Exercise 1: Let the discrete random variable X be described by a probability mass function $p_X(x) = \Pr(X = x)$. The current state of a Metropolis–Hastings Markov chain is x_t, which is generated from the same distribution as X. Demonstrate that the next state x_{t+1} will also be drawn from the same distribution as X.

Exercise 2 (chapter 7.4): Let the random variable X follow a Laplace distribution with location $\mu = 0$ and scale parameter $\sigma = 2$. The density of the Laplace distribution is

$$f_X(x) = \frac{1}{2\sigma} \exp\left\{ -\frac{|x-\mu|}{\sigma} \right\} \qquad \sigma > 0\,.$$

1. Implement an independent Metropolis–Hastings sampler with a Normal$\left(0, \sigma_1^2\right)$ proposal distribution.
2. Implement a random walk Metropolis–Hastings sampler based on Normal$\left(0, \sigma_2^2\right)$ noise.
3. Compare the performance of both samplers in terms of $\mathbb{E}[X]$ and $\mathbb{V}[X]$ for various values of σ_1^2 and σ_2^2. What value of σ_2^2 is required to achieve an acceptance rate of about 40% in case of the random walk Metropolis–Hastings sampler?

Exercise 3 (chapter 7.4): Let $\{Y_i\}_{i=1}^3$ be independent and identically distributed random variables that follow a Cauchy distribution with location μ and scale parameter $\sigma = 1$. The density of the Cauchy distribution is

$$f_Y(y) = \frac{1}{\pi} \left[\frac{\sigma}{\sigma^2 + (y-\mu)^2} \right] \qquad \sigma > 0\,.$$

The prior density of the location parameter is $p(\mu) \propto \exp\left\{-\mu^2/100\right\}$.

1. Show that the posterior density has three modes when $Y_1 = 0, Y_2 = 5$ and $Y_3 = 9$.
2. Implement a random walk Metropolis–Hastings sampler based on Cauchy$\left(0, \sigma_1^2\right)$ and Normal$\left(0, \sigma_2^2\right)$ noise.
3. Compare the performance of both samplers in terms of $\mathbb{E}[\mu \,|\, y_1, y_2, y_2]$ and monitor convergence using trace plots.

Exercise 4: Let X and Y be discrete random variables with support $\{x_1, \ldots, x_n\}$ and $\{y_1, \ldots, y_m\}$. Denote the joint probability mass function of X and Y by $p_{X,Y}(x, y) = \Pr(X = x, Y = y)$. Using a Gibbs sampler, assume that convergence to the distribution of (X, Y) has occurred. Demonstrate that the next state (x_{t+1}, y_{t+1}) will also be drawn from the same distribution as (X, Y).

Exercise 5 (chapter 7.5): Let the vector $\boldsymbol{X} = [X_1, X_2]^{\mathrm{T}}$ follow a bivariate Normal distribution with zero mean vector and covariance matrix $\begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$ with $|\rho| < 1$.

1. Implement Monte Carlo simulation and Gibbs sampling to compute marginal expectations and variances.
2. Use $\rho = 0$ and generate 500 samples. Compare both methods in terms of bias.
3. Use $\rho = 0.5, 0.9, 0.99, 0.999$ and generate again 500 samples. Create trace plots and explain how the correlation affects Gibbs sampling.
4. Repeat 2. and 3. by generating 10 000 samples. Explain how Gibbs sampling improves in terms of bias when generating more samples.

Exercise 6 (chapter 7.5): Let $\{y\}_{i=1}^n$ be observations from a counting process where

$$y_i \mid \mu_1, \mu_2, \lambda \sim \begin{cases} \text{Poisson}(\mu_1) & \text{if } i \leq \lambda \\ \text{Poisson}(\mu_2) & \text{if } i > \lambda \end{cases}$$

and λ denotes a changepoint. Let the priors be

$$\begin{aligned} \mu_1 &\sim \text{Gamma}(\alpha_1, \beta_1) \\ \mu_2 &\sim \text{Gamma}(\alpha_2, \beta_2) \\ \lambda &\sim \text{Uniform}(1, 2, \dots, n) \end{aligned} \quad .$$

1. Find the likelihood and joint posterior density for the changepoint model.
2. Find all full conditional densities to implement a Gibb sampler.
3. Use the Gibbs sampler and the following data to perform changepoint detection:

$$\begin{aligned} &4, 4, 3, 1, 3, 2, 1, 0, 11, 11, 12, 4, 4, 7, 9, 6, 9, 12, 13, 15, 12, 10, 10, 6, 6, 7, 12, 11, \\ &15, 5, 11, 8, 11, 7, 11, 12, 14, 12, 8, 11, 9, 10, 6, 14, 14, 8, 4, 7, 10, 3, 14, 10, 17, 7, \\ &16, 9, 12, 11, 7, 11, 5, 11, 13, 9, 7, 9, 7, 11, 12, 13, 6, 9, 10, 13, 8, 18, 6, 16, 8, 4, 16, \\ &8, 9, 5, 7, 9, 10, 11, 13, 12, 9, 11, 7, 9, 6, 7, 6, 11, 8, 5 \end{aligned}$$

Exercise 7 (chapter 7.8): Let $\{X_i\}_{i=1}^n$ be correlated random variables with $\mathbb{V}[X_i] = \sigma^2$ for all $i = 1, \dots n$ and $\text{Cov}[X_i, X_{i+k}] = \sigma_k$ for all i, k. Consider the sample mean $\bar{X} = n^{-1} \sum_{i=1}^n X_i$ and find its variance $\mathbb{V}\big[\bar{X}\big]$.