Bayesian Analysis (2010) **5**, Number 1, pp. 39–40

Comment on Article by Hogg et al.

John Skilling* and Devinder Sivia†

This is a beautifully illustrated and clear account of a particular data analysis problem in the physical sciences. It is particularly helpful to see this sort of problem in the literature of statistics, whose techniques can serve any of the sciences.

Different sciences import different outlooks. To oversimplify, social science data is "soft" and modelled by smooth distributions, physical science models are tightly defined with the data defining a likelihood function that can be arbitrarily digital and rough, while biological sciences are nowadays dominated by size and complexity. Statistics used to be oriented towards the social sciences, with physical sciences being an almost trivial (usually frequentist and wrong) aside. That has changed, and it's valuable that statisticians be exposed to the differing needs. The particular topic of neutron scattering measurement of magnetism is specific, but the authors' approach is general. And it raises basic points.

Theory suggests a model for magnetism in solids, described by a 4-parameter function of radial separation r;

$$G(r) = \Sigma_L \exp(-r/\xi_L) + \Sigma_S \, r^{-1} \exp(-r/\xi_S)$$

Neutron scattering measures the Fourier transform of G, being the intensity scattered sideways as a beam traverses the magnetised solid. Of course, matters aren't *quite* so simple when instrumentation is involved, and the data finally involve 16 partially known parameters, not 4. And minor complexity is added because the sample is embedded in magnetic fields H of three different strengths (0, 0.2 and 1 Tesla), so there are three cases to consider. Welcome to the real world.

Bayesian analysis demands careful attention to priors as well as to likelihood, but the prize is that "the richness of information available [in the posterior] enables novel forms of presentation, capable of conveying deep insight into uncertainty and correlation at a glance". Quite so!

Yet there remains scope for improvement. All too often, Bayesian computation remains expensive, but that may be due to algorithm inefficiency more than intrinsic difficulty. Eyeballing these results suggests that posterior parameters are here located to a few percent accuracy, and there are only 16 of them, so the prior-to-posterior compression is no more than a factor $30^{16} = 2^{80}$. A good algorithm, compressing geometrically, should not need more than 80 iterates to accomplish this. Then, having a seed point, 100 random samples from the posterior should be more than enough to acquire any chosen property, along with its uncertainty, each needing a few MC proposals from the seed. So 1000 or so iterates should suffice. Yet the authors use a million, which takes 4 days and seems to be inefficient by 3 orders of magnitude. Better off-the-shelf algorithms please!

*Maximum Entropy Data Consultants, Kenmare, Ireland, mailto:skilling@eircom.net
†St John's College, Oxford, UK mailto:devinder.sivia@sjc.ox.ac.uk

 DOI:10.1214/10-BA501B

More fundamentally, and this is a distressingly common omission, there is no mention of the *evidence* (*a.k.a.* marginal likelihood) value. Bayesian inference produces this fundamental number *as well as* the posterior. As a matter of courtesy, that number should always be calculated and stated. Otherwise, researchers with a different model or a different prior are forced to reproduce this calculation as well as performing their own before they can reach the Bayes factor that guides the decision as to which model is better. Evidence please!

But let us not seem churlish by applying future standards to present practice. We like the paper.