Analysis of the Gibbs Sampler for hierarchical and crossed-effect Gaussian models via multigrid decomposition

PRESENTING AUTHOR[1], Omiros Papaspiliopoulos[2], Gareth Roberts[3]

[1]*Bocconi University, Italy*
[2]*Universitat Pompeu Fabra, Spain*
[3]*University of Warwick, UK*

Abstract

We study the convergence properties of the Gibbs Sampler in the context of Gaussian hierarchical and crossed-effect models. We develop a novel methodology based on multi-grid decompositions to derive analytic expressions for the convergence rates of the algorithm, extending significantly the class of conditionally Gaussian models amenable to direct analysis. In the hierarchical context, our work gives a rather complete understanding of the Gibbs Sampler behavior for symmetric models (with arbitrary depth), while providing approximations and bounds for the non-symmetric cases. The theoretical results give rise to simple and easy-to-implement guidelines to optimize practical implementations of the Gibbs samplers on such models. While the good performances of the Gibbs Sampler in hierarchically-structured models is renowned, the context of crossed-effect models is drastically different. Here hierarchical centering is not possible and the convergence of commonly implemented Gibbs Sampler strategies deteriorates as the data-size increases, resulting in super-linear computational complexity (potentially even quadratic) in the number of data-points. We show how to leverage the negatively-correlated structure of crossed-effect models to design easy-to-implement collapsed Gibbs Samplers whose complexity matches the one of hierarchical scenarios.