




SEMINAR

Relaxation of Projected Prior with Continuous Gap Shrinkage

by Leo Li Duan
University of Florida



Thursday 19 Nov 2026, 2.30 pm

Departement of Statistical Sciences - Room Benvenuti

Projected priors were originally introduced to accommodate parameter constraints, but have recently regained popularity due to their ability to assign probability mass to low-dimensional parameter sets, such as the spaces of sparse vectors, directed acyclic graphs, or transport plans. When employed as a transformation of random variables, projection is especially useful, since its contraction property not only preserves probability concentration, but also often preserves differentiability for gradient-based posterior computation. On the other hand, unless the projection can be obtained by some non-iterative algorithm, posterior computation can be expensive because it requires nesting an iterative optimization routine within each Markov chain Monte Carlo iteration. In this article, inspired by the success of continuous shrinkage models as replacements for discrete spike-and-slab priors, we propose a continuous relaxation of projected priors. The key idea is to quantify the duality gap between the primal projection loss and the dual objective, and impose a probabilistic prior that shrinks this gap toward zero. The resulting gap-shrinkage prior has a tractable form, does not require running an optimization subroutine inside each posterior update, and puts probability mass near the exact projection. We demonstrate useful properties of gap-shrinkage priors, including connections to global-local shrinkage priors, broad applicability to generalized projection functions, and competitive performance in posterior contraction. We apply the gap-shrinkage model to a marketing data analysis aimed at identifying important predictor effects on multivariate grocery-shopping decisions.