Analysis of Victoria's COVID-19 Delta wave using particle-marginal MCMC

Chelsea Pham
School of Mathematics and Statistics, University of Melbourne

1. An Important Problem

In June 2021, Victoria saw an uptick of COVID-19 Delta-strain cases which led to several lockdowns. We aim to understand the dynamics of COVID-19 under the lockdown restrictions. In this project, we explore the use of Particle-Marginal Markov Chain Monte Carlo (pMCMC) to infer dynamics of Victoria's COVID-19 Delta wave from June to mid-November 2021. Particularly we are interested in the parameters:

- Reproduction number (R_0) - the average number of secondary infections caused by an infected host,
- Observation probability (P_{obs}) - the probability that an infected individual becomes a case, and
- Infectious period ($\frac{1}{\gamma}$) - the average duration of an infection.

2. Model

We implemented an exact Bayesian approach using particle Markov chain Monte Carlo (pMCMC) [1] with a bootstrap particle filter (PF) to estimate the likelihood of observations. The PF [Fig. 1] acts to estimate the likelihood by generating p particles which can be thought of as 'random' hypotheses of points belonging to the posterior density which are recursively weighted and resampled.The likelihood estimate from the PF is then incorporated into the usual MCMC algorithm to make inference on parameters of interest.

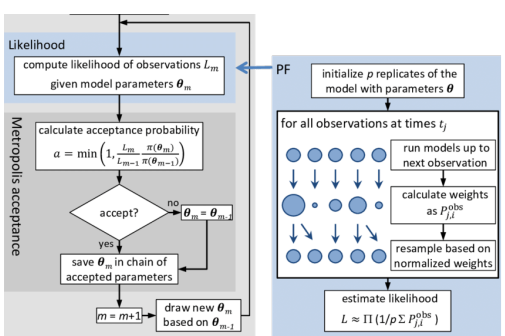

Fig. 1: Flowchart of a pMCMC routine (left) using a bootstrap particle filter(right). *Note: P^{obs} here is defined as the probability of observing the particles, it is not the same as the P_{obs} referred to in Section 1. Sourced from [3]

We assumed the underlying epidemic-model of the COVID-19 Delta strain as an SIR model with a Poisson observation process where:

$$C_t \sim \text{Poisson}((\Delta R_t + \epsilon)P_{obs}),$$

where C_t is the number of cases on day t, and ΔR_t is the number of newly recovered individuals. In this project, we also included a small constant ϵ added to ΔR_t to prevent numerical errors in the case where ΔR_t is zero.
This observation process reflects the plausible situation that the average observed counts of those reported infectious are less than the true values. This is often referred to as compensating the *under-reporting rate*.

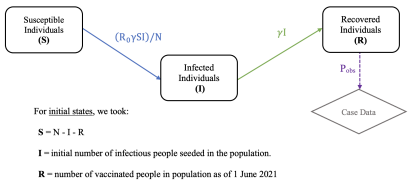

Fig. 2: SIR Model and demonstration of P_{obs} being the relationship between compartment R and the reported daily Delta cases. N is the VIC population in June 2021

Due to the time-frame of our data being short and the Victorian border restrictions reinforcing minimal change in state population, our SIR model does not account for population turnover. Vaccines were assumed 'perfect'. i.e. vaccinated people were counted as recovered population.

3. Results

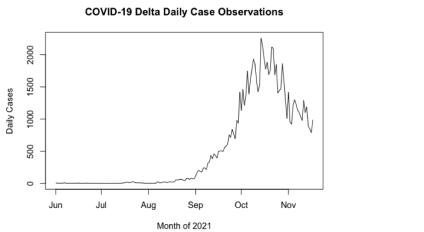

Fig. 3: COVID-19 Delta daily cases from June to mid-November 2021

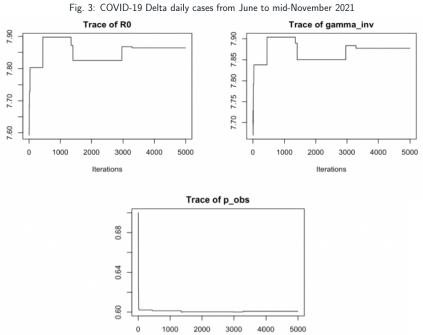

Fig. 4: Trace plots of parameters

Our model was run with 50 particles and 5000 iterations, occupying a run-time of 45 hours. These simulations provided parameter estimates:

Parameter	Estimate (Median)	Standard Deviation
R_0	7.864	0.0193
$\frac{1}{\gamma}$	7.877	0.0137
P_{obs}	0.600	0.0002

Our results show that on average, a single infected person can be expected to transmit COVID-19 to just under 8 people, a person with the disease is contagious for just over a week and the probability that we observe an infected individual as recovered is 60%. We acknowledge the limited number of iterations and particle use in these results; which was a consequence of the extremely high computational-cost of pMCMC with large scale population sizes.

4. Remarks

- Our trace plot results reveal the difficulty in mixing from pMCMC.
- P_{obs} was lower than we expected, leading us to slightly question if there could be an identifiability issue [2].
- We witnessed that configuration of the particle filter can significantly effect results; it was important to have a large number of particles for the acceptance rates in MCMC to avoid becoming too low.
- It is recommended that a super computer and ample time for running code is used for future re-runs of this method to allow for greater numbers of iterations and particles. Due to our restricted performance, we would like to highlight the findings in this project as preliminary insights at best.

5. Acknowledgements

I would like to thank my supervisor, James Walker, for introducing me to the interesting world of epidemiology and for their unwavering support and guidance throughout this project. Also a special thank you to the School of Mathematics & Statistics, especially Roy Ridgeway and Thomas Quella, for this valuable scholarship program which has offered me and many other aspiring students a first glimpse into mathematical research and its countless opportunities.

6. References

[1] Akira Endo, Edwin van Leeuwen, and Marc Baguelin. "Introduction to particle Markov-chain Monte Carlo for disease dynamics modellers". In: *Epidemics* 29 (2019), p. 100363. issn: 1755-4365. doi: https://doi.org/10.1016/j.epidem.2019.100363. url: https://www.sciencedirect.com/science/article/pii/S1755436519300301.

[2] Andreas Raue et al. "Addressing parameter identifiability by model-based experimentation". In: *IET systems biology* 5.2 (2011), pp. 120–130.

[3] Jonas Šukys and Mira Kattwinkel. "SPUX: Scalable Particle Markov Chain Monte Carlo for Uncertainty Quantification in Stochastic Ecological Models". In: *Advances in Parallel Computing* 32 (Nov. 2017).