Lecture 4, October 2nd, 2025

1. Beta-binomial hierarchical model + rstan implementation
2. Bayesian analysis of SIR model of infectious disease dynamics

Beta-binomial hierarchical model

Let's consider a generalization of our problem from lecture 1, where we wanted to estimate a population proportion from a finite sample. Now, we want to estimate multiple proportions and we suspect that these proportions may be similar.

Let $y^T = (y_1, \dots, y_m)$, where

$$y_i \mid \theta_i \sim \text{Binomial}(n_i, \theta_i), \quad i = 1, \dots, m$$

and y_is are independent given θ_is. We put the following prior on θ_is:

$$\theta_1, \dots, \theta_m \overset{iid}{\sim} \text{Beta}(\alpha, \beta)$$

But here is a big difference: α and β are unknown parameters that get their own priors:

$$\alpha \sim \text{Exp}(\lambda_\alpha)$$

$$\beta \sim \text{Exp}(\lambda_\beta)$$

As you can see, there 2 layers of parameters $\Rightarrow$

multi-level

=> the word "hierarchical" in the name

In statistics, we often describe data generating models via a graphical model / directed acyclic graph (DAG):

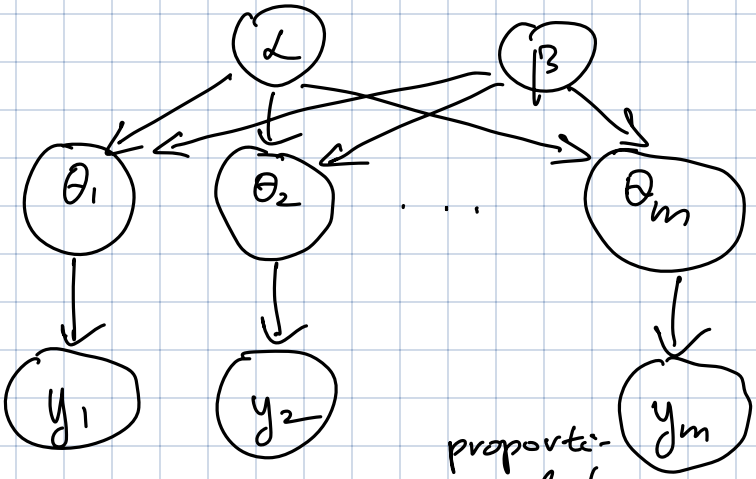

Posterior: $p(\theta_1, \ldots, \theta_m, \alpha, \beta \mid \underline{y}) \propto \prod_{i=1}^{m} Pr(y_i \mid \theta_i) \times \prod_{i=1}^{m} p(\theta_i \mid \alpha, \beta) \times p(\alpha) \times p(\beta)$

proportional to

likelihood | prior 1 | prior 2 | joint prior

In our computer demo, n_i = # rats in a rat colony under study and y_i = # of tumor occurences in rat colony i.

Hamiltonian Monte Carlo (HMC)

HMC is a particular type of Rosenbluth-Hastings algorithm, coupled with data augmentation. Uses Hamiltonian dynamics to generate R-H proposals. Works well in high dimensions, especially the no-u-turn-sampler (NUTS) variant of HMC.

Implemented in

- Stan (C++, but can be called from R, Python, and Julia

- Pyro (Python)
- Turing (Julia)
- ⋮
- many more

computer demo here

Fitting ODE-based SIR model to case data

When population size $N \to \infty$, many stochastic models converge to deterministic models.

Susceptible-Infectious-Removed (SIR) model

$$\frac{dS(t)}{dt} = -\beta S(t) I(t)$$

$$\frac{dI(t)}{dt} = \beta S(t) I(t) - \gamma I(t)$$

$$\frac{dR(t)}{dt} = \gamma I(t)$$

$$\frac{dN_{SI}(t)}{dt} = \beta S(t) \cdot I(t)$$

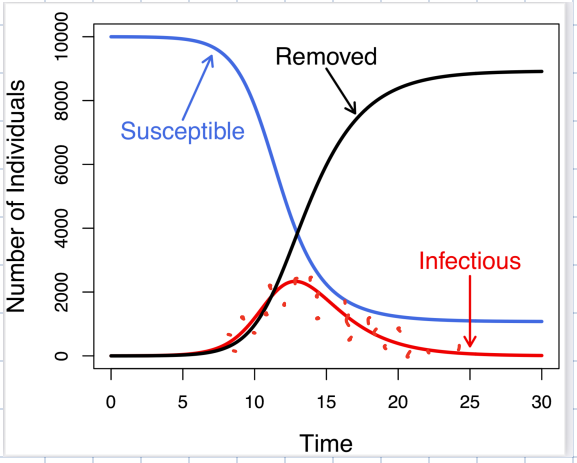

$S(0), I(0)$ ← initial conditions

$S(t) + I(t) + R(t) = N$

Data: $y = (y_1, \dots, y_n)$, where

y_i = number of new cases in interval $[t_i, t_{i+1})$

$N_{SI}(t)$ = cumulative incidence or the total number of infections up to time t

$\Delta N_{SI}^{(i)} = N_{SI}(t_{i+1}) - N_{SI}(t_i)$ = # of new infections (observed and unobserved) during interval $[t_i, t_{i+1})$

We assume that given model parameters y_is are independent and

$$y_i \sim \text{Negative-Binomial}\left(\mu = \rho \times \Delta N_{SI}^{(i)},\ \sigma^2 = \mu\left(1 + \frac{\mu}{\varphi}\right)\right)$$

ρ – case detection probability

φ – overdispersion parameter

μ – mean of N-B .

Q: Why use Negative-Binomial?

A: Convenient for count data and more flexible than Poisson distribution

Likelihood: $\Pr(y_1, \ldots, y_n \mid \beta, \gamma, \rho, \varphi, S(0), I(0)) = \prod_{i=1}^{n} \Pr(y_i \mid \beta, \gamma, \rho, \varphi, S(0), I(0))$

N-B pmf

Posterior: $p(\beta, \gamma, \rho, \varphi, S(0), I(0) \mid y_1, \ldots, y_n) \propto$

$\propto \Pr(y_1, \ldots, y_n \mid \beta, \gamma, \rho, \varphi, S(0), I(0)) \times p(\beta, \gamma, \rho, \varphi, S(0), I(0))$

need to specify priors

computer demo here

Further reading materials