Highly Structured Models, High-Energy Astrophysics, and Efficient Computation

David A. van Dyk
Department of Statistics
University of California, Irvine

In recent years, a trend has been growing in applied statistics—it is becoming ever more feasible to build application specific models which are designed to account for the hierarchical and latent structures inherent in any particular data generation mechanism. Such multi-level models have long been advocated on theoretical grounds, but recently the development of new computational tools (e.g., hardware, software, and algorithms) for statistical analysis has begun to bring such model fitting into routine practice. Of course, much remains to be done. The flexibility of methods that rely on Markov chain Monte Carlo, for example, comes at a cost—they require problem specific coding, long computation times, and present difficulties in ascertaining convergence, especially when a posterior distribution is multi-modal.

This talk both illustrates the power of multi-level models to solve real scientific problems and illustrates some of the inherent computational difficulties. In particular, we describe a highly structured model for image reconstruction using data obtained with the Chandra X-ray Space Observatory. The resulting reconstructed images are more valuable in terms of the substantive science than other available reconstructions.

The second part of the talk discusses computational issues involved in fitting such highly structured models. In particular, we discuss constructing a Gibbs sampler with incompatible conditional distributions. The conditional distributions $p(x|y)$ and $p(y|x)$ are said to be compatible if there exists a joint distribution, $p(x, y)$, such that $p(x|y) = p(x, y)/\int p(x, y)dx$ and $p(y|x) = p(x, y)/\int p(x, y)dy$. In this case, a Gibbs sampler can be constructed by iteratively sampling from the two conditional distributions and under certain regularity conditions the stationary distribution of the resulting Markov chain is $p(x, y)$. Much less is known about the behavior of a Gibbs sampler that is constructed using incompatible conditional distributions. We explore this possibility and show that careful choice of such incompatible distribution can lead to a Gibbs sampler with known stationary distribution and with a better geometric rate of convergence than the standard sampler. The method of marginal data augmentation gives a prescription for such a choice. We illustrate our methods using an example from high-energy Astrophysics and demonstrate the improved rate of convergence of the sampler, some of its unexpected behavior, and warnings to potential users.