

Solutions for module 6

The Gibbs sampler

1 IQ test

1. The joint distribution of $(\underline{x}, \underline{\mu}, \mu_G, \tau_G)$ has pdf

$$\begin{aligned}
 \pi(\underline{x}, \underline{\mu}, \mu_G, \tau_G) &= \pi(\underline{x}|\underline{\mu}, \mu_G, \tau_G)\pi(\underline{\mu}, \mu_G, \tau_G) \\
 &= \pi(\underline{x}|\underline{\mu}, \mu_G, \tau_G)\pi(\underline{\mu}|\mu_G, \tau_G)\pi(\mu_G, \tau_G) \\
 &= \left(\prod_{i=1}^n \pi(x_i|\mu_i) \right) \left(\prod_{i=1}^n \pi(\mu_i|\mu_G, \tau_G) \right) \pi(\mu_G)\pi(\tau_G) \\
 &= \left(\prod_{i=1}^n \sqrt{\frac{\tau_i}{2\pi}} \exp\left(-\frac{1}{2}\tau_i(x_i - \mu_i)^2\right) \right) \left(\prod_{i=1}^n \sqrt{\frac{\tau_G}{2\pi}} \exp\left(-\frac{1}{2}\tau_G(\mu_i - \mu_G)^2\right) \right) \times \\
 &\quad \left(\sqrt{\frac{\tau_0}{2\pi}} \exp\left(-\frac{1}{2}\tau_0(\mu_G - \mu_0)^2\right) \right) \left(\frac{1}{\Gamma(\alpha)\beta^\alpha} \tau_G^{\alpha-1} e^{-\tau_G/\beta} \right).
 \end{aligned}$$

2. The pdf for the full conditional distribution of μ_i is

$$\begin{aligned}
 \pi(\mu_i|\mu_1, \dots, \mu_{i-1}, \mu_{i+1}, \dots, \mu_n, \mu_G, \tau_G, \underline{x}) \\
 &\propto \pi(x_i|\mu_i)\pi(\mu_i|\mu_G, \tau_G) \\
 &= \sqrt{\frac{\tau_i}{2\pi}} \exp\left(-\frac{1}{2}\tau_i(x_i - \mu_i)^2\right) \sqrt{\frac{\tau_G}{2\pi}} \exp\left(-\frac{1}{2}\tau_G(\mu_i - \mu_G)^2\right) \\
 &\propto \exp\left(-\frac{1}{2}(\tau_i + \tau_G)\mu_i^2 + (\tau_i x_i + \tau_G \mu_G)\mu_i\right),
 \end{aligned}$$

which we recognise as the unnormalised density of a normal distributed random variable with mean $(\tau_i x_i + \tau_G \mu_G)/(\tau_i + \tau_G)$ and precision $(\tau_i + \tau_G)$. Hence the full conditional distribution of μ_i is

$$\mu_i|\mu_1, \dots, \mu_{i-1}, \mu_{i+1}, \dots, \mu_n, \mu_G, \tau_G, \underline{x} \sim N\left(\frac{\tau_i x_i + \tau_G \mu_G}{\tau_i + \tau_G}, \tau_i + \tau_G\right).$$

The pdf for the full conditional distribution of μ_G is

$$\begin{aligned}
 \pi(\mu_G|\underline{\mu}, \tau_G, \underline{x}) &\propto \prod_{i=1}^n \pi(\mu_i|\mu_G, \tau_G)\pi(\mu_G) \\
 &= \left(\prod_{i=1}^n \sqrt{\frac{\tau_G}{2\pi}} \exp\left(-\frac{1}{2}\tau_G(\mu_i - \mu_G)^2\right) \right) \sqrt{\frac{\tau_0}{2\pi}} \exp\left(-\frac{1}{2}\tau_0(\mu_G - \mu_0)^2\right) \\
 &\propto \exp\left(-\frac{1}{2}\tau_G \sum (\mu_i - \mu_G)^2\right) \exp\left(-\frac{1}{2}\tau_0(\mu_G - \mu_0)^2\right) \\
 &= \exp\left(-\frac{1}{2}(n\tau_G + \tau_0)\mu_G^2 + (\tau_G n\bar{\mu} + \tau_0 \mu_0)\mu_G - \frac{1}{2}(\tau_G + \tau_0) \sum \mu_i^2\right).
 \end{aligned}$$

Again we recognise this an unnormalised density of a normal distributed random variable obtaining

$$\mu_G|\underline{\mu}, \tau_G, \underline{x} \sim N\left(\frac{n\tau_G \bar{\mu} + \tau_0 \mu_0}{n\tau_G + \tau_0}, n\tau_G + \tau_0\right).$$

Finally, the full conditional distribution for τ_G is given by

$$\begin{aligned}\pi(\tau_G | \underline{\mu}, \mu_g, \underline{x}) &\propto \prod_{i=1}^n \pi(\mu_i | \mu_G, \tau_G) \pi(\tau_G) \\ &= \left(\prod_{i=1}^n \sqrt{\frac{\tau_G}{2\pi}} \exp\left(-\frac{1}{2}\tau_G(\mu_i - \mu_G)^2\right) \right) \tau_G^{\alpha-1} \beta^\alpha e^{-\tau_G/\beta} \\ &\propto \tau_G^{n/2+\alpha-1} \exp\left(-\tau_G \left(\frac{1}{2} \sum_{i=1}^n (\mu_i - \mu_G)^2 + 1/\beta\right)\right),\end{aligned}$$

which can be recognised as the unnormalised density of a gamma distributed random variable, hence

$$\tau_G | \underline{\mu}, \mu_g, \underline{x} \sim \text{Gamma}\left(\frac{n}{2} + \alpha, \left(\frac{1}{2} \sum_{i=1}^n (\mu_i - \mu_G)^2 + 1/\beta\right)^{-1}\right).$$

3. First step in a Gibbs sampler is to chose initial values. For $\mu_1^{(0)}, \dots, \mu_n^{(0)}$ and μ_G any values will, in principle, do. The initial value $\tau_G^{(0)}$ should be positive. A more “clever” choice could be to set $\mu_i^{(0)} = x_i$ for $i = 1, \dots, n$, $\mu_G^{(0)} = \bar{x} \equiv \frac{1}{n} \sum_{i=1}^n x_i$ and $\tau_G^{(0)} = n(\sum_{i=1}^n (x_i - \bar{x})^2)^{-1}$.

A Gibbs sampler could then proceed as follows

For $t = 1 \dots T$ do

• For $i = 1 \dots n$ do

• Generate $\mu_i^{(t)}$ from $N\left(\frac{\tau x_i + \tau_G^{(t-1)} \mu_G^{(t-1)}}{\tau^{(t-1)} + \tau_G^{(t-1)}}, \tau^{(t-1)} + \tau_G^{(t-1)}\right)$

• Generate $\mu_G^{(t)}$ from $N\left(\frac{n\tau_G^{(t-1)} \bar{\mu}^{(t)} + \tau_0 \mu_0}{n\tau_G^{(t-1)} + \tau_0}, n\tau_G^{(t-1)} + \tau_0\right)$

• Generate $\tau_G^{(t)}$ from $\text{Gamma}\left(\frac{n}{2} + \alpha, \left(\frac{1}{2} \sum_{i=1}^n (\mu_i^{(t)} - \mu_G^{(t)})^2 + 1/\beta\right)^{-1}\right)$

2 Radiocarbon dating

Note that

$$\pi(\mu_1, \mu_2, \mu_3) \propto \mathbf{1}[0 < \mu_1] \mathbf{1}[\mu_1 < \mu_2] \mathbf{1}[\mu_2 < \mu_3] \mathbf{1}[\mu_3 < k],$$

where $\mathbf{1}[\cdot]$ is the indicator function:

$$\mathbf{1}[\text{“expression”}] = \begin{cases} 1 & \text{if “expression” is true} \\ 0 & \text{otherwise.} \end{cases}$$

1. The joint posterior pdf is given by

$$\begin{aligned}\pi(\mu_1, \mu_2, \mu_3 | x_1, x_2, x_3) &\propto \pi(x_1, x_2, x_3 | \mu_1, \mu_2, \mu_3) \pi(\mu_1, \mu_2, \mu_3) \\ &= \pi(x_1 | \mu_1) \pi(x_2 | \mu_2) \pi(x_3 | \mu_3) \pi(\mu_1, \mu_2, \mu_3) \\ &= \left(\prod_{i=1}^3 \sqrt{\frac{\tau}{2\pi}} \exp\left(-\frac{1}{2}\tau(x_i - \mu_i)^2\right) \right) \times \\ &\quad \mathbf{1}[0 < \mu_1] \mathbf{1}[\mu_1 < \mu_2] \mathbf{1}[\mu_2 < \mu_3] \mathbf{1}[\mu_3 < k].\end{aligned}$$

2. The full conditional for μ_1 has pdf

$$\begin{aligned}\pi(\mu_1 | \mu_2, \mu_3, x_1, x_2, x_3) &\propto \pi(x_1 | \mu_1) p(\mu_1, \mu_2, \mu_3) \\ &= \sqrt{\frac{\tau}{2\pi}} \exp\left(-\frac{1}{2}\tau(x_1 - \mu_1)^2\right) \times \\ &\quad \mathbf{1}[0 < \mu_1] \mathbf{1}[\mu_1 < \mu_2] \mathbf{1}[\mu_2 < \mu_3] \mathbf{1}[\mu_3 < k] \\ &\propto \exp\left(-\frac{1}{2}\tau(x_1 - \mu_1)^2\right) \mathbf{1}[0 < \mu_1 < \mu_2],\end{aligned}$$

which is the density of a normal distribution $N(x_1, \tau)$ restricted to the open interval $(0, \mu_2)$. In a similar way we show that $\pi(\mu_2|\mu_1, \mu_3, x_1, x_2, x_3)$ is the density of $N(x_2, \tau)$ restricted to (μ_1, μ_3) and that $\pi(\mu_3|\mu_1, \mu_2, x_1, x_2, x_3)$ is the density of $N(x_3, \tau)$ restricted to (μ_2, k) .

3. To generate a sample μ_1 from $\pi(\mu_1|\mu_2, \mu_3, x_1, x_2, x_3)$ we proceed as follows: Generate proposals μ_1 from $N(x_1, \tau)$ until $\mu_1 \in (0, \mu_2)$. When this happens return μ_1 as a sample from $\pi(\mu_1|\mu_2, \mu_3, x_1, x_2, x_3)$. This procedure is an example of so-called rejection sampling.