On the relationship between Adaptive Monte Carlo and Online Learning

Yee Whye Teh

University of Oxford and DeepMind, UK

Monte Carlo sampling underpins a large range of statistical analysis, but can be computationally costly. Adaptive methods aim to improve computational efficiency of Monte Carlo sampling procedures, by using previous samples to gain insights about the target distribution and adapt the procedures based on these insights. However, adaptation can also lead to biased inferences, e.g. where the stationary distribution deviates from the desired target. This process of sequentially learning about a target distribution and adapting the sampling procedure is reminiscent of online learning, where the community has developed a rich set of methods to develop online learning algorithms and theoretical tools to analyse the convergence of the algorithms.

In this talk I will describe an initial step on this path. Specifically, we propose and study a partition-based adaptive importance sampling (AIS) procedure called Daisee as an online learning problem. Borrowing ideas from the online learning community, we argue for the importance of the trade-off between exploration and exploitation in this adaptation, and show a $T^{-1/2}$ convergence rate. We then extend Daisee to adaptively learn a hierarchical partitioning of the sample space for more efficient sampling and confirm the performance of both algorithms empirically.

This is joint work with Xiaoyu Lu.