

Statistical computing and data analysis with complex models

Petascale Hierarchical Modeling via Parallel Execution

Andrew Gelman
Columbia University
New York, N.Y. 10027

Abstract

We have developed a fast and flexible open-source C++ program, Stan, that automatically performs Bayesian inference in high-dimensional continuous spaces. Stan (sampling through adaptive neighborhoods) incorporates automatic analytical differentiation, adaptation of tuning parameters, and the no U-turn sampler, a combination of Hamiltonian dynamics and slice sampling that allows the algorithm to move efficiently through high-dimensional continuous spaces.

In this talk I will outline our motivations in constructing Stan, consider directions for improvement, and discuss connections between computation, modeling, and data analysis, based on our own applied experiences.

This work is joint with Bob Carpenter, Matt Hoffman, and others in our research group at Columbia.