Data 102 (Fall 2022) **Names:**
Discussion Worksheet 5
09/16/2022

1. You are searching for your phone and you are not sure where it is. It could be in your bedroom, your living room, or your car. We represent this as a prior over possible locations, which for simplicity we've projected to a 1-dimensional line:

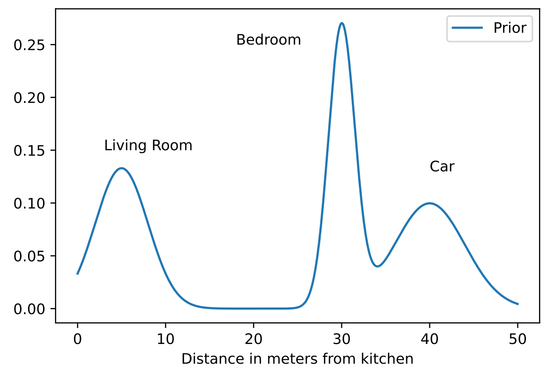

You use the Find Your Phone app, which gives you the GPS location of the phone. The GPS is accurate up to some small error, which we model as a Gaussian with a standard deviation of 2 meters. After observing the GPS signal, you use rejection sampling to sample from your posterior distribution, and observe the following samples:

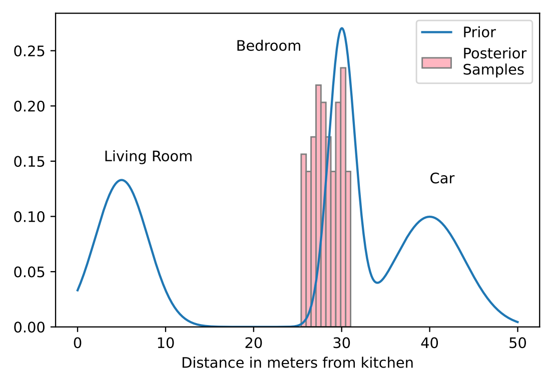

(a) Approximately where was the GPS signal? Draw it on the x-axis.

(b) Suppose you had instead seen the following posterior samples:

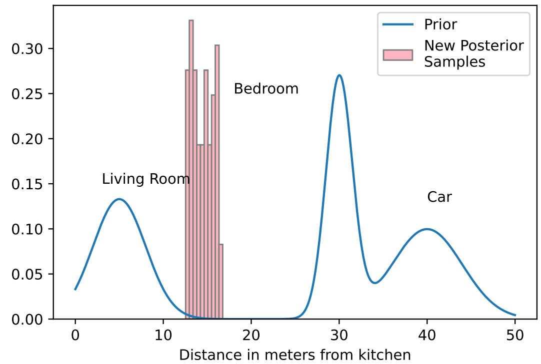

How is this result possible? What does it say about your prior?

(c) For the posterior from part (b): If you used rejection sampling with the prior as your proposal distribution, would the acceptance rate be high or low? Justify your answer.

2. Which of the following statements are true about this graphical model?

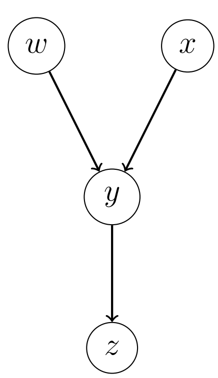

(a) $x \perp\!\!\!\perp w$

(b) $x \perp\!\!\!\perp w|y$

(c) $w \perp\!\!\!\perp z|y$

3. Given the following distribution, rank the Metropolis-Hastings proposal distributions from best to worst. Explain your answers.

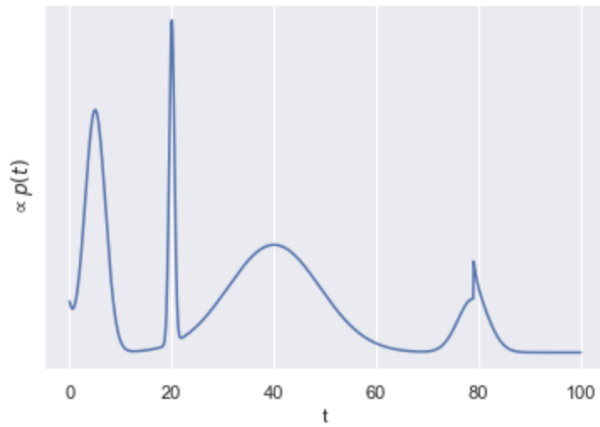

(a) A uniform distribution of width 1 centered at the current t (i.e. Uniform$[t-0.5, t+0.5]$)

(b) A shifted exponential distribution starting at the current t with $\lambda = 1$

(c) A normal distribution with mean at the current t and standard deviation 200

(d) A normal distribution with mean at the current t and standard deviation 40

4. In this problem, we'll use a Bayesian hierarchical model to model the number of failures, X_i, for each of n power plant pumps[1]. Consider the following Gamma-Poisson model,

$$\begin{aligned} \beta &\sim \text{Gamma}(m, \alpha) \\ \theta_i \mid \beta &\sim \text{Gamma}(k, \beta), \quad i = 1, \dots, n \\ X_i \mid \theta_i &\sim \text{Poisson}(\theta_i), \quad i = 1, \dots, n, \end{aligned}$$

where the θ_i are independent of each other and represent the rate of failures for each power plant pump. The parameters β and θ_i are unknown, and m, α, and k are fixed and known.

We'd like to infer the parameters β and θ_i from the data X. That is, we'd like to sample from the posterior distribution $\mathbb{P}(\beta, \theta \mid X)$. We will do so using Gibbs sampling.

(a) Draw a graphical model that represents the specified Gamma-Poisson model.

To run Gibbs sampling, we need to sample each parameter conditional on the current values of the other parameters. We'll do this in the next two steps.

(b) *Conditional Distribution of* β. What is the distribution $\mathbb{P}(\beta \mid \theta_{1:n}, X_{1:n})$? (*Hint:* The answer lies within a common distribution family.)

(c) *Conditional Distribution of* θ_i. What is the conditional distribution $\mathbb{P}(\theta_i \mid \beta, \theta_{1:i-1}, \theta_{i+1:n}, X_{1:n})$? (*Hint:* This also lies within a common distribution family.)

[1] E I George, U E Makov, and A F M Smith. Conjugate likelihood distributions. Scandinavian Journal of Statistics, 20:147–156, 1993.

(d) Using the results from the last two parts, write out pseudocode for Gibbs sampling from the posterior distribution.

5. (Challenge Question)

(a) Consider a Markov chain over the integers $\{1, \ldots, n\}$ such that i transitions to $\max(1, i-1)$ or $\min(n, i+1)$ with equal probability. How fast does the mixing time grow as a function of n? (I.e., is it $O(n)$, $O(n^2)$, $\ldots$?)

(b) Construct a Markov chain over $\{1, \ldots, n\}$ whose mixing time grows exponentially in n.

Feedback Form

On a scale of 1-5, where 1 = much too slow and 5 = much too fast, how was the pace of the discussion section?

1 2 3 4 5

Which problem(s) did you find most useful?

Which were least useful?

Any other feedback?