STANFORD UNIVERSITY

DEPARTMENT OF STATISTICS

STATISTICS SEMINAR

4:15 p.m., Tuesday, May 4, 2004
Sequoia Hall Room 200
(Cookies at 3:45 in the 1st Floor Lounge)

Ed George
University of Pennsylvania

Improved Minimax Prediction Under Kullback-Leibler Loss

Abstract:

Let $X|\mu \sim N_p(\mu, v_x I)$ and $Y|\mu \sim N_p(\mu, v_y I)$ be independent p-dimensional multivariate normal vectors with common unknown mean μ, and let $p(x|\mu)$ and $p(y|\mu)$ denote the conditional densities of X and Y. Based on only observing $X = x$, we consider the problem of obtaining a predictive distribution $\hat{p}(y|x)$ for Y that is close to $p(y|\mu)$ as measured by Kullback-Leibler loss. The natural straw man for this problem is the best invariant predictive distribution, the Bayes rule $p_U(y|x)$ under the uniform prior $\pi_U(\mu) \equiv 1$, which is seen to be minimax. We show that $p_U(y|x)$ is dominated by any Bayes rules for which the square root of the marginal distribution is superharmonic. This yields wide classes of dominating predictive distributions including Bayes rules under superharmonic priors. These dominating predictive shrinkage distributions can be constructed to adaptively shrink $p_U(y|x)$ towards arbitrary points or subspaces. Those procedures corresponding to superharmonic priors can be further combined to obtain minimax multiple shrinkage predictive distributions that adaptively shrink $p_U(y|x)$ towards an arbitrary number of points or subspaces. Fundamental similarities and differences with the parallel theory of estimating a multivariate normal mean under quadratic loss are described throughout. (This is joint work with Feng Liang and Xinyi Xu).