

Bayes Theorem

Contents

- 4.1. A test result
- 4.2. Bayes Theorem

We'll begin by challenging our common intuition about the concept of “accuracy” specifically when accuracy is applied to a diagnostic test.

4.1. A test result

Carl goes to the doctor and takes a test to see if he has an infection. The general incidence rate of the infection in the local population is 1%. The test he is taking to determine the presence of the infection has an accuracy of 97%. If Carl tests positive, how concerned should he be that he has the infection?

To measure Carl's “worry coefficient” we want to know $P(I|+)$, or, the probability that he has the infection, given that he had a positive test result. We know $P(+|I) = 0.97$, because the test is 97% accurate so if he has it, the test will be positive 97% of the time.

Logically the ratio needed is the number of people who have a positive test result and are infected, divided by the number of people who have a positive test result total. Out of 1000 people, 10 of them is infected. 97% of the 10 will get a positive test result, and 3% of the 990 uninfected will get a positive test result. This says

$$\#\{\text{positive test cases}\} = (0.97)(10) + (0.03)(990) = 9.7 + 29.7 = 39.4$$

and the number of infected with a positive test result is

$$\#\{\text{true positive test results}\} = (10)(0.97) = 9.7$$

Giving a worry coefficient of $9.7/39.4 = 0.246$, or a 24.6% chance of being infected given a positive test result.

So what can be made of the “97%” accuracy claim? Well it sounds great, but it's wrong 3% of the time. And only 1% of the population has an infection. So the size of the error is actually larger than the number of infected people. High accuracy on rare events can be misleading and requires you to know more context to properly interpret. This is an early example of a critical lesson in machine learning: **accuracy alone is not an informative, and is occasionally a deceptive, metric.**

Before we delve too deeply into variations on accuracy, let's pause to look at Bayes' Theorem, which tells us how to compute $P(I|+)$ if we know $P(+|I)$.

4.2. Bayes Theorem

Bayes' Theorem gives us a way to *invert* conditional probabilities. The formula comes from the definition of conditional probability

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

this implies the following

$$P(A \cap B) = P(A|B)P(B) = P(B|A)P(A)$$

Solving for $P(A|B)$ we get

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

Though this is the final form, in practice you will need to compute $P(B)$ using the following

$$P(B) = P(B|A)P(A) + P(B|\bar{A})P(\bar{A})$$

which says the probability of B is the sum of the probability of B given A and B given not A . (A is either true or false so these are the only two options)

Applied to the above example

$$\begin{aligned} P(I|+) &= \frac{P(+|I)P(I)}{P(+)} \\ &= \frac{(0.97)(0.01)}{P(+|I)P(I) + P(+|\bar{I})P(\bar{I})} \\ &= \frac{(0.97)(0.01)}{(0.97)(0.01) + (0.03)(0.99)} \\ &= 0.2462 \end{aligned}$$

In the accompanying lab you will explore Bayes' Theorem and the ways in which the final result depends on both the accuracy of the test and the incidence rate of the infection.