Numerical Methods in Bayesian Inference

U. von Toussaint[1]
(1) Max-Planck-Institute for Plasmaphysics,
Boltzmannstrasse 2, 85748 Garching, Germany

May 31, 2011

Abstract

Bayesian inference is very simple from a conceptual point of view: Once the likelihood and prior distributions are specified Bayes' theorem allows to derive the posterior probability for every specified parameter vector. However, in most situations the posterior distribution is required primarily for the purpose of evaluating expectation values of a function of interest $f(\boldsymbol{\theta})$ with respect to the posterior,

$$\langle f(\boldsymbol{\theta})\rangle = \int \mathrm{d}\boldsymbol{\theta}\, f(\boldsymbol{\theta})\, p(\boldsymbol{\theta}|\mathbf{D}, I) = \int \mathrm{d}\boldsymbol{\theta}\, f(\boldsymbol{\theta})\, \frac{p^*(\boldsymbol{\theta})}{Z} \tag{1}$$

The normalization constant of the unnormalized distribution $p^*(\boldsymbol{\theta})$ is given by

$$Z = \int \mathrm{d}\boldsymbol{\theta}\, p^*(\boldsymbol{\theta})\,. \tag{2}$$

These integrals over the parameter space are commonly high-dimensional and analytically intractable, except in very rare circumstances, so that typically neither the expectation value nor the normalization constant are at hand - the latter the key quantity for Bayesian model comparison. Also the marginalization of parameters requires integration in often high-dimensional spaces. There are two different ways to proceed. Either the integrant of Eq. (1) is approximated by a different, more easily accessible function or the integral itself is approximated by numerical integration or by sampling (MCMC) techniques. In the tutorial the key concepts and algorithms to evaluate these integrals are presented and their respective merits are compared using real-world examples.

Key Words: Bayesian Data Analysis, Numerical Methods, Markov Chain Monte Carlo