Problem Sheet for Chapter 5 – Week 9

1. The density function of the general Pareto distribution Pareto(α, β) with scale $\alpha > 0$ and shape $\beta > 0$ is given by

$$\pi(y \mid \alpha, \beta) = \frac{\beta \alpha^{\beta}}{y^{\beta+1}} \mathbf{1}_{[\beta,\infty)}(y), \quad y \in \mathbb{R}$$

Suppose that $Y_1, \ldots, Y_n \mid (\alpha, \beta) \overset{\text{iid}}{\sim} \text{Pareto}(\alpha, \beta)$ and denote the observed data by $\boldsymbol{y} = \{y_1, \ldots, y_N\}$.

   (a) Using independent prior distributions $\beta \sim \Gamma(a, b)$ and $\alpha \sim \text{Exp}(c)$, derive the posterior distribution and the full conditional distributions.

   (b) Construct an MCMC algorithm to sample from the posterior distribution.

   (c) How would you check and improve the mixing of this MCMC algorithm?

2. Consider a model with two parameters, θ_1 and θ_2. By using conjugate prior distributions, the full conditional distributions, $\pi(\theta_1 \mid \boldsymbol{y}, \theta_2)$ and $\pi(\theta_2 \mid \boldsymbol{y}, \theta_1)$, have closed forms. This means a Gibbs sampler can be used.

   (a) Construct an MCMC algorithm to generate samples from the posterior distribution.

   (b) By writing down the acceptance probability for θ_1, show that a Gibbs sampler is a form of Metropolis–Hastings algorithm where the acceptance probability is always 1.