Lifted samplers for partially ordered discrete state-spaces

Philippe Gagnon.
University of Montreal.

30th October 2020, 4pm UK time

Abstract

Lifting is a well known technique employed to avoid that the Markov chains used as samples backtrack too often. It consists in *lifting* the state-space to include *direction* variables for guiding the Markov chains. Its implementation is direct when the probability mass function targeted is defined on a totally ordered set, such as that of a integer-valued random variables In this paper, we adapt this technique to the situation where only a partial order can be established on the sampling space and explore its benefits. Important applications include the simulation of systems formed from binary variables, such as those arising in the Ising model, and variable selection when the posterior model probabilities can be evaluated, up to a normalising constant. To accommodate for the situation where one does not have access to these marginal model probabilities, a lifted trans-dimensional sampler for partially ordered model spaces is introduced. We show through theoretical analyses and empirical results that the lifted samplers outperform their non-lifted counterparts in some situations, but not always, achieving this at no extra computational cost. The code to reproduce all experiments is available online.