In Vivo: Basics of Bayes (a.k.a. “Afterthoughts”)

Joachim Vandekerckhove^{a,1}

This is a preprint of a manuscript to appear in *Computational Modeling of Cognition and Behavior* (Cambridge University Press, UK).

For many psychology students, Bayesian statistics remains shrouded in mystery. At the undergraduate level, Bayes’ theorem may be taught as part of probability theory, but the link between probability theory and scientific inference is almost never made. This is unfortunate, as this link—first made almost a century ago—provides a mathematically elegant and robust basis for the quantification of scientific knowledge.

Prior | Likelihood | Bayes’ theorem | Bayesian inference

As argued by Wrinch and Jeffreys (3) and later in the works of Harold Jeffreys, *probability theory is extended logic*. Jaynes (2) calls it “the logic of science.” Indeed, it is easy to see how probability theory maps directly to propositional logic if all statements are fully “true” or “false” – that is, all probabilities are either 0 or 1. Take for example the statement $P(A|B) = 1$. “If B is true, then the probability of A is 1” is simply another way of saying that “ B implies A ” ($B \rightarrow A$). Similarly, $P(A|B) = 0$ is the same as $B \rightarrow \neg A$. Probability theory extends this concept to include uncertainty, but the rules of probability have the same status as the rules of logic: they can be used to derive statements that are guaranteed to be correct if the premises are correct. Paraphrasing Edwards et al. (1, p. 194): Probability is orderly uncertainty and inference from data is revision of uncertainty in the light of relevant new information. Bayesian statistics, then, is nothing more—and nothing less—than the application of probability theory to real problems of inference.

The close relationship of probability theory and logic leads to further fertile insights. For example, a common misunderstanding regarding Bayesian methods is that they are somehow invalidated by the fact that conclusions may depend on the prior probabilities assigned to parameter values or hypotheses. Translated to the terminology of formal logic, this claim is that logical deduction is somehow invalidated because conclusions depend on premises. Clearly, an inferential procedure is not pathological because its conclusions depend on its assumptions – rather the inverse is true. Conclusions that do not depend on assumptions may be robust, but they cannot be rational any more than conclusions that do not depend on observations.

However, the dependence on prior probabilities involves another dimension that is often misunderstood: At first glance, it appears that the prior introduces the analyst’s beliefs—an element of subjectivity—into the inference, and this is clearly undesirable if we are to be objective in our scientific inquiries. Two observations address this issue. First, it is important to emphasize—lest we forget—that “subjective” is not synonymous with “arbitrary.” Rather than beliefs, we may think of probability as conveying information. It is not at all peculiar to say that relevant information may be subjective – after all, not all humans have access to the same information. Accordingly, the information that is encoded in probability

distributions may be subjective, but that does not mean it is elective. Belief—in the sense in which it is used in probability theory—is not an act of will, but merely a state in which the individual passively finds itself. Accordingly, different scientists using different sources of information can rationally reach different conclusions.

The second observation regarding the subjectivity of the prior follows from inspection of Bayes’ theorem:

$$P(\theta|y) = \frac{P(y|\theta)P(\theta)}{P(y)}.$$

In the right hand side numerator appears the product $P(y|\theta)P(\theta)$: likelihood and prior side by side determine the relative density of all possible values of θ . In a typical cognitive-modeling scenario, researchers will specify these distributions with some care – much defense and reasoning will often go into the selection of which prior to use, possibly using arguments from previous literature and graphical exploration of the prior predictive distribution; criticism of prior decisions is common and expected. The likelihood is defined also.

The way the components of Bayes’ theorem are specified is somewhat reminiscent of the Biblical description of the creation of the heavens, in which “God made two great lights; the greater light to rule the day, and the lesser light to rule the night: he made the stars also” (Gen 1:16, KJV). Much like how in this verse the billions upon trillions of stars are created as an afterthought, far less argument is usually deemed necessary for the definition of the likelihood function even though it is usually much more consequential than the definition of the prior – after all, given even moderate amounts of data the prior will typically wash out in favor of the likelihood. To see argument at all for the choice of likelihood is not typical and the tacit assumptions of sequential independence and normally-distributed residuals are ubiquitous. Jaynes (2) writes that “if one fails to specify the prior information, a problem of inference is just as ill-posed as if one had failed to specify the data” (p. 373), but the emphasis here can apply to both factors in the RHS numerator of Bayes’ theorem: if we fail to question the likelihood it is as if we fail to question the prior.

In some contexts, questioning the likelihood is common: we ask whether this or that is the “right model for the data.” For example, in the reaction time modeling world, we might wonder if a set of observations is best described by a standard linear ballistic accumulator or by some stochastic variant. In more conventional scenarios, we sometimes worry if a t test with equal variances is appropriate, or an unequal-variance procedure should be used

^aUniversity of California, Irvine

¹To whom correspondence should be addressed. E-mail: joachim@uci.edu.

instead. This invites a question: What if we want to estimate the magnitude of some manipulation effect but are unwilling to commit to model E (equal variance) or U (unequal variance)? Perhaps unsurprisingly, probability theory has an answer. If the posterior distribution of the effect size assuming some model M ($M \in \{E, U\}$) is $p(\delta|y, M)$ and the posterior probability that E is the correct model of the two is $P(E|X) = 1 - P(U|X)$, then the posterior distribution of δ , *averaged over these two models*, is immediately given by the sum rule of probability:

$$p(\delta|y) = p(\delta|y, E)P(E|y) + p(\delta|y, U)P(U|y).$$

One interpretation of this equation is that the exact identity of the model is a nuisance variable, and we can “integrate it out” by taking an average weighted by the posterior probability of each model. It provides a posterior distribution of δ that does not assume that model E is true or that model U is true, only that one of

them is. This technique of marginalizing over models is a direct consequence of probability theory that is often called Bayesian model averaging. It can be applied in a staggering variety of circumstances.

While most psychologists readily draw conclusions that are based on an often arbitrary and tenuously appropriate likelihood, one who is uncomfortable with any of the assumptions can apply Bayesian model averaging to assuage their concerns. This way, we can avoid having to commit to any particular set of model assumptions function by averaging over likelihood functions – and so it goes with priors also.

- 1 Ward Edwards, Harold Lindman, and Leonard J Savage. Bayesian statistical inference for psychological research. *Psychological Review*, 70(3):193, 1963.
- 2 Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge University Press, 2003.
- 3 Dorothy Wrinch and Harold Jeffreys. On certain fundamental principles of scientific inquiry. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 42(249): 369–390, 1921.