Solutions to Computer Exercise 8 (R)

1.

For the weak prior, there is autocorrelation for the `Type:Pot` variance component, and also some for the `Pot` component, which makes us a bit unhappy with this analysis (but increasing the `nitt`, `burnin`, and `thin` sampling parameters fixes these problems). For the stronger prior, there is no problem with with autocorrelation. From the posterior distributions one can detect an influence of the prior, in particular on the `Type:Pot` component, which is larger for the stronger prior. Most likely this will not influence our conclusions about the fixed effects, but it is not really ideal (so it is preferable to increase the sampling parameters). Having very strong priors that "make" the variance components small can have the effect of giving an underestimate of the uncertainty in fixed effect estimates (it corresponds to underestimating the "noise" that is compared with the "signal"). As we saw in exercise 6, it could be better just to drop the `Type:Pot` random effect from the model, which is supported by a lower DIC value. Because the `Type:Pot` interaction can be ignored, the difference between the two weed-types is consistent among pots. We can also note that the variance component for `Pot` is smaller than the residual variance (compare the Pot and units components in the summary output). If we would analyze a model without the `Pot` random effect (perhaps you can try this), that model has a larger DIC, so we should keep the `Pot` effect in the model. Finally, there is an effect of Type ($p < 0.001$), with Type N having a log length-width ratio that is about 0.31 greater than Type G.

2.

a) Yes, there is strong evidence for sexual size dimorphism: female legs measure on average 6.58, male legs are 0.83 (0.79-0.86) units longer. You have to change the prior quite a lot to affect estimates of the random effects in an important way that would influence the estimates of the fixed effect, because the data set here is large, and the difference in body size between males and females is large.

b) The overall genetic variance among lines explains 8.5% (3.7-14.7%) of the total variance in size within the sexes. However, there is about an equal amount of genetic variance for sexual size dimorphism, explaining 8.3% (4.2-13.4%). The rest of the variance (100-8.5-8.3 = 83.2%) is error variance.

c) If one compares these estimates and their confidence limits between the models using different priors one finds that they are very similar, thus changing the prior in this case did not dominate the data too much. It makes sense because it is a big dataset. For this large dataset, results produced by lmer() are also very similar.

3. Increasing the thinning interval and burn-in time helped to decrease the correlation between our stored estimates for the random effects. This is obvious from plots of autocorrelations and by looking at the number of efficient (independent) samples of the random effects in the model output. There are still some problems though, so increasing the strength of the prior in combinations with increasing the number of simulations and thinning interval may be a solution. If we are to trust this model, both the main effect of temperature and the interaction between sex and temperature is highly significant ($p < 0.001$). Females decrease in size by 0.49 units when developing in 31˚C whereas males decrease 0.80 units (0.49+0.31). Even in relative terms this is a larger decrease for the males. Perhaps males that grow faster and attain larger sizes under normal circumstances are more sensitive to the quality of the environment and cannot maintain their, relative to females, faster growth in 31˚C.