFire \times ProbOfFire + ProbOfAlarmIfNoFire \times ProbOfNoFire}. \tag{4}$$

Multiplication rule

I showed in class the multiplication rule, which is

$$P(A \cap B) = P(A \mid B) \times P(B) = P(B \mid A) \times P(A) \tag{5}$$

if and only if A and B are not independent, that is if they share an area on the Venn diagram (you can also think about it as some kind of correlation $\neq 0$).

On the board I wrote that

$$P(B \mid A) = \frac{P(A \cap B)}{P(B)} = \frac{P(A \mid B) \times P(B)}{P(B)} = \frac{P(B \mid A) \times P(A)}{P(B)} \tag{6}$$

and Theo pointed out that we have on the nominator what we want to calculate. But I failed to notice that on the nominator we *do not* have $P(B \mid A)$ but actually $P(B \mid A) \times P(A)$ which in plain words is $ProbOfFireWhenAlarmRings \times ProbOfAlarmRinging$. That element doesn't make much sense from what we now about the test and the probabilities we can gather. It would imply using something that we don't yet know ($P(B \mid A)$) to compute itself. Thus, equation 6 is just a mathematical equivalence which is correct but doesn't help us solve the Bayesian equation.