MSA101/MVE187 2017 Lecture 7

Petter Mostad

Chalmers University

September 25, 2018

Gibbs sampling

- Gibbs sampling for a density $\pi(x) = \pi(x_1, \ldots, x_n)$ over n variables iterates between using the j different proposal functions ($j = 1, \ldots, n$)

$$q_j(x^* \mid x) = q_j(x_1^*, \ldots, x_n^* \mid x_1, \ldots, x_n)$$

where $q_j(x^* \mid x) = 0$ unless $x_i^* = x_i$ for $i \neq j$ and

$$q_j(x_1, \ldots, x_j^*, \ldots, x_n \mid x_1, \ldots, x_j, \ldots, x_n) = \pi(x_j^* \mid x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)$$

- The quotient in the acceptance probability becomes

$$\frac{\pi(x^*)q(x \mid x^*)}{\pi(x)q(x^* \mid x)} = \frac{\pi(x_1, \ldots, x_j^*, \ldots, x_n)\pi(x_j \mid x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)}{\pi(x_1, \ldots, x_j, \ldots, , x_n)\pi(x_j^* \mid x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)}$$

$$= \frac{\pi(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)}{\pi(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)} = 1$$

- So: Gibbs sampling changes one coordinate of x at a time, simulating from the conditional densities.
- These conditional densities are in many cases easy to derive.

Hierarchiclal models

- Sometimes, observed data have dependencies that can best be described using a hierarchy.
- Example: Test results for students may depend on the class they are in, the school they attend, and the country they live in.
- A statistical model for the data should then contain a variable for each "source of infuence"; they would depend on each other in a hierarchy, which can be drawn as an upside-down tree.
- When making computations, the tree structure can be very useful, for example in Gibbs sampling.

Hierarchical models example

- Example in Chapter 7 of Albert: Estimating the mortality rates due to heart transplants in 94 hospitals: $\lambda_1, \ldots, \lambda_{94}$.
- We have counts y_i and "exposures" e_i, and assume $y_i \sim \text{Poisson}(\lambda_i e_i)$. Questions: What is the probability that the next transplant patient at hospital i will die? What is the probability that hospital i has a lower mortality rate than hospital j?
- Possibility 1: All λ_i are independent. Problematic.
- Possibility 2: All λ_i are equal. Does the model fit the data?
- Possibility 3: A hierarchical model.