Stat 534: Fall 2021. Exercises from the Intro to rjags document.

This document extracts the exercises from the Intro to rjags document:

1. Download and save in an R working directory the grizzly.csv data file, the grizz0.r R code and the lreg.bug BUGS code. Run the grizz0.r code and look at the various bits of output. Ask if you don't understand something or want something additional.

2. Re model 0b, the state-space linear regression with measurement error model.

   The bugs and r code to fit this model are in lregb.bug and grizz0b.r. The r code is almost exactly the same as that in grizz0.r Make sure you understand how lregb.bug works.

   What is the 95% credible interval for r?

3. Re model 1, the stat-space model with process error.

   The bugs and r code to fit this model are in exp.bug and grizz1.r. The first part of the r code is almost exactly the same as that in grizz0.r. One difference is that we also ask jags to return the posterior distributions for each N_t. That way we can easily plot the distribution of $\hat{N}(t)$ over time.

   What is the 95% credible interval for r?

   Plot the distribution of $\hat{N}(t)$ over time.
   I suggest a boxplot for each time. An alternative is the mean and +/- 1 se bars.

Optional extensions: After fitting the exponential growth model with process error and being (at least somewhat) comfortable with the output, feel free to modify the code (grizz1.r, exp.bug, or both) so that:

4. Observations conditional on the latent population size have a Poisson distribution
   The BUGS specification for a Poisson distribution is dpois(mean), where mean is the expected value for that observation.

5. Observations conditional on the latent population size have an overdispersed Poisson distribution.

6. The process model allows the population growth rate to vary smoothly over time, i.e., a local linear trend model.

7. The process model allows the population growth rate to depend on current population size, i.e., a Ricker density-dependence model or something like it.