BUGS
Bayes Using Gibbs Sampling

Goal: Posterior Inference

Posterior proportional Likelihood*Priors

Normalizing constant is a @#*!$

BUT, Conditional Distributions are a breeze!
Easy to generate samples from the conditional distributions

Gibbs Sampling

Let $(\theta_1, \theta_2, \ldots \theta_k) = (\alpha_1, \ldots, \alpha_{16}, \beta_1, \ldots, \beta_{16}, \alpha_s, \beta_s, \tau, \tau_\alpha, \tau_\beta)$

$$p(\theta_1 | \theta_2, \theta_3, \ldots, \theta_k, Y)$$
$$p(\theta_2 | \theta_1, \theta_3, \ldots \theta_k, Y)$$
$$\vdots$$
$$p(\theta_k | \theta_{(k)}, Y)$$

- *Full* Conditional Distributions
- Cycle thru drawing from each distribution to obtain a complete vector
- Repeat a large number of times...

Theory

- This defines a Markov Chain *(continuous version)*
- Full conditional distributions define transition probabilities for moving from one state of the chain to another at iteration *t*
- The stationary distribution *(limiting distribution as t goes to infinity)* of the Markov Chain is the ***joint posterior distribution***
- For large *t*, samples from the Markov Chain can be viewed as a sample (approximately) from the desired posterior distribution

BUGS

- You specify
  - Model file
  - Data file
  - Initial Starting Values
- Model Menu: Specification
  - Double click *model* in model file, then click **Check Model**
  - Double click *list* in data file, then click **Load Data**
  - Double click *model,* then click **Compile**
  - Double click *list*, then click **Load Inits**
  - Click **Gen Inits** if there are still uninitialized nodes
- Model Menu: Updates
- Inference Menu: Samples

Creating Data Files

- Rectangular format *(all variable must be same length)*
- List format *(useful for objects of varying length)*
  - **`dput(list(MERCURY=fish$MERCURY, LENGTH=fish$LENGTH, S=fish$STATION, N=dim(fish)[1]), file="fish-R.dat")`**
  - Do not add variables not referenced in model!
  - Same format works for creating file of starting values!

Initial Starting Values

- Sometimes can generate form model (with informative prior distributions)
- Use MLEs to start
- Fit linear fixed effects model (same parameterization!)
  - Use MLEs of alpha, beta, tau = 1/sigma^2
  - Second stage:
    - mean(alpha), mean(beta)
    - 1/var(alpha), 1/var(beta)

OUTPUT

- Inference Menu:Samples
- CODA button: Creates output files that can be analyzed with
  - CODA
  - BOA
  - R/Splus
- Two files:
  - CODA Index: order of variables and indices in output
  - CODA Chain #: two columns, index and value of variables strung out in to one long vector.
  - See example for reading into R/Splus