STA402L: HOMEWORK 6

DUE: 11:59 PM ON WEDNESDAY, APRIL 22

Instructions. Solutions must be submitted to Gradescope as a single PDF. Programming exercises must be completed in R, should be clearly presented, and include all R code. Lab questions are restated here for convenience, but you should refer to the lab itself for details.

Total points. Book exercises: 8; Lab exercises 20; Overall: 28.

BOOK EXERCISES

B1. (5 points) Suppose we consider two Bayesian linear regression models:

$$m_1 : Y = \beta_0 + \beta_1 x + \varepsilon$$
$$m_2 : Y = \beta_0 + \varepsilon.$$

Assume the posterior model probabilities are

$$p(m_1 \mid y, x) = 0.7 \quad \text{and} \quad p(m_2 \mid y, x) = 0.3,$$

and suppose the posterior predictive means for a new observation at x_{n+1} are

$$\mathbb{E}[y_{n+1} \mid x_{n+1}, m_1, y, x] = 10 \quad \text{and} \quad \mathbb{E}[y_{n+1} \mid x_{n+1}, m_2, y, x] = 8.$$

(a) Write the Bayesian model averaging formula for $\mathbb{E}[y_{n+1} \mid y, x]$.

(b) Compute the Bayesian model averaged prediction.

(c) Briefly explain why model averaging can be preferable to selecting a single model.

B2. (3 points) Consider three predictors x_1, x_2, and x_3. Suppose we fit Bayesian regression models corresponding to all subsets of predictors and obtain the following posterior model probabilities:

Model	Posterior probability
x_1	0.20
x_2	0.10
x_3	0.05
x_1, x_2	0.25
x_1, x_3	0.15
x_2, x_3	0.10
x_1, x_2, x_3	0.10
Null model (intercept only)	0.05

Compute the posterior inclusion probability for each predictor. Which predictor appears most important?

Lab exercises

L1. (2 points) Interpret the significant coefficients (at the 0.05 significance level).

L2. (2 points) How do the estimated coefficients compare in this glm model to those from the model fit using 'lm'? You can use the 'summary' function on the 'stan.glm' object to see the posterior summaries.

L3. (2 points) How do the credible intervals and standard errors of the coefficients compare to the confidence intervals and standard errors from the model fit using 'lm'?

L4. (2 points) Interpret the significant coefficients (at the 0.05 significance level).

L5. (2 points) What do our choice of priors say about our beliefs? How do we interpret these normal priors?

L6. (2 points) How do the estimated coefficients compare in this model to those from the model fit using 'glm'?

L7. (2 points) How do the credible intervals and standard errors of the coefficients compare to the confidence intervals and standard errors from the model fit using 'glm'?

L8. (2 points) Which model is better? Why?

L9. (2 points) How does posterior inference for the coefficients compare to when we used the weakly informative Normal prior above?

L10. (2 points) How do the two models compare in terms of predictive performance? Consider using the 'loo' function as we have been doing.