MSA101/MVE187 2017 Lecture 15

Petter Mostad

Chalmers University

October 17, 2017

Approximate Bayesian Computations (ABC)

- In our Bayesian inference methods so far, simulation from the posterior $\pi(\theta \mid x)$ is based on being able to compute, for various θ, $\pi(x \mid \theta)\pi(\theta)$, (at least up to a constant).
- What if we do not have a formula for the likelihoood $\pi(x \mid \theta)$?
- Example: Our stochastic "model" could be some very complex stochastic computer simulation program $R(\theta)$ producing a value for x given a value for θ.
- Idea for simulating from the posterior: Simulate θ from the prior, and keep only those θ with $R(\theta) = x$.

- Example:
  - θ is binary with $P(\theta = 1) = 0.6$
  - x is binary with $\Pr(x = 1 \mid \theta = 1) = 0.9$, $\Pr(x = 1 \mid \theta = 0) = 0.1$
  - If the data is $x = 1$ then simulated values $\theta = 1$ would be kept with probability 0.9, simulated values $\theta = 0$ would be kept with probability 0.1.
  - We see the result corresponds to simulating $\theta = 1$ with probability $0.54/0.58 = 0.93$; correct according to Bayes formula.
- For continuous variables x we would get zero acceptance probability unless we replace the acceptance criterion $R(\theta) = x$ with $R(\theta) \approx x$.
- The most basic ABC algorithm defines a distance function ρ on the set where x lives, and an acceptance threshold ϵ. Then $\theta_1, \dots, \theta_n$ are simulated from the prior, and those θ_i with $\rho(\theta_i, x) \leq \epsilon$ are accepted.

ABC: Using sufficient statistic or similar

- In any (Bayesian) analysis, the likelihood $\pi(x \mid \theta)$ can be replaced by the corresponding likelihood $\pi(S(x) \mid \theta)$ of a sufficient statistic $S(x)$.
- Simple example: The likelihood of data $x = (x_1, \ldots, x_n)$, where $x_i \sim \text{Normal}(\theta, 1)$ can be replaced with the likelihood of $S(x) = \overline{x} \sim \text{Normal}(\theta, 1/n)$.
- If we can only simulate $x = S(\theta)$ we are unlikely to know a sufficient statistic. HOWEVER, we may know a function S that "summarizes" the features of the data that depend on θ, then replace x with $S(x)$.

ABC: Modelling the likelihood

- In realistic examples the acceptance rate or the accuracy still becomes too low.
- A solution: Try to simulate the "correct" θ:
  - Example: If $R(\theta_1)$ and $R(\theta_2)$ are "on either side of x", maybe $(\theta_1 + \theta_2)/2$ will result in a value closer to x.
  - (Illustration)
- Note: Targeting the simulation of θ in this way means the acceptance must be adjusted accordingly.