02 Bayesian Data Analysis exercises

Date: September 3, 2019

1 Deriving Bayes' rule

Let A and B be two observable events. P(A) is the probability that A occurs, and P(B) is the probability that B occurs. $P(A|B)$ is the conditional probability that A occurs given that B has happened. $P(A, B)$ is the joint probability of A and B both occurring.

You are given the definition of conditional probability:

$$P(A|B) = \frac{P(A,B)}{P(B)} \text{ where } P(B) > 0 \tag{1}$$

Using the above definition, and using the fact that $P(A, B) = P(B, A)$ (i.e., the probability of A and B both occurring is the same as the probability of B and A both occurring), derive an expression for $P(B|A)$. Show the steps clearly in the derivation.

2 Conjugate forms 1

1. Suppose you are given a vector of data x consisting of 1's and 0's, coming from a Binomial(n,θ) distribution. 1 represents success, and 0 failure. Example data are shown below, generated with probability of success $\theta = 0.5$, just for illustration:

   ```
   ## data:
   x<-rbinom(n=10,size=1,prob=0.5)
   x

   ##  [1] 0 1 0 1 0 0 0 0 0 1

   ## k:
   sum(x)

   ## [1] 3
   ```

   Here, n represents the number of trials, and k the number of successes. **The above code and output is just an example, and is no longer relevant for the question below.**

Given k successes in n trials coming from a Binomial distribution, we define a Beta(a,b) prior on the parameter θ.

Write down the Beta distribution that represents the posterior, in terms of a,b, n, and k.

2. We ask 10 yes/no questions from a participant, and the participant returns 0 correct answers. We assume a Binomial likelihood function for these data. Also assume a Beta(1,1) prior on the parameter θ, which represents the probability of success. Use the result you derived in the preceding part of this question to write down the posterior distribution of the θ parameter.

3 Conjugate forms 2

Suppose you have n independent and identically distributed data points from a distribution that has the likelihood function $f(x|\theta) = \theta(1-\theta)^{\sum_{i=1}^{n} x_i}$, where the data points x can have values 0,1,2,. . . . Let the prior on θ be Beta(a,b), a Beta distribution with parameters a,b. The posterior distribution is a Beta distribution with parameters a* and b*. Determine these parameters in terms of a, b, and $\sum_{i=1}^{n} x_i$.

4 Conjugate forms 3

The Gamma distribution is defined in terms of the parameters a, b: Ga(a,b). The probability density function is:

$$Ga(a,b) = \frac{b^a \lambda^{a-1} \exp\{-b\lambda\}}{\Gamma(a)} \tag{2}$$

We have data $x_1, \ldots, x_n$, with sample size n that is exponentially distributed. The exponential likelihood function is:

$$f(x_1, \ldots, x_n; \lambda) = \lambda^n \exp\{-\lambda \sum_{i=1}^{n} x_i\} \tag{3}$$

It turns out that if we assume a Ga(a,b) prior distribution and the above likelihood, the posterior distribution is a Gamma distribution. Find the parameters a' and b' of the posterior distribution.