



THE UNIVERSITY OF
CHICAGO

Department of Statistics

STATISTICS COLLOQUIUM

JAMES JOHNDROW

Department of Statistics
Stanford University

Bayes Shrinkage at GWAS Scale

THURSDAY, January 31, 2019 at 1:00 PM
Kent 120, 1020 E. 58th St.

ABSTRACT

Bayesian analysis in high-dimensional sparse regression settings differs from typical frequentist methods in part by eschewing selection of a single model. While this simplifies some aspects of inference, computation for these models is considerably less scalable than popular procedures for selection, such as the lasso. Global-local shrinkage priors, such as horseshoe, were proposed in part as a computationally scalable alternative to classical spike-and-slab mixture priors, but the promised computational scalability has failed to materialize. We propose an MCMC algorithm for computation with the horseshoe prior that permits analysis of genome-wide association study (GWAS) data with hundreds of thousands of covariates and thousands of subjects. The algorithm is shown empirically to outperform alternatives by orders of magnitude. Among other enhancements, our algorithm employs certain approximations to an expensive matrix operation. We give general results on the accuracy of time averages obtained from approximating Markov chains, with an application to our horseshoe MCMC sampler.