

Statistics 651: Introduction to Applied Bayesian Methods

Mini-Project # 4

Due: November 13, 2009, 5pm

Student evaluations of professors provide information to administrators who are making tenure decisions for faculty members. Evaluations are made on a 7 point scale (1 = poor, ..., 7 = outstanding). The “benchmark” for faculty performance is the overall teacher rating. One measure of faculty teaching is the average score for that question (averaged across students). The website <http://madison.byu.edu/bayes/faculty.dat> represents a sample of Fall 2002 faculty members. We made a strong assumption on HW # 2, that the evaluations came from the same distribution. Relax that assumption and allow each faculty member their own mean. That is, reproduce the analysis you did in HW # 2 where you allow $Y_i \sim N(\theta_i, \sigma^2)$, instead of $Y_i \sim N(\theta, \sigma^2)$.

- Choose a prior distribution for the parameters of the likelihood function. *Justify* this prior distribution!
- Use MCMC techniques to:
 1. Calculate the posterior distribution of the parameters of the model (all of them).
 2. Calculate $E[\Theta|\mathbf{Y}]$, the posterior mean (the whole thing).
 3. Calculate $V[\Theta|\mathbf{Y}]$, the posterior variance (how are you going to summarize this???)
 4. Calculate the predictive distribution of the “next” average faculty evaluation.
 5. In general, it is considered “good” to score a 5 or better. Calculate the probability that a randomly selected professor gets a 5 or better in her next evaluation.