Project IV 2026/27

Modern Bayesian computation for stochastic dynamical systems

Supervisor: Andrew Golightly
Project research area: Statistics

Background

Models defined by stochastic differential equations (SDEs) allow for the representation of random variability in dynamical systems. SDEs can be used to model many continuous-time processes such as stock price, predator-prey dynamics and epidemics and cellular processes. Solutions of nonlinear SDEs are rarely available in closed form. Consequently, performing simulation and inference for multidimensional SDE models is challenging. One way to tackle this problem is to use a numerical solver (e.g. Euler-Maruyama, Milstein etc) to generate approximate solutions over a fine time grid. When observations are available at discrete times, Bayesian inference proceeds via a suitable prior specification combined with the observed data likelihood. The latter typically involves an intractable integral, precluding the application of standard Metropolis-Hastings techniques. This project will therefore leverage recent posterior sampling techniques including pseudo-marginal Metropolis-Hastings and approximate Bayesian computation.

This project will focus on these modern simulation-based techniques for fitting discretised SDEs to discrete-time data. At least one of these methods will be implemented in R and applied to real and synthetic data sets. Potential application areas will depend on the interests of the project student but could involve stochastic volatility models of financial time series data, stochastic SIR models of epidemic data, predator-prey (Lotka-Volterra) models for data arising from population ecology or models of intra-cellular processes.

Anticipated Outcomes

By the end of this project, you will have an understanding of:

- Ito integrals and their role in stochastic differential equations.
- Simple discretisation schemes such as Euler-Maruyama.
- The observed data likelihood and its role in Bayesian inference for SDEs.
- Pseudo-marginal Metropolis-Hastings methodology and the application thereof to toy problems and to SDEs.
- Various approximate Bayesian computation schemes and their application to SDEs.
- The challenges and limitations of these inference methodologies in real world settings.

By the end of the project, you will be able to:

- Discuss existing approaches to Bayesian inference for SDEs with focus on either approximate Bayesian computation or pseudo-marginal Metropolis-Hastings, illustrated with simple examples.
- Write a suite of bespoke R functions for simulating and fitting nonlinear, multivariate SDEs to discrete-time data.
- Apply of one or both inference techniques to real and/or synthetic data.
- Compare and contrast inference approaches in terms of efficiency.

Mode of operation in evidence of learning

The project will involve learning through reading and programming in R. Students will demonstrate their understanding by comparing theory to simulation results, writing R code to implement core methodology, analysing simulated and real data sets, and clearly communicating the material in both written and oral formats.

Essential prior modules: MATH3421 Bayesian Computation and Modelling III

Resources (indicative)

- Overview of Bayesian inference for SDEs
- ABC for SDEs
- Pseudo-marginal MH for SDEs

Further information: contact Andy Golightly (andrew.golightly@durham.ac.uk)