Multifidelity approximate Bayesian computation

Ruth Baker.
University of Oxford, UK.

16 July 2020, 11.30 UK time

Abstract
A vital stage in the mathematical modelling of real-world systems is to calibrate a model's parameters to observed data. Likelihood-free parameter inference methods, such as Approximate Bayesian Computation, build Monte Carlo samples of the uncertain parameter distribution by comparing the data with large numbers of model simulations. However, the computational expense of generating these simulations forms a significant bottleneck in the practical application of such methods. We identify how simulations of cheap, low-fidelity models have been used separately in two complementary ways to reduce the computational expense of building these samples, at the cost of introducing additional variance to the resulting parameter estimates. We explore how these approaches can be unified so that cost and benefit are optimally balanced, and we characterise the optimal choice of how often to simulate from cheap, low-fidelity models in place of expensive, high-fidelity models in Monte Carlo ABC algorithms. The resulting early accept/reject multifidelity ABC algorithm that we propose is shown to give improved performance over existing multifidelity and high-fidelity approaches.