STATS 531 Homework 6

Due Sunday 3/22

Please submit your homework report to Canvas, including both the Quarto (qmd) source file and an pdf compiled from it. Also submit other files needed for your qmd file to compile, but you do not need to submit a copy of the data.

Your report should contain a reference section listing sources. The grader should be able to clearly identify where the sources were used, for example using reference numbers in the text. Anything and anyone consulted while you are working on the homework counts as a source and should be credited. AI has become omnipresent in coding, so it is not practical to credit every time you consult an AI. It is good scholarship to describe which AI you used, and provide useful relevant information about how it was used. The homework will be graded following the grading scheme in the syllabus.

Getting started with Pypomp

Computation time is an unavoidable consideration when working with simulation-based inference, for all but small datasets and simple models.

Pypomp builds on JAX for parallelization, compilation and automatic differentiation. Some notes to get you started are on the course website. Time spent reading the JAX documentation at https://docs.jax.dev/ will be useful.

Pypomp is a new library. Please don't hesitate to contact the GSI or instructor if you notice errors or room for improvement. Recall that extra credit is available for posting relevant issues on the course GitHub site.

For this homework, you are required to use Pypomp rather than the previous R-pomp. Much online material is available for R-pomp. R-pomp has a similar structure, where the user must specify `rprocess`, `dmeasure`, etc. However, R-pomp reqquires the user to write these in shippets of C code which must then be compiled. In Pypomp, all the compilation is handled by JAX, so the user writes the model in Python. Model structures and data analysis strategies can be learned from R-pomp examples.

Question 6.1. Exploring behavior of a POMP model: simulating an SIR process.

Write a solution to Exercise 2.3 from Chapter 13 (Simulation of stochastic dynamic models). Note the following:

- We are working toward formal inference for POMP models. Nevertheless, playing with your model by plotting simulations at various parameter values is a useful exercise for getting to understand how your model behaves. It is not enough to know just what parameter value maximizes the likelihood, we also want to understand enough about the model to be able to interpret this MLE. What types of behavior can the model exhibit? How could we describe the behaviors that are consistent with the data?
- Worked solutions are available online. You can look at these if you get stuck, but you should avoid becoming over-dependent. You may like to look at them after solving the homework independently. Your solution is welcome to discuss the relationship between your investigation of the model and the posted solutions.

Question 6.2. Modifying a POMP model: Adding a latent period to the SIR model

Write a solution to Exercise 2.4 from Chapter 13 (Simulation of stochastic dynamic models).

Acknowledgements

The questions derive from material in a short course on Simulation-based Inference for Epidemiological Dynamics using Pypomp, which in turn derives from a previous version of this course using R-pomp.