

Bayesian Hierarchical Modeling

David Draper

Department of Applied Mathematics and Statistics
University of California, Santa Cruz
draper@ams.ucsc.edu

Introductory Tutorial I1

ISBA 2004: Viña del Mar, Chile

Sunday 23 May 2004, 10.30am–12.15pm

This material provides coverage of introductory topics arising in the formulation, fitting, and checking of *hierarchical* or *multilevel* models from the Bayesian point of view. Hierarchical models (HMs) arise frequently in at least the following four kinds of applications:

- HMs are common in fields such as health and education, in which data—both outcomes and predictors—are often gathered in a *nested* or hierarchical fashion, e.g., patients within hospitals, or students within classrooms within schools. HMs are thus also ideally suited to the wide range of applications in government and business in which single- or multi-stage *cluster samples* are routinely drawn, and offer a unified approach to the analysis of *random-effects* (variance-components) and *mixed models*.
- A different kind of nested data arises in *meta-analysis* in, e.g., medicine and the social sciences. In this setting the goal is combining information from a number of studies of essentially the same phenomenon, to produce more accurate inferences and predictions than those available from any single study. Here the data structure is subjects within studies, and as in the clustered case above there will generally be predictors available at both the subject and study levels.
- For simplicity people often try to model data as (conditionally) IID at a fairly high level of aggregation—for instance, by pretending that all the subjects in a sampling experiment are drawn homogeneously from a single population. In fact, heterogeneity is often the rule rather than the exception, and frequently the available predictor variables do not “explain” this heterogeneity sufficiently. With recent computational advances in *Markov chain Monte Carlo* (MCMC) methods it’s becoming increasingly straightforward to at least describe such heterogeneity with *mixture models* that employ *latent variables* (unobserved predictors) in a hierarchical structure.
- Hierarchical modeling also provides a natural way to treat issues of *model selection* and *model uncertainty* with all types of data, not just cluster samples. For example, in regression, if the data appear to exhibit residual variation that changes with the predictors, you can expand the model that assumes constant variation, by embedding it in a family of models that span a variety of assumptions about residual variation. In this way, instead of having to choose one of these models and risk making the wrong choice, you can work with several models at once, weighting them in proportion to their plausibility given the data.

In this tutorial the principles of Bayesian hierarchical modeling will be developed, with emphasis on practical rather than theoretical issues, and illustrated with real data drawn from case studies involving applications of HMs in clinical trials, cluster sampling, and latent-variable modeling (and covering all four of the categories of HMs above). Extensive details required for carrying out the analyses will be provided, including hardcopy of a number of sessions with R, Maple, and WinBUGS.