Multi-fidelity Bandit Optimisation*

Kirthevasan Kandasamy

Abstract

In many scientific and engineering applications, we are tasked with the optimisation of an expensive to evaluate black box function. Traditional methods for this problem assume just the availability of this single function. However, in many cases, cheap approximations may be available. For example, in optimal policy search in robotics, the expensive real world behaviour of a robot can be approximated by cheap computer simulations. We formalise this task as a multi-fidelity bandit problem where the target function and its approximations are sampled from a Gaussian process. We develop a method based on upper confidence bound (UCB) techniques and prove that it uses the approximations to eliminate low function value regions and uses the expensive evaluations mostly in a small region containing the optimum. For instance, in the above robot control example, our method would use the simulations to quickly eliminate suboptimal policies, while reserving real world trials for a small set of promising candidates.

I will begin this talk with reviews on UCB methods for bandits and Gaussian processes before proceeding to multi-fidelity optimisation. A preliminary version of the paper is available at www.cs.cmu.edu/~kkandasa/pubs/kandasamy16mfbo.pdf.

*Internal Talk, Gatsby Unit, July 12, 2016.