

Objective Bayesian Model Selection in Gaussian Graphical Models

Carlos Carvalho
University of Chicago

January 24, 2008

This paper presents a default model-selection procedure for Gaussian graphical models that involves two new developments. First, we develop an objective version of the hyper-inverse Wishart prior for restricted covariance matrices, called the HIW g-prior, and show how it corresponds to the implied fractional prior for covariance selection using fractional Bayes factors. Second, we apply a class of priors that automatically handles the problem of multiple hypothesis testing implied by covariance selection. Numerical experiments show that these priors strongly control the number of false edges included in the model, thereby automatically rewarding sparsity. We demonstrate our methods on a variety of simulated examples, concluding with a real example analyzing covariation in mutual-fund returns. These studies reveal that the combined use of a multiplicity-correction prior on graphs with the hyper-inverse Wishart g-prior on covariance matrices yields better performance than conventional covariance selection methods.