MIXING TIME OF THE CONDITIONAL BACKWARD SAMPLING PARTICLE FILTER

ANTHONY LEE

The conditional backward sampling particle filter (CBPF) is a powerful Markov chain Monte Carlo sampler for general state space hidden Markov model smoothing. It was proposed as an improvement over the conditional particle filter, which is known to have an O(T^2) computational time complexity under a general `strong' mixing assumption, where T is the time horizon. We provide the first proof that the CBPF admits an O(T logT) time complexity under strong mixing, complementing strong empirical evidence of the superiority of the CBPF in practice. In particular, the CBPF's mixing time is upper bounded by O(logT), for any sufficiently large number of particles N that depends only on the mixing assumptions and not T. We show that an O(logT) mixing time is optimal. The proof involves the analysis of a novel coupling of two CBPFs, which involves a maximal coupling of two particle systems at each time instant. The coupling is implementable, and thus can also be used to construct unbiased, finite variance, estimates of functionals which have arbitrary dependence on the latent state's path, with a total expected cost of O(TlogT). We also investigate other couplings, and we show some of these alternatives have improved empirical behaviour.