APPM 5720: Computational Bayesian Statistics

Final Exam Review Problems

1. Suppose that $X_1, X_2, \ldots, X_n \stackrel{iid}{\sim} Poisson(\lambda)$. Suppose that λ has a $\Gamma(\alpha, \beta)$ prior. Find the posterior distribution for λ.

2. Suppose that $X_1, X_2, \ldots, X_n \stackrel{iid}{\sim} Bernoulli(\theta)$ and that θ has a $Beta(a, b)$ prior.
   (a) Find the posterior Bayes estimator of θ. Show that it is a weighted average of the sample mean and the prior mean.
   (b) Find the predictive distribution for X_{n+1}. Give a point estimate of X_{n+1} given the previous data.

3. Wishart Distribution:
   (a) Define the Wishart distribution. Don't give the pdf. What I am looking for is something like:

   Let $\vec{X}_1, \vec{X}_2, \ldots, \vec{X}_n \stackrel{iid}{\sim} MVN_p(\vec{0}, V)$ for $n \geq p$. etc....

   (b) What is the Wishart distribution used for in Bayesian statistics?

4. The lifetime X of a machine component, in days, is known to have an exponential distribution with rate λ. Suppose that λ is modelled as having an exponential rate 2 prior distribution. Suppose that, for a random sample of 5 components, the total lifetime is observed to be 3 days.
   (a) Find the posterior distribution for λ. Give a sketch of the pdf. On your sketch, draw symbolic endpoints of a 90% credible interval for λ. (i.e. You would really need to do this numerically, but don't.)
   (b) Indicate on your sketch how you would find a 90% highest posterior density region for λ. Does this correspond to the shortest 90% credible interval? Explain.
   (c) Find the predictive density for a 5th machine component. Explain how you might used simulation to estimate the probability that this component will have a lifetime of at least 0.25 days.

5. Suppose that $X_1, X_2, \ldots, X_n \stackrel{iid}{\sim} Bernoulli(\theta)$ and that θ has a $Beta(a, b)$ prior.
   (a) Find the Bayes factor for comparing the two models

   $$\begin{array}{rcl} M_1 & : & \theta = 1/2 \\ M_2 & : & \theta \neq 1/2 \end{array}$$

   (b) Find the Bayes factor for comparing the two models

   $$\begin{array}{rcl} M_1 & : & \theta \leq 1/2 \\ M_2 & : & \theta > 1/2 \end{array}$$

   (c) Without going into the specific cutoffs we discussed for Bayes factors, fill, which model is generally supported by a large Bayes factor?

6. Assume that you have n measurements of a response variable Y, denoted by $y_1, y_2, \ldots, y_n$. Further assume that you have n measurements of two different possible predictor variables X_1 and X_2, denoted $x_{11}, x_{12}, \ldots, x_{1,n}$ and $x_{21}, x_{22}, \ldots, x_{2n}$.

   Suppose that the response variable is thought to be related to each of the predictor variables as follows.

   $$y_i = \beta_1 x_{1i} + \varepsilon_{1i}, \quad i = 1, 2, \ldots, n$$

   or

   $$y_i = \beta_2 x_{2i} + \varepsilon_{2i}, \quad i = 1, 2, \ldots, n$$

   where $\varepsilon_{1i} \stackrel{iid}{\sim} N(0, \sigma_1^2)$ and $\varepsilon_{2i} \stackrel{iid}{\sim} N(0, \sigma_2^2)$ are independent.

   These are the only two models you will consider. Call them M_1 and M_2.

   In what follows, do no actual computation and stay pretty "generic". For example, your answers should include terms like $f(\vec{y}|\beta_j, \sigma_j^2)$ and not the actual normal pdf written out.

   (a) What priors need to be set up?

   (b) Write down an integral expression for the likelihood of the data given a particular model. (That is, write down an integral expression for $f(\vec{y}|M_j)$.)

   (c) Write down an expression for the posterior model probabilities given the data.

   (d) Write down an expression for the posterior odds ratio for model 1 versus model 2. What does a large posterior odds ratio generally support?

   (e) Suppose that you make a additional x observations. Further suppose that you decide to go with model 1. Write down an expression for your prediction of a new response variable y_{n+1} given all previous data.

   (f) Suppose that you do not choose one model over the other. Give a model averaging approach to the prediction of y_{n+1}.

7. What is the invariance problem that the Jeffreys prior aims to solve?

8. Let $X_1, X_2, \ldots, X_n \stackrel{iid}{\sim} f(x|\theta)$. Assume a prior $f(\theta)$ for θ. Show that the Bayes rule under squared error loss is the posterior Bayes estimator.

9. Suppose that X has a Poisson distribution with rate parameter λ. In what follows, use squared error loss and an consider decision rules of the form $\delta(X) = cX$.

   (a) Calculate the frequentist risk, $R_\delta(\lambda)$.

   (b) Show that δ is inadmissible if $c > 1$.

   (c) Is there a minimax decision rule? If so find it. If not, explain.

   (d) Find the Bayes rule using an exponential rate 1 prior for λ.

   (e) Find the PBE (posterior Bayes estimator) for λ. Note that it is not the Bayes rule. Explain why this does not contradict the result of Problem 8.

This should get you started– a few extra problems will be coming this weekend!