Science as a Bayesian algorithm

There is a new way to think about science and its power and limits, which is gaining favor among scientists and philosophers alike, and which I think beautifully clears up many common misunderstandings that both scientists and the general public seem to have about the nature of science. I am referring to an old and until recently rather obscure way to think about probability invented by Reverend Thomas Bayes back in 1763.

Bayes realized that when we think of the probability of a hypothesis to be true we base our judgment on our previous knowledge about the phenomenon under study (i.e., we use induction). We then assess new information in the light of this prior probability and modify our belief (meant as degree of confidence, not as blind faith) in the hypothesis based on the new information. This process can be repeated indefinitely, so that the degree of trust we have in any hypothesis is always due to the current (and ever changing) balance between what we knew before and the new knowledge that additional data bring in.

Here is the fundamental equation of Bayesian statistics, known as Bayes rule:

$$P(H|D)=\frac{P(D|H)*P(H)}{P(D|H)*P(H)+P(D|\sim H)*P(\sim H)}$$

Where P(H|D) (which reads "the probability of H given D") is the probability that our hypothesis is correct given the available data; P(D|H) is the probability that the data would be observed given the hypothesis; P(H) is the unconditional probability of the hypothesis (i.e., its probability before we knew of the new data); and the denominator is a product of the numerator plus an equivalent term which includes the probability to observe the data if the hypothesis is actually wrong (or, in the case of multiple hypotheses, the probabilities of observing the data if each additional hypothesis is correct). The denominator of the right hand of the equation is also known as the *likelihood* of all hypotheses being considered. The left hand of the equation is called the *posterior probability* of the hypothesis in question; the left part of the numerator on the right side of the equation is known as the *conditional likelihood* of the hypothesis in question; and the right part of the same numerator is called the *prior probability* of the hypothesis being considered. This sounds very complicated until we examine a particular example, so bear with me for a few more minutes.

A family has plans to go fishing on a Sunday afternoon, but their plans are dependent on the weather at noon on Sunday: if it is sunny, then there is a 90 % chance that they will go fishing; if it is cloudy, then the probability that they will go fishing drops to 50 %; and if it is raining, the chances drop to 15 %. The weather prediction at the point we first consider the situation call for a 10 % chance of rain, a 25 % chance of clouds, and a 65 % chance of sunshine. The question is: given that we know that the family eventually did go fishing, was the weather sunny, cloudy, or rainy? You will probably have your intuitions about this, and they may well be correct. But science goes beyond intuition to empirically-based reasoning. Here is how Bayes would solve the problem:

Reverend Thomas Bayes, 1702-1761.

- First, let's plug our preliminary assessment of the situation into Bayes rule: the probability of fishing given that it is sunny, P(F|S) = 0.90; the probability of fishing given that it is cloudy, $P(F|C) = 0.50$; and the probability of fishing given that it is rainy, $P(F|R) = 0.10$.

- Second, the probability of each kind of weather given the predictions of the weather report can be summarized as: probability of sunny weather, $P(S) = 0.65$; probability of cloudy weather, $P(C) = 0.25$; and probability of rainy weather, $P(R) = 0.10$.
- Third, notice that the sum of the probabilities of each weather condition is 100%: $P(S) + P(C) + P(R) = 0.65 + 0.25 + 0.10 = 1.00$ and these hypotheses are mutually exclusive (in the sense that it was either sunny, or cloudy, or rainy, but not a combination of them).
- Fourth, the overall likelihood of going fishing (the denominator of the right side of Bayes rule), $P(F)$ is $0.725 = P(F|S)*P(S) + P(F|C)*P(C) + P(F|R)*P(R) = 0.90*0.65 + 0.50*0.25 + 0.15*.10$.
- We can now get to the new conclusions about our hypotheses on the weather, given the prior and new information (the latter being that the family *did* go fishing):
    - The probability that the weather was sunny given that the family went fishing is, according to Bayes rule: $P(S|F) = P(F|S)*P(S) / P(F) = 0.90*0.65 / 0.725 = 0.807$.
    - The probability that the weather was cloudy given that the family went fishing is, according to Bayes rule: $P(C|F) = P(F|C)*P(C) / P(F) = 0.50*0.25 / 0.725 = 0.172$.
    - The probability that the weather was rainy given that the family went fishing is, according to Bayes rule: $P(F|R)*P(R) / P(F) = 0.15*0.10 / 0.725 = 0.021$.
- Finally, note that $P(S|F) + P(C|F) + P(R|F) = 0.807 + 0.172 + 0.021 = 1.00$, because one of the hypotheses *must* be true (it either was sunny, or cloudy, or rainy, no other possibilities are in the game).

Bayes rule, therefore, tells you that—given the prior knowledge of the situation we had and the new information that the family did go fishing, the likelihood that the weather was sunny was the highest. Well, you could have guessed that, no? Yes, in this simple case. But notice that Bayes rule gives you additional information: first, it tells you what the *best* available estimates of the probabilities of all three hypotheses are; consequently, it tells you how much confident you can be that the weather was sunny (which is better than simply saying "it's more likely"); also, it is clear from the equations that the probability of the hypothesis that the weather was sunny went *up* with the new information (from 0.65 to 0.807); finally, Bayes theorem reminds you that your degree of confidence in any hypothesis is never either zero or one hundred per cent, although it can get very close to those extremes.

Bayesian statistical analysis is a good metaphor (some philosophers of science would say a good *description*) of how science really works. More, it is a good description of how *any* logical inquiry into the world goes if it is based on a combination of hypotheses and data. The scientist (and in general the rationally thinking person) is always evaluating several hypotheses based on her previous understanding and knowledge on the one hand and on new information gathered by observation or experiment on the other hand. Her judgment of the validity of a theory therefore changes constantly, although very rarely it does so in a dramatic fashion. Bayesian analysis, therefore, clearly shows why the extremes of anti-science and scientism are naïve positions: they correspond respectively to having certainty into the hypotheses that science never works (e.g., creationism) or always works (i.e., scientism). But attaching a probability of zero or one to a given hypothesis, as the Bayesian framework makes clear, is the same as saying that our conclusions are valid *no matter what the data say*, i.e., we take them on faith.

———

Text by Massimo Pigliucci
Distributed by the Pigliucci Lab
University of TN—www.nature-nurture.org