Intro to ABC

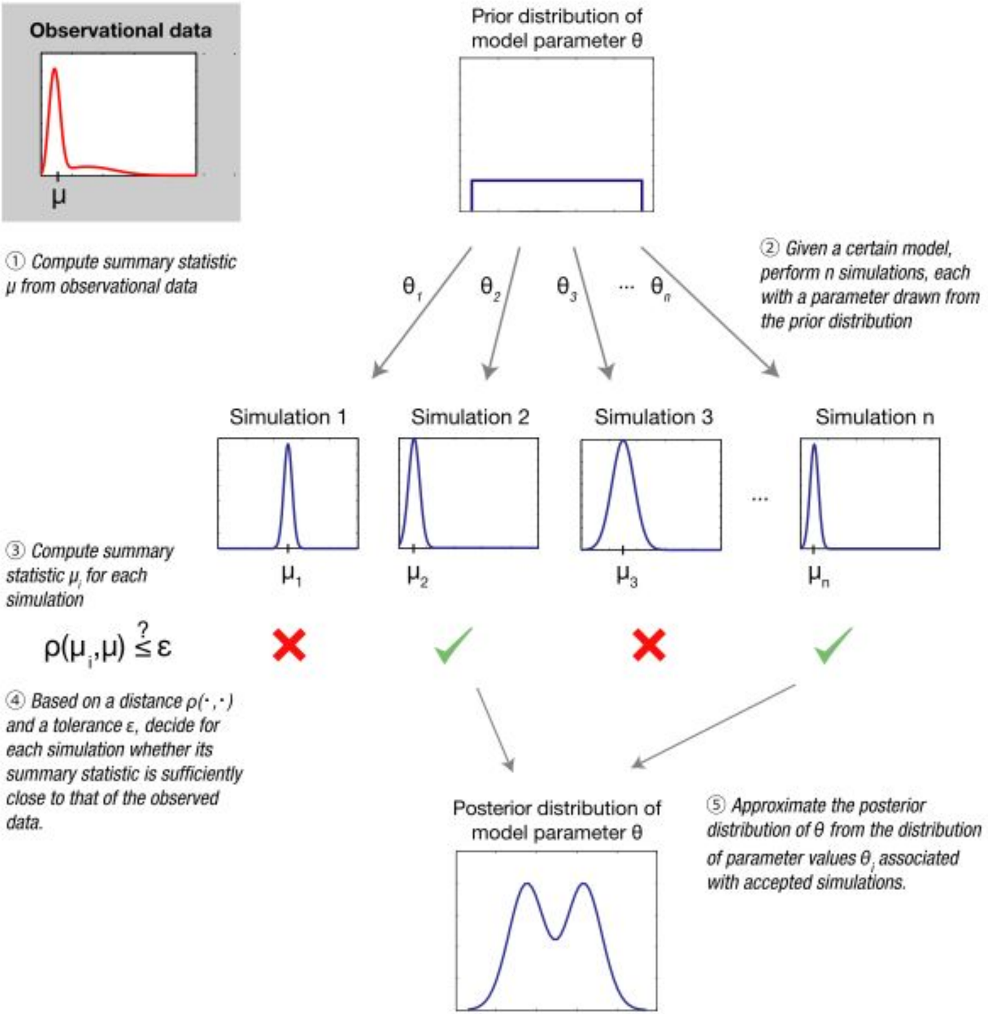
Observational data
μ
① Compute summary statistic μ from observational data
Prior distribution of model parameter θ
θ1 θ2 θ3 ··· θn
② Given a certain model, perform n simulations, each with a parameter drawn from the prior distribution
Simulation 1
Simulation 2
Simulation 3
Simulation n
μ1 μ2 μ3 μn
③ Compute summary statistic μi for each simulation
ρ(μi,μ) ≤? ε
④ Based on a distance ρ(·,·) and a tolerance ε, decide for each simulation whether its summary statistic is sufficiently close to that of the observed data.
Posterior distribution of model parameter θ
⑤ Approximate the posterior distribution of θ from the distribution of parameter values θi associated with accepted simulations.

Likelihood-Free Rejection Sampling Algorithm

Inputs:

- A target posterior density $\pi(\theta|y_{obs}) \propto p(y_{obs}|\theta)\pi(\theta)$, consisting of a prior distribution $\pi(\theta)$ and a procedure for generating data under the model $p(y_{obs}|\theta)$.
- A proposal density $g(\theta)$, with $g(\theta) > 0$ if $\pi(\theta|y_{obs}) > 0$.
- An integer $N > 0$.

Sampling:
For $i = 1, \ldots, N$:

1. Generate $\theta^{(i)} \sim g(\theta)$ from sampling density g.
2. Generate $y \sim p(y|\theta^{(i)})$ from the likelihood.
3. If $y = y_{obs}$, then accept $\theta^{(i)}$ with probability $\frac{\pi(\theta^{(i)})}{Kg(\theta^{(i)})}$, where $K \geq \max_\theta \frac{\pi(\theta)}{g(\theta)}$. Else go to 1.

Output:
A set of parameter vectors $\theta^{(1)}, \ldots, \theta^{(N)}$ which are samples from $\pi(\theta|y_{obs})$.

Likelihood-Free Rejection Sampling Algorithm

3. If $\|y - y_{obs}\| \leq h$, then accept $\theta^{(i)}$ with probability $\frac{\pi(\theta^{(i)})}{Kg(\theta^{(i)})}$, where $K \geq \max_\theta \frac{\pi(\theta)}{g(\theta)}$.

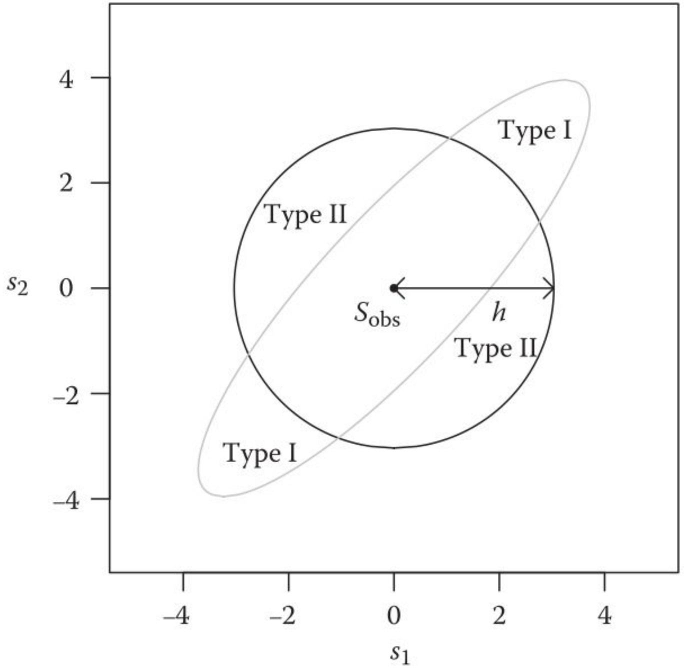

Kernel functions instead of indicators

rejecting θ if $\|y - y_{obs}\| \leq h$

$$I(\|y - y_{obs}\| \leq h)p(y|\theta)g(\theta)$$

As an alternative, use a kernel function:

$$K_h(u) = \frac{1}{h}K\left(\frac{u}{h}\right)$$

$$u = \|y - y_{obs}\|$$

Common kernel functions

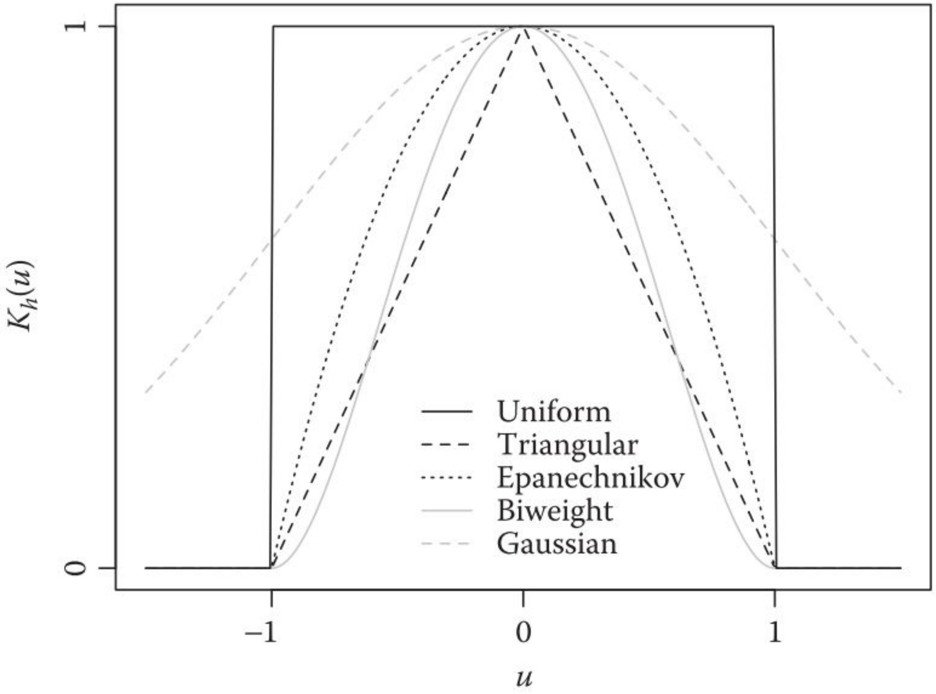

Curse of dimensionality - Summary statistics

Sampling:

For $i = 1, \ldots, N$:

1. Generate $\theta^{(i)} \sim g(\theta)$ from sampling density g.
2. Generate $y \sim p(y|\theta^{(i)})$ from the likelihood.
3. Compute summary statistic $s = S(y)$.
4. Accept $\theta^{(i)}$ with probability $\frac{K_h(\|s - s_{obs}\|)\pi(\theta^{(i)})}{Kg(\theta^{(i)})}$ where $K \geq K_h(0)\max_\theta \frac{\pi(\theta)}{g(\theta)}$. Else go to 1.

Output:

A set of parameter vectors $\theta^{(1)}, \ldots, \theta^{(N)} \sim \pi_{ABC}(\theta|s_{obs})$.

Summary questions:

- Choice of `h` and distance measure
- Choice of kernel
- Summary statistics
- Approximations due to other ABC techniques
- Sensitivity to prior
- Applications