STA2311 (Fall 2023) - Practice Problems for Class 7 (MCMC Basics)

1. Suppose that $Y_{i,j} \mid \theta_i \stackrel{iid}{\sim} \mathcal{N}(\theta_i, \sigma_y^2)$ for $i = 1, \ldots, K$ and $j = 1, \ldots, J_i$, and $\theta_i \stackrel{iid}{\sim} \mathcal{N}(\mu, \sigma_\theta^2)$ for $i = 1, \ldots, K$, where $\mu \in \mathbb{R}$ and $\sigma_y^2, \sigma_\theta^2 > 0$ are unknown parameters. We are Bayesians and place an InvGamma(α_y, β_y) prior on σ_y^2, an InvGamma$(\alpha_\theta, \beta_\theta)$ prior on σ_θ^2, and a $\mathcal{N}(\alpha_\mu, \beta_\mu)$ prior on μ. Given observations $y_{1,1}, \ldots, y_{K,J_K}$, derive a Gibbs sampler for sampling from the posterior distribution of $(\mu, \sigma_y^2, \sigma_\theta^2, \theta_1, \ldots, \theta_K)$.

2. Derive and implement a Metropolis-within-Gibbs sampler for sampling from the Bayesian mixture of exponentials posterior from Homework 1.

3. Show that the multiple-try Metropolis algorithm transition rule satisfies the detailed balance condition.

4. Show by example that we can have a model for some $\boldsymbol{X} \in \mathbb{R}^d$ where the *full conditionals* (those used in the Gibbs sampler) are easy to sample from, but all the other conditioned marginals, such as $X_1 \mid X_3$, are not. *Hint: consider the case where the full marginals are exponentially distributed*

5. Show that the Gibbs sampler is a composition of d Metropolis-Hastings algorithms.

6. In lecture, we saw that the component-wise transitions of the (systematic scan) Gibbs sampler satisfy the detailed balance condition. However, show by example that the combined transitions (i.e., all/any components together) do *not* satisfy the detailed balance condition.