EGU24-16036, updated on 20 May 2024
https://doi.org/10.5194/egusphere-egu24-16036
EGU General Assembly 2024

An Independent State sampler for Trans-dimensional Bayesian Inference

Malcolm Sambridge[1], Andrew Valentine[2], and Juerg Hauser[3]
[1]Research School of Earth Sciences, Australian National University, Canberra, Australia (malcolm.sambridge@anu.edu.au)
[2]Dept. of Earth Sciences, Durham University, South Road, Durham, DH1 3LE, UK (andrew.valentine@durham.ac.uk)
[3]CSIRO, Mineral Resources, Black Mountain, Clunies Ross Street, Acton, Australian Capital Territory, Australia (Juerg.Hauser@csiro.au)

Over the past twenty years, Trans-dimensional Bayesian Inference has become a popular approach for Bayesian sampling. It has been applied widely in the geosciences when the best class of model representation, e.g. of the subsurface, is not obvious in advance, or the number of free variables undecided. Making arbitrary choices in these areas may result in sub-optimal inferences from data. In trans-D, one typically defines a finite number of model *states*, with differing numbers of unknowns, over which Bayesian Inference is to be performed using the data.

A key attraction of Trans-D Bayesian Inference is that it is designed to *let the data decide* which state, as well as which configurations of parameters within each state, are preferred by the data, in a probabilistic manner. Trans-D algorithms may hence be viewed as a combination of fixed dimensional *within-state* sampling and simultaneous *between-state* sampling where Markov chains visit each state in proportion to their support from the data.

In theory, each state may be completely independent, involving different classes of model parameterization, with different numbers of unknowns, data noise levels, and even different assumptions about the data-model relationship. Practical considerations, such as convergence of the finite length Markov chains between states, usually mean that each state must be closely related to each other, e.g. differing by a single layer in a 1-D seismic Earth model. In addition, since the form of the necessary Metropolis-Hastings balance condition depends on the mathematical relationship between the unknowns in each state, then implementations are often bespoke to each class of model parameterization and data type. To our knowledge there exists no automatic trans-D sampler where one can define arbitrary independent states, together with a prior and Likelihood, and simply pass to a generalised sampling algorithm, as is common with many fixed dimensional MCMC algorithms and software packages.

A second limitation in trans-D sampling is that since implementations are bespoke within a class of model parameterizations, within-state sampling is typically performed with simplistic and often dated algorithms, e.g. Metropolis-Hastings or Gibbs samplers, thereby limiting convergence rates. Over the past 30 years fixed dimensional sampling has advanced considerably with numerous efficient algorithms available and many conveniently translated into user friendly software packages, almost all of which have not been used within a trans-D framework due to a lack of a way to conveniently deploy them in a trans-D setting.

In this presentation we will address all of these issues by describing the theory under-pinning an 'Independent State' (IS) Trans-D sampler, together with some illustrative examples. In this algorithm class, sampling may be performed across states that are completely independent, containing arbitrary numbers of unknowns and parameter classes. In addition, the IS-sampler can conveniently take advantage of any fixed dimensional sampler without the need to derive and re-code bespoke Markov chain balance conditions, or specify mechanisms for transitions between model parameters within different states. In this sense it represents a general purpose automatic trans-D sampler.