

Abstract

Bayesian inverse problems have become an important part of comprehensive parameter studies for complex models. However, the construction of posterior samples can be computationally expensive, especially if one model run is costly and/or the dimension of the parameter space is high. Markov chain Monte Carlo methods are popular techniques used to construct (correlated) samples of an arbitrary distribution. Many of them, however, have deteriorating properties in high-dimensional spaces that can indeed be solved but only through additional computational cost. Thus, finding low-dimensional structure can be one approach to make the solution of a Bayesian inverse problem more efficient. The particular approach that is demonstrated is the active subspace method. Active subspaces are structures along which a function of interest is only mildly changing, on average (or not at all). If they are present in a particular inference problem, only a few directions in the parameter space drive the update from the prior to the posterior. Exploiting this fact by running Markov chains in this lower dimensional subspace, and thus improving their mixing behavior, can speed up the construction of posterior samples. We shall see advantages of this approach in a real case study of a methane hydrate model, but will also discuss possible hurdles and weaknesses of the technique in general.