

# DS 102 Discussion 5

Monday, March 2, 2020

The past few lectures have looked at how to perform Bayesian inference using Markov Chain Monte Carlo sampling methods, such as Gibbs sampling. Today, we'll practice deriving one iteration of Gibbs sampling for the Gamma-Poisson model.

The goal of Bayesian inference is to get the posterior distribution of the parameters given the data,  $\mathbb{P}(\theta \mid X)$ . Often this is difficult to derive in closed form, so instead we'll try to sample from it. However, sampling from the posterior over *all* the parameters  $\theta$  is also often difficult. The main insight behind Gibbs sampling is that it can be much easier to sample the posterior over just a *single* parameter,  $\mathbb{P}(\theta_i \mid X, \theta_{-i})$  (where we use the index  $-i$  to mean all indices except for  $i$ ). Gibbs sampling then iterates through each parameter  $\theta_i$  and samples from  $\mathbb{P}(\theta_i \mid X, \theta_{-i})$ . This loop is repeated, each time conditioning on the newly sampled values.

Iterating through each parameter  $\theta_i$  and sampling from  $\mathbb{P}(\theta_i \mid X, \theta_{-i})$  is not the same thing as sampling from  $\mathbb{P}(\theta \mid X)$ . However, the good news is that given enough iterations, the former converges to the latter.

## 1. Gibbs sampling for Gamma-Poisson model.

Consider the hierarchical Bayes model where

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

where the  $\theta_i$  are independent of each other and the  $X_i$  are independent of each other. The  $\beta$  and  $\theta_i$  are unknown parameters, and  $m$ ,  $\alpha$ , and  $k$  are fixed and known.

We'd like to infer the parameters  $\beta$  and the  $\theta$  from the data  $X$ . That is, we'd like to sample from the posterior distribution  $\mathbb{P}(\beta, \theta \mid X)$  using Gibbs sampling. This entails deriving the posterior of each parameter, conditioned on the data and all the other parameters.

- We'll start with  $\beta$ . Derive  $\mathbb{P}(\beta \mid \theta_1, \dots, \theta_n, X_1, \dots, X_n)$ .
- Next, we'll look at each  $\theta_i$ . Derive  $\mathbb{P}(\theta_i \mid \beta, \theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n, X_1, \dots, X_n)$
- Using the posteriors you derived in the last two parts, write out the algorithm for the Gibbs sampler.