Title: Bayesian inference for high-dimensional nonstationary Gaussian processes

Mark Risser
Lawrence Berkeley National Laboratory

Abstract: In spite of the diverse literature on nonstationary spatial modelling and approximate Gaussian process (GP) methods, there are no general approaches for conducting fully Bayesian inference for moderately sized nonstationary spatial data sets on a personal laptop. For statisticians and data scientists who wish to conduct posterior inference and prediction with appropriate uncertainty quantification, the lack of such approaches and software is a limitation. In this work, we develop methodology for implementing formal Bayesian inference for a general class of nonstationary GPs. Our novel approach uses pre-existing frameworks for characterizing nonstationarity in a new way while utilizing via modern GP likelihood approximations. Posterior sampling is implemented using flexible MCMC methods, with nonstationary posterior prediction conducted as a post-processing step. We demonstrate our novel methods on three data sets, ranging from several hundred to over 50,000 locations. All of our methods are implemented in the freely available BayesNSGP software package for R.