
Anytime Monte Carlo

Lawrence Murray*¹

¹Department of Statistics, University of Oxford – United Kingdom

Abstract

Monte Carlo algorithms typically simulate some fixed number of samples, n , with the real time taken to do so a random variable, $T(n)$. For the purposes of real-time deadlines, particularly in a distributed computing context, an alternative is to fix the real time, t , and allow the number of samples drawn in this time to be a random variable, $N(t)$. Naive estimators constructed from these $N(t)$ samples are not necessarily consistent, however, and in general exhibit length bias with respect to compute time. This talk will introduce a framework for dealing with the length bias for both iid and Markov chain Monte Carlo samplers, and demonstrate the utility of the approach on a large scale sequential Monte Carlo deployment on the Amazon EC2 cloud computing infrastructure.

*Speaker