Gaussian process accelerated ABC[*]

Richard Wilkinson

Abstract

Approximate Bayesian computation (ABC) is a class of algorithms used for doing Bayesian inference when you do not have access to the likelihood function. Instead, all simulation is done using realisations from the simulation. These methods are widely applicable, easy to implement, and consequently have become popular in many applied scientific disciplines. One of the major challenges for ABC methods is dealing with the computational cost that arises from needing to repeatedly run the simulator.

In this talk, I will discuss several approaches for overcoming this cost. There are two main approaches: inverse modelling, where regression models (random forests, neural nets etc) are used to directly learn a mapping from the (high dimensional) simulator output to the input parameter; and surrogate modelling approaches, where we instead seek to approximate the behaviour of the forward simulator (i.e. from parameter to simulator output), before inverting it to learn the parameters using Bayes theorem. I will concentrate on the latter, but discuss the pros and cons of both approaches and discuss how approaches from machine learning have led to advances in the field.

[*] Machine Learning External Seminar, Gatsby Unit, September 21, 2016.