

Workshop 11: Introduction to JAGS

This workshop will introduce [JAGS](#). Before proceeding, you will need to install [JAGS](#) which is available [here](#). Then you'll need to install and load the [rjags](#) and [R2jags](#) libraries.

You can download individual files as you go, or download all the files you will need at once [here](#). Solutions are available here: [lions](#), [repeatability](#), [mouse](#).

Getting familiar with [JAGS](#) (Lions)

Let's start by fitting a simple linear model in [JAGS](#). We're going to use the lion nose color data-set that we analyzed in our Linear Regression workshop (see that workshop for details about the data). The script posted [here](#) contains everything you need to run the model, except the construction of the data and the likelihood for the [JAGS](#) analysis.

1. The first step in running an analysis in [JAGS](#) is to construct a list containing all the data that you will pass in to the model. The items in this list should contain names. For example, if you wanted to pass in `a=5` and `b=1:10` as data, you would create the list as:

```
my.data <- list(a=5, b=1:10)
```

Construct the data you will pass in to [JAGS](#). Your list should contain three things.

2. Fill in the likelihood and run the model.
3. Use the relevant coefficients from the summary which you can access with the following (assuming that your output from [JAGS](#) is stored in an object names `res`):

```
res$BUGSoutput$summary
```

to plot the model fit line (use a different colour).

4. Use the [traceplot](#) command to generate a traceplot for the `b.black` parameter. Do your chains look well-mixed?
5. You can extract the chains manually using:

```
res$BUGSoutput$sims.array[, , 'b.black']
```

See if you can plot the three chains manually using (the [matplot](#) command might be helpful).

6. Generate a posterior distribution using these chains.
7. See if you can manually calculate the mean and 95% BCI from these chains (you might find the `quantile` function helpful).

Repeatability of a sexual signal trait

Next, we will fit a model with a random effect. To do this, we'll re-do the analysis we did using the data-set from the “Repeatability of a sexual signal trait” example (Mixed Models workshop; see that workshop for details about the data). Download the [script](#).

1. Run the analysis in [JAGS](#) using the provided script. This model should fit an intercept only. You will likely get warnings indicating that you have provided [JAGS](#) with data that was not used. This is because our data contained [bird](#) and [nbird](#), neither of which we used (we'll need these for the random effect).
2. Next, we will incorporate a random effect of bird. To do this, you will need to give each bird its own intercept. The most natural way to do this is to create a [tau.phi.bird](#) variable, the same way we create [tau.phi.residual](#), and then use this to draw [nbird](#) values from a normal distribution with [dnorm\(0, tau.phi.bird\)](#). These [nbird](#) values could be indexed in a vector [b.bird\[j\]](#), where [j](#) goes from [1](#) to [nbird](#). Run the model with this random effect. Compare the estimate for [sigma.phi.bird](#) to the estimated standard deviation for the bird random effect from [lme4](#). Do they agree? *Note: Don't forget to track the new parameters you create!*

Mouse mis-match

Here we will go further with the mouse example we covered in lecture. The [script](#) posted here runs the analysis, as presented in lecture.

1. Create a plot of the posterior distribution for one of the parameters.
2. Color this histogram so that the 95% BCI is visible.
3. The goal of this study was to test whether survival differed for different mouse phenotypes on different soil types. Create four new variables inside the [JAGS](#) model that correspond to the survival probability of each mouse phenotype on each soil type. Track these new variables and re-run the model. These new variables should now show up in your summary table (so you should be able to see the mean and 95% BCI for each of them). Plot these four means (\pm their BCIs). What do these tell you about survival of mouse phenotypes on each of the soil types?
4. This existing analysis does not account for the fact that mice in the same enclosure might have similar rates of survival for reasons that are specific to that enclosure. The enclosure a mouse is in is already part of the data passed in to [JAGS](#)

```
my.data$site
```

as is the number of sites

```
my.data$nsite
```

Incorporate a random effect of [site](#) into the survival component of the model. You should be able to follow an almost identical set of steps as you did in the previous example. Did including this random effect change any of your earlier parameter estimates?

5. Evaluate the convergence of the new parameter that corresponds to the standard deviation of your random effect.