

Department of Mathematical Sciences  
The Johns Hopkins University

SEMINAR

Edward I. George  
Department of Statistics  
The Wharton School  
University of Pennsylvania

November 14, 2002  
304 Whitehead Hall  
Refreshments: 3:30 p.m.  
Seminar: 4:00 p.m.

DILUTION PRIORS FOR MODEL UNCERTAINTY

ABSTRACT

In model uncertainty setups, some models may be very similar to others. For example, this would occur with the uncertainty about which of several highly correlated variables to include in a linear regression. In such setups, it is tempting to characterize ignorance by a uniform prior which assigns equal probability to each model under consideration. Unfortunately, such a prior may not be uniform on “neighborhoods of models,” especially when there are regions of high redundancy in the model space. By skewing the posterior, such priors can be problematic for MCMC sampling of the posterior for model search and for model averaging. Dilution priors rectify this problem by assigning probability more uniformly to model neighborhoods.

This talk will focus on the construction of dilution priors for two setups: Bayesian variable selection for the linear model and Bayesian CART model selection. For the linear model, several potential dilution prior constructions will be presented. These include a construction based on a Voronoi tessellation of the space of all models and a construction based on metrics over the space of all models. For Bayesian CART, dilution priors are seen to arise naturally with tree-generating process priors.