

1 An Ecological Modeler's Primer on JAGS

2 N. Thompson Hobbs

3 May 19, 2014

4 Natural Resource Ecology Laboratory, Department of Ecosystem Science and
5 Sustainability, and Graduate Degree Program in Ecology, Colorado State University,
6 Fort Collins CO, 80523

7 Contents

8	1 Aim	4
9	2 Introducing MCMC Samplers	4
10	3 Introducing JAGS	5
11	4 Installing JAGS	8
12	4.1 Mac OS	8
13	4.2 Windows	8
14	4.3 LINUX	8
15	5 Running JAGS	9
16	5.1 The JAGS model	9
17	5.2 Technical notes	10
18	5.2.1 The model statement	10
19	5.2.2 for loops	11
20	5.2.3 Specifying priors	14
21	5.2.4 The <- operator	14
22	5.2.5 Vector operations	14
23	5.2.6 Keeping variables out of trouble.	15
24	5.3 Running JAGS from R	15
25	6 Output from JAGS	21
26	6.1 coda objects	21
27	6.1.1 Summarizing coda objects	21
28	6.1.2 The structure of coda objects (MCMC lists)	22
29	6.1.3 Manipulating coda objects	24
30	6.2 JAGS objects	25

31	6.2.1	Why another object?	25
32	6.2.2	Summarizing the JAGS object	26
33	6.2.3	The structure of JAGS objects (MCMC arrays)	27
34	6.2.4	Manipulating JAGS objects	28
35	6.2.5	Converting JAGS objects to coda objects	30
36	7	Which object to use?	30
37	8	Checking convergence using the coda package	30
38	8.1	Trace and density plots	31
39	8.2	Gelman and Rubin diagnostics	31
40	8.3	Heidelberger and Welch diagnostics	32
41	8.4	Raftery diagnostic	32
42	9	Monitoring deviance and calculating DIC	33
43	10	Differences between JAGS and WinBUGS / OpenBUGS	34
44	11	Troubleshooting	34
45	12	Answers to exercises	36
46		Literature Cited	40

1 Aim

The purpose of this Primer is to teach the programming skills needed to estimate the marginal posterior distributions of parameters and derived quantities of interest in ecological models using software implementing Monte Carlo Markov chain methods. Along the way, I will reinforce some of the ideas and principals that we have learned in lecture. The Primer is organized primarily as a tutorial and contains only a modicum of reference material ¹. There is an important supplement to this primer, excised from the JAGS users manual, that covers functions and distributions.

2 Introducing MCMC Samplers

WinBugs, OpenBUGS, and JAGS are three systems of software that implement Markov chain Monte Carlo sampling using the BUGS language. BUGS stands for Bayesian Analysis Using Gibbs Sampling, so you can get an idea what this language does from its name. Imagine that you took the MCMC code you wrote for a Gibbs sampler and tried to turn it into an R function for building chains of parameter estimates. Actually, you know enough now to construct a very general tool that would do this. However, you are probably delighted to know that accomplish the same thing with less time and effort using the BUGS language.

The BUGS language is currently implemented in three flavors of software: OpenBUGS, WinBUGS, and JAGS. OpenBUGS and WinBUGS run on Windows operating systems, while JAGS was specifically constructed to run multiple platforms, including Mac OS and Unix. Although all three programs use essentially the same syntax, OpenBUGS and WinBUGS run in an elaborate graphical user interface, while JAGS only runs from the command line of a Unix shell or from R. However, all three can be easily called from R, and this is the approach I will teach. My experience is that that the GUI involves far to much tedious

¹Other good references on the BUGS language are the WinBUGS manual (<http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml>, look for the manual .pdf link) which has lots of detailed treatment of functions and syntax as well as McCarthy (2007). The JAGS manual can be a bit confusing because it is written as if you were going to use the software stand alone, that is, from a UNIX command line.

70 pointing and clicking and doesn't' provide the flexibility that is needed for serious work.

71 3 Introducing JAGS

72 In this course we will use JAGS, which stands somewhat whimsically for “Just another Gibbs
73 Sampler.” There are three reasons I have chosen JAGS as the language for this course. First
74 and most important, is because my experience is that JAGS is *far* less fussy than WinBUGS
75 (or OpenBUGS) which can be notoriously difficult to debug. Second is that JAGS runs
76 on all platforms which makes collaboration easier. Finally, JAGS has some terrific features
77 and functions that are absent from other implementations of the BUGS language. That
78 said, if you learn JAGS you will have no problem interpreting code written for WinBugs
79 or OpenBUGS (for example, the programs written in McCarthy 2007) . The languages are
80 almost identical except that JAGS is better.²

81 This tutorial will use a simple example of regression as a starting point for teaching the
82 BUGS language implemented in JAGS and associated R commands. Although the problem
83 starts simply, it builds to include some fairly sophisticated analysis. The model that we will
84 use is the a linear relationship between the per-capita rate of population growth and the the
85 size a population, which, as you know, is the starting point for deriving the logistic equation.
86 For the ecosystem scientists among you, this problem is easily recast as the mass specific rate
87 of accumulation of nitrogen in the soil; see for example, Knops and Tilman (2000). Happily,
88 both the population example and the ecosystem example can use the symbol N to represent
89 the state variable of interest. Consider the model,

$$\frac{1}{N} \frac{dN}{dt} = r - \frac{r}{K} N, \quad (1)$$

²There is also software called GeoBUGS that is specifically developed for spatial models, but I know virtually nothing about it. However, if you are interested in landscape ecology otherwise have an interest in spatial modeling, I urge you to look into it after completing this tutorial. The manual can be found at <http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml>

90 which, of course, is a linear model with intercept r and slope $\frac{r}{K}$. Note that these quantities
 91 enjoy a sturdy biological interpretation; r is the intrinsic rate of increase, $\frac{r}{K}$ is the strength of
 92 the feedback from population size to population growth rate, and K is the carrying capacity,
 93 that is, the population size (o.k., o.k., the gm N per gm soil) at which $\frac{dN}{dt} = 0$. Presume
 94 we have some data consisting of observations of per capita rate of growth of N paired with
 95 observations of N . The vector \mathbf{y} contains values for the rate and the vector \mathbf{x} contains aligned
 96 data on N , i.e., $y_i = \frac{1}{N_i} \frac{dN_i}{dt}$, $x_i = N_i$. To keep things simple, we start out by assuming that
 97 the x_i are measured without error. A simple Bayesian model specifies the joint distribution
 98 of the parameters and data as

$$\begin{aligned}
 \mu_i &= r - \frac{rx_i}{K} \\
 [r, K, \tau \mid \mathbf{y}] &\propto \prod_{i=1}^n [y_i \mid \mu_i, \tau] [r] [K] [\tau] \\
 [r, K, \tau \mid \mathbf{y}] &\propto \prod_{i=1}^n \text{normal}(y_i \mid \mu_i, \tau) \times \qquad \qquad \qquad (2) \\
 &\quad \text{gamma}(K \mid .001, .001) \text{gamma}(\tau \mid .001, .001) \text{gamma}(r \mid .001, .001),
 \end{aligned}$$

99 where the priors are uninformative. Note that I have used the precision (τ) as a argument
 100 to the normal distribution rather than the variance ($\tau = \frac{1}{\sigma^2}$) to keep things consistent with
 101 the code below. Now, I have full, abiding confidence that with a couple of hours worth of
 102 work, perhaps less, you could knock out a Gibbs sampler to estimate r , K , and τ . However,
 103 I am all for doing things nimbly in 15 minutes that might otherwise take a sweaty hour of
 104 hard labor, so, consider the code in algorithm 1, below.

105 This code illustrates the purpose of JAGS (and other BUGS software): to translate the
 106 numerator of Bayes theorem (a.k.a., the joint, e.g., equation 2) into a specification of an
 107 MCMC sampler. JAGS parses this code, sets up proposal distributions and steps in the
 108 Gibbs sampler and returns the MCMC chain for each parameter. These chains form the
 109 basis for estimating posterior distributions and associated statistics, i.e., means, medians,
 110 standard deviations, and quantiles. As we will soon learn, it easy to derive chains for other

111 quantities of interest and their posterior distributions, for example, $K/2$ (What is $K/2?$),
112 N as a function of time or dN/dt as a function of N . It is easy to construct comparisons
113 between of the growth parameters of two populations or among ten of them. If this seems
114 as if it might be useful to you, you should continue reading.

Algorithm 1 Linear regression example

```
##Logistic example for Primer
model{
  #priors
  K~dgamma(.001,.001)
  r~dgamma(.001,.001)
  tau~ dgamma(.001,.001) #precision
  sigma<-1/sqrt(tau) #calculate sd from precision
  #likelihood
  for(i in 1:n){
    mu[i] <- r - r/K * x[i]
    y[i] ~ dnorm(mu[i],tau)
  }
} #end of model
```

115 JAGS is a compact language that includes a lean but useful set of scalar and vector functions
116 for creating deterministic models as well as a full range of distributions for constructing the
117 stochastic models. The syntax closely resembles R, but there are differences and of course,
118 JAGS is far more limited. Detailed tables of functions and distributions can be found in
119 the supplementary material, JAGS functions and distributions.pdf, taken from the JAGS
120 manual (Plummer, 2011). Rather than focus on these details, this tutorial presents general
121 introduction JAGS models, how to call them from R, how to summarize their output, and
122 how to check convergence.

123 4 Installing JAGS

124 4.1 Mac OS

125 Update your version of R to the most recent one. Go to the package installer under Packages
126 and Data on the toolbar and check the box in the lower right corner for install dependencies.
127 Install the `rjags` package from a CRAN mirror of your choice. Now go to [http://www-ice.
128 iarc.fr/~martyn/software/jags/](http://www-ice.iarc.fr/~martyn/software/jags/) and look in the section under downloads. Click on the
129 files page link and then click on Download JAGSdist.dmg (4.7 MB) where `_____` is the
130 number of the latest version to get the disk mounting image. Install as you would any other
131 Mac software.

132 4.2 Windows

133 Update your version of R to the most recent one. Go to the package installer under Packages
134 and Data on the toolbar and check the box in the lower right corner for install dependencies.
135 Install the `rjags` package from a CRAN mirror of your choice. Check the version number
136 of `rjags`. Go to <http://sourceforge.net/projects/mcmc-jags/files/JAGS/>. Click on
137 3x then JAGS-3.3.0.exe.

138 Occasionally, students using windows operating systems have problems loading `rgags`
139 from R after everything has been installed properly. In all cases I have encountered, this
140 problem occurs because they have more than one version of R resident on their computers
141 (wisely, Mac OS will not allow that). So, if you can't seem to get `rjags` to run after a proper
142 install, then uninstall all versions of R, reinstall the latest version, install the latest version
143 of `rjags` and the version of JAGS that matches it.

144 4.3 LINUX

145 There is a link to the path for binaries found at <http://mcmc-jags.sourceforge.net/>
146 . If you want to compile from source code, there are detailed instructions at <http://>

147 `yusung.blogspot.com/2009/01/install-jags-and-rjags-in-fedora.html` There are tar
148 files found at `http://sourceforge.net/projects/mcmc-jags/files/JAGS/3.x/Source/`.
149 You want `JAGS-3.0.0.tar.gz`. My guess is that you will need to download the `rjags`
150 package in R before installing JAGS.

151 Here is a note on using the Ubuntu Software Center, compliments of Jean Fleming:

152 “Elsie and I both use Ubuntu which is a specific linux distribution, it is one
153 of the more commonly used distributions (it is user friendly!) so it is likely that
154 many linux users in the future will be able to use this advice. If anyone does not
155 have Ubuntu they may need to use the steps you described in the primer.

156 I installed the `rjags` package following the directions in the primer. Ubuntu
157 comes with a Software Center where you can search for and download most open
158 source software, so to download and install JAGS I just opened up Software
159 Center, searched for JAGS, and installed it.”

160 5 Running JAGS

161 5.1 The JAGS model

162 Study the relationship between the numerator of Bayes theorem (equation 2) and the code
163 (algorithm 1). Although this model is a simple one, it has the same general structure as all
164 Bayesian models in JAGS:

- 165 1. code for priors,
- 166 2. code for the deterministic model,
- 167 3. code for the likelihood(s).

168 The similarity between the code and equation 2 should be pretty obvious, but there are a few
169 things to point out. Priors and likelihoods are specified using the \sim notation that we have

170 seen in class. For example, remember that

$$y_i \sim \text{normal}(\mu_i, \tau)$$

171 is the same as

$$\text{normal}(y_i \mid \mu_i, \tau).$$

172 So, it is easy to see the correspondence between the mathematical formulation of the model
173 (i.e., the numerator of Bayes theorem, equation 2) and the code. In this example, I chose
174 uninformative gamma priors for r , K and τ because they must be positive. I chose a normal
175 likelihood because the values of y and μ are continuous and can take on positive or negative
176 values.

177 **Exercise: always plot your priors** Plot priors for each parameter, scaling the x axis
178 appropriately for each value— r should be about .2, K about 1200, and τ should be about
179 2500. Discuss with you lab mates if $\text{gamma}(\theta \mid .001, .001)$ is vague for all parameters, i.e.,
180 $\theta = r, K, \tau$. Be sure to include lots of x points in your plots to get a good interpolation, at
181 least 1000.

182 5.2 Technical notes

183 5.2.1 The model statement

184 Your entire model must be enclosed in

```
185     model{  
186     .  
187     .  
188     .  
189     .  
190     } #end of model
```

191 I am in the habit of putting a hard return (a blank line) after the } #end of model state-
192 ment. If you fail to do so, you may get the message #syntax error, unexpected NAME,
193 expecting \$end. (This may have been fixed in the newer versions of JAGS, but just to
194 be safe....)

195 5.2.2 for loops

196 Notice that the for loop replaces the $\prod_{i=1}^n$ in the likelihood. Recall that when we specify an
197 *individual* likelihood, we ask, what is the probability (actually, probability density) that we
198 would obtain this data point conditional on the value of the parameter(s) of interest? The
199 total likelihood is the product of the individual likelihoods. Recall in the Excel example
200 for the light limitation of trees that you had an entire column of likelihoods adjacent to a
201 column of deterministic predictions of our model. If you were to duplicate these “columns”
202 in JAGS you would write

```
203     mu[1] <- r - r/K * x[1]
204     y[1] ~ dnorm(mu[1],tau)
205     mu[2] <- r - r/K * x[2]
206     y[2] ~ dnorm(mu[3],tau)
207     mu[3] <- r - r/K * x[3]
208     y[3] ~ dnorm(mu[3],tau)
209     .
210     .
211     .
212     mu[n] <- r - r/K * x[n]
213     y[n] ~ dnorm(mu[n],tau)
```

214 Well, presuming that you have something better to do with your time than to write out
215 statements like this for every observation in your data set, you may substitute

```

216     for(i in 1:n){
217         mu[i] <- r - r/K * x[i]
218         y[i] ~ dnorm(mu[i],tau)
219     }

```

220 for the line by line specification of the likelihood. Thus, the for loop specifies the elements
 221 in the product of the likelihoods.

222 Note however, that the for structure in the JAGS language is subtly different from what
 223 you have learned in R. For example the following would be legal in R but not in the BUGS
 224 language:

```

225     #WRONG!!!
226     for(i in 1:n){
227         mu <- r - r/K * x[i]
228         y[i] ~ dnorm(mu,tau)
229     }

```

230 If you write something like this in JAGS you will get a message that complains about multiple
 231 definitions of node mu. If you think about what the for loop is doing, you can see the reason
 232 for this complaint; the incorrect syntax translates to

```

233     #Wrong
234     mu <- r - r/K * x[1]
235     y[1] ~ dnorm(mu,tau)
236     mu <- r - r/K * x[2]
237     y[2] ~ dnorm(mu,tau)
238     mu <- r - r/K * x[3]
239     y[3] ~ dnorm(mu,tau)
240     .

```

```

241     .
242     .
243     mu <- r - r/K * x[n]
244     y[n] ~ dnorm(mu,tau),

```

245 which is *nonsense* if you are specifying a likelihood because μ is used more than once in a
246 likelihood for different values of y . This points out a fundamental difference between R and
247 the JAGS language. In R, a `for` loop specifies how to repeat many operations in sequence. In
248 JAGS a `for` construct is a way to specify a product likelihood or the distributions of priors
249 for a vector. One more thing about the `for` construct. If you have two product symbols
250 in the conditional distribution with different indices, that is $\prod_{i=1}^n \prod_{j=1}^m \dots$ then this dual
251 product is specified in JAGS using nested `for` loops, i.e.,

```

252     for(i in 1:n){
253         for(j in 1:m){
254             expression[i,j]
255         } #end of j loop
256     } #end of i loop

```

257 As an alternative to giving an explicit argument for the number of iterations (e.g., `n` and `m`
258 above), you can use the `length()` function. For example we could use

```

259     for(1 in 1:length(x[])){
260         mu[i] <- r - r/K * x[i]
261         y[i] ~ dnorm(mu[i],tau)
262     }

```

263 **Exercise: using for loops** Write a code fragment to set vague normal priors (`dnorm(0,10e-6)`)
264 for 5 regression coefficients stored in the vector `B`.

265 5.2.3 Specifying priors

266 We specify priors in JAGS as `parameter ~ distribution(shape1, shape2)`. See the sup-
267 plementary material for available distributions. Note that in the code (algorithm 1), the
268 second argument to the normal density function is `tau`, which is the precision, defined as
269 the reciprocal of the variance. This means that we must calculate `sigma` from `tau` if we
270 want a posterior distribution on `sigma`. Be very careful about this—it is easy to forget that
271 you must use the precision rather than the standard deviation as an argument to `dnorm` or
272 `dlnorm`. Failing to do this is a source of immense suffering. (I know from experience.) For
273 the lognormal, it is the precision on the log scale. If you would like, you can express priors
274 on σ rather than τ using code like this:

275

```
276     sigma~dunif(0,100) #presuming this more than brackets the posterior of sigma  
277     tau <- 1/sigma^2
```

278 There are times when this seems to work better than the gamma prior for `tau`.

279 5.2.4 The `<-` operator

280 Note that, unlike R, you do not have the option in JAGS to use the `=` sign in an assignment
281 statement. You must use `<-`.

282 5.2.5 Vector operations

283 I don't use any vector operations in the example code, but JAGS supports a rich collection
284 of operations on vectors. You have already seen the `length()` function—other examples in-
285 clude means, variances, standard deviations, quantiles, etc. See the supplementary material.
286 However, you cannot form vectors using syntax like `c()`. If you need a specific-valued vector
287 in JAGS, read it in as data.

288 5.2.6 Keeping variables out of trouble.

289 Remember that all of the variables you are estimating will be sampled from a broad range of
290 values, at least initially, and so it is often necessary to prevent them from taking on undefined
291 values, for example logs of negatives, divide by 0, etc. You can usually use JAGS' `max()`
292 and `min()` functions to do this. For example, to prevent logs from going negative, I often
293 use something like:

```
294     mu[i]<- log(max(.0001,expression))
```

295 **Exercise: Coding the JAGS script** Carefully write out all of the code in the Logistic
296 example (algorithm 1) into a program window in R. You may save this code in any directory
297 that you like and may name it anything you like. I use names like `logistic exampleJAGS.R`
298 which lets me know that the file contains JAGS code. Using an R extension allows me to
299 search these files easily with Spotlight.

300 5.3 Running JAGS from R

301 We implement our model using R (algorithm 2.) We will go through the R code step by
302 step. We start by bringing the growth rate data into R as a data frame. Next, we specify
303 the initial conditions for the MCMC chain in the statement `inits = . . .`. This is exactly the
304 same thing as you did when you wrote your MCMC code and assigned a guess to the first
305 element in the chain. There are two important things to notice about this statement.

Algorithm 2 R code for running logistic JAGS script.

```
setwd("/Users/Tom/Documents/Ecological Modeling Course/JAGS Primer")
rm(list=ls())
pop.data=(read.csv("Logistic Data II.csv"))
names(pop.data)=c("Year","Population Size", "Growth Rate")
inits=list(
list(K=1500, r=.2, tau=2500)
)#chain 1

n.xy = nrow(pop.data)
data=list(
  n=n.xy,
  x=as.double(pop.data$"Population Size"),
  y=as.double(pop.data$"Growth Rate")
)

library(rjags)
##call to JAGS
library(rjags)
##call to JAGS
n.adapt=5000
n.update = 10000
n.iter = 10000
jm=jags.model("Logistic example JAGS.R",data=data,inits,n.chains=length(inits),
n.adapt = n.adapt)
#Burnin the chain.
update(jm, n.iter=n.update)
#generate coda object
zm=coda.samples(jm,variable.names=c("K", "r", "sigma"), n.iter=n.iter, thin=1)
```

306 First, initial conditions must be specified as as “list of lists”, as you can see in the code.

307 If you create a single list, rather than a list of lists, i.e.,

```
308 inits= list(K=1500, r=.5, tau=2500) #WRONG
```

309 you will get an error message when you execute the `jags.model` statement and your code
310 will not run. Second, this statement allows you to set up multiple chains³, which are needed

³I start my work with a single chain. Once everything seems to be running, I add additional ones.

311 for some tests of convergence and to calculate DIC (more about these tasks later). For
312 example, if you want three chains, you would use something like:

```
313     inits=list(  
314     list(K=1500, r=.5, tau=1500), #chain 1  
315     list(K=1000, r=.1, tau=1000), #chain 2  
316     list(K=900, r=.3, tau=900) #chain 3  
317     ) #end of inits list
```

318 Now it is really easy to see why we need the “list of lists” format—there is one list for each
319 chain; but remember, you require the same structure for a single set of initial conditions,
320 that is, a list of lists.

321 Which variables in your JAGS code require initialization? Anything you are estimating
322 must be initialized, which means anything on the right hand side of a conditioning symbol
323 (except, of course, data) Think about it this way. When you were writing your own Gibbs
324 sampler, every chain required a value as the first element in the vector holding the chain.
325 That is what you are doing when you specify initial conditions here. You can get away
326 without explicitly specifying initial values—JAGS will choose them for you if you don’t specify
327 them—however, I strongly urge you to provide explicit initial values, particularly when your
328 priors are vague. This habit also forces you to think about what you are estimating.

329 The next couple of statements,

```
330     n.xy = nrow(pop.data)  
331     data=list(n=n.xy,  
332             x=as.double(pop.data$"Population Size"),  
333             y=as.double(pop.data$"Growth Rate"))
```

334 specify the data that will be used by your JAGS program. Notice that you can assign data
335 vectors on the R side to different names on the JAGS side. For example, the bit that reads

```
336     x=as.double(pop.data$"Population Size")
```

337 says that the x vector in your JAGS program (algorithm 1) is composed of the column in
338 your data frame called `Population Size` and the bit that reads

```
339     y=as.double(pop.data$"Growth Rate")
```

340 creates a y vector required by the JAGS program from the column in your data frame called
341 `Growth Rate` (pretty cool, I think). Notice that if I had named the variable `Growth.Rate`
342 instead of `Growth Rate`, the quotes would not be needed. It is important for you to un-
343 derstand that the left hand side of the = corresponds to variable name for the data in the
344 JAGS program and the right hand side of the = is what they are called in R. Also, note
345 that because `pop.data` is a data frame I used `as.double()`⁴ to be sure that JAGS received
346 real numbers instead of characters or factors, as can happen with data frames. This can
347 be particularly important if you have missing data in the data. The `n` is required in the
348 JAGS program to index the `for` structure (algorithm 2) and it must be read as data in
349 this statement⁵. By the way, you don't need to call this list "data"—it could be anything
350 ("apples", "bookshelves", "xy" etc.)

351 Now that you have a list of data and initial values for the MCMC chain you make calls
352 to JAGS using the following:

```
353     library(rjags)
354     ##call to JAGS
355     n.adapt=5000
356     n.update = 10000
357     n.iter = 25000
358     jm=jags.model("Logistic example JAGS.R",data=data,init=init,n.chains=length(init),
```

⁴This says the number is real and is stored with double precision, i.e., 64 bits in computer memory. This varies with the type of number being stored, but a good rule of thumb is that 16 decimal places can be kept in memory. This is usually sufficient for ecology!

⁵You could hard code the `for` index in the JAGS code, but this is bad practice.

```

359         n.adapt = n.adapt)
360
361     #Burnin the chain.
362     update(jm, n.iter=n.update)
363     #generate coda object
364     zm=coda.samples(jm,variable.names=c("K", "r", "sigma"), n.iter=n.iter, thin=1)

```

364 There is a quite a bit to learn here, so if your attention is fading, go get an espresso or come
365 back to this tomorrow. First, we need to get the library `rjags`. We then specify 3 scalars,
366 `n.adapt`, `n.update`, and `n.iter`. These tell JAGS the number of iterations in the chain
367 for adaptation (`n.adapt`), burn in (`n.udpate`) and the number to keep in the final chain
368 (`n.iter`). The first one, `n.adapt`, may not be familiar– it is the number of iterations that
369 JAGS will use to choose the sampler and to assure optimum mixing of the MCMC chain.
370 The second, `n.update`, is the number of iterations that will be discarded to allow the chain
371 to converge before iterations are stored (aka, burn in) . The final one, `n.iter`, is the number
372 of iterations that will be stored in the chain as samples from the posterior distribution–it
373 forms the “rug.”

374 The `jm=jags.model....` statement sets up the MCMC chain. Its first argument is the
375 name of the file containing the BUGS code. Note that in this case, the file resided in the
376 current working directory, which I specified at the top of the code (algorithm 2). Otherwise,
377 you would need to specify the full path name. (It is also possible to embed the BUGS code
378 within your R script, see Algorithm 3,). The next two expressions specify where the data
379 come from, where to get the initial values, and how many chains to create (i.e., the length
380 of the list `inits`). Finally, it specifies the “burn-in” how many samples to throw away before
381 beginning to save values in the chain. Thus, in this case, we will throw away the first 10,000
382 values.

383 The second statement (`zm=coda.samples...`) creates the chains and stores them as
384 an MCMC list (more about that soon). The first argument (`jm`) is the name of the jags
385 model you created in the `jags.model` function. The second argument (`variable.names`)

Algorithm 3 Example of code for inserting BUGS code within R script. This should be placed above the `jags.model()` statement (algorithm). You must remember to execute the code starting at `sink` and ending at `sink` every time you make changes in the model.

```
sink("logisticJAGS.R")
#This is the file name for the bugs code
cat(" model{

    K~dgamma(.001,.001)
    r~dgamma(.001,.001)
    tau~ dgamma(.001,.001)
    sigma<-1/sqrt(tau)
    #likelihood
    for(i in 1:n){
        mu[i] <- r - r/K * x[i]
        y[i] ~ dnorm(mu[i],tau)
    } #end of i for

} #end of model
",fill=TRUE)
sink()
```

386 tells JAGS which variables to “monitor.” These are the variables for which you want poste-
387 rior distributions. Finally, `n.iter=n.iter` says we want 25000 elements in each chain and
388 `n.thin` specifies how many of these to keep. For example, if `n.thin = 10`, we would store
389 every 10th element. Sometimes setting `n.thin > 1` is a good idea to reduce the size of the
390 data files that you will analyze.

391 **Exercise: Coding the logistic regression** Write R code (Algorithm 2) to use the JAGS
392 model to estimate the parameters, r , K and σ . When your model is running without error
393 messages, proceed to get output, as described below.

394 6 Output from JAGS

395 6.1 coda objects

396 6.1.1 Summarizing coda objects

397 The `zm` object produced by the statement

```
398   zm=coda.samples(jm,variable.names=c("K", "r", "sigma"), n.iter=n.iter,n.thin=1)
```

399 is a “coda” object, or more precisely, an MCMC list. Assuming that the coda library is
400 loaded [i.e. `library(coda)`], you can obtain a summary of statistics from MCMC chains
401 contained in a coda object using `summary(objectname)`. All of the variables in the
402 `variable.names=c()` argument to the `coda.samples` function will be summarized. For
403 the logistic example, `summary(zm)` produces:

```
404   Iterations = 15001:25000
405   Thinning interval = 1
406   Number of chains = 3
407   Sample size per chain = 10000
408   1. Empirical mean and standard deviation for each variable,
409      plus standard error of the mean:
410
411           Mean          SD Naive SE Time-series SE
412   K      1.313e+03 1.180e+02 6.811e-01      1.244e+00
413   r      1.998e-01 1.101e-02 6.359e-05      1.113e-04
414   sigma 2.538e-02 5.204e-03 3.004e-05      3.604e-05
415
416   2. Quantiles for each variable:
417
418           2.5%      25%      50%      75%      97.5%
419   K      1.125e+03 1.235e+03 1.300e+03 1.374e+03 1.583e+03
420   r      1.776e-01 1.928e-01 1.999e-01 2.070e-01 2.213e-01
421   sigma 1.759e-02 2.169e-02 2.460e-02 2.814e-02 3.773e-02
```

419 Each of the two tables above has the properties of a matrix⁶. You can output the cells of
420 these tables using syntax as follows. To get the mean and standard deviation of r,

```
421 > summary(zm)$stat[2,1:2]
422           Mean           SD
423 0.19980128 0.01101439
```

424 To get the upper and lower 95% quantiles on K,

```
425 > summary(zm)$quantile[1,c(1,5)]
426           2.5%           97.5%
427 1124.539 1582.647
```

428 **Exercise: Manipulating coda summaries** Build a table that contains the mean, stan-
429 dard deviation, median and upper and lower 2.5% CI for parameter estimates from the
430 logistic example. Output your table with 3 significant digits to .csv file readable by Excel
431 (hint, see the `signif()` function).

432 6.1.2 The structure of coda objects (MCMC lists)

433 So, what is a coda object? Technically, the coda object is an MCMC list. It looks like this:

```
434 [[1]]
435 Markov Chain Monte Carlo (MCMC) output:
436 Start = 60001
437 End = 60010
438 Thinning interval = 1
439           K           r           sigma
```

⁶Consider `m=summary(zm)`. The object `m` is a list of two matrices, one for the table of means and the other for the table of quantiles. As with any list, you can access these tables with `m[[1]]` and `m[[2]]` or the syntax shown above. Try it.

```

440 [1,] 1096.756 0.1914722 0.02889710
441 [2,] 1196.326 0.2088859 0.03155777
442 [3,] 1401.511 0.1804327 0.02553913
443 [4,] 1471.539 0.1754886 0.03589013
444 [5,] 1245.909 0.1567580 0.04248644
445 [6,] 1134.738 0.2114307 0.04151478
446 [7,] 1105.661 0.2303630 0.03141035
447 [8,] 1108.569 0.2169765 0.03708956
448 [9,] 1134.755 0.1964426 0.02660658
449 [10,] 1161.750 0.2152418 0.03700475
450 .
451 .
452 .

```

453 as many rows as you have thinned iterations

454 So, the output of coda is a list of matrices (or tables if you prefer) where each matrix contains
455 the output of the chains for each parameter to be estimated. Parameter values are stored in
456 the columns of the matrix; values for one iteration of the chain are stored in each row. So,
457 the example above is a case where we had 10 iterations of one chain. If we had 2 chains, 5
458 iterations each, the coda object would look like:

```

459 [[1]]
460 Markov Chain Monte Carlo (MCMC) output:
461 Start = 10001
462 End = 10005
463 Thinning interval = 1
464           K           r           sigma
465 [1,] 1070.013 0.2126878 0.02652204
466 [2,] 1085.438 0.2279789 0.02488036

```

```

467     [3,] 1170.086 0.2259743 0.02331958
468     [4,] 1094.564 0.2228788 0.02137309
469     [5,] 1053.495 0.2368199 0.03209893
470     [[2]]
471     Markov Chain Monte Carlo (MCMC) output:
472     Start = 10001
473     End = 10005
474     Thinning interval = 1
475           K           r           sigma
476     [1,] 1137.501 0.2657460 0.04093364
477     [2,] 1257.340 0.1332901 0.04397191
478     [3,] 1073.023 0.2043738 0.03355776
479     [4,] 1159.732 0.2339060 0.02857740
480     [5,] 1368.568 0.2021042 0.05954259
481     attr(,"class")
482     [1] "mcmc.list"

```

483 **Exercise: Understanding coda objects:** Modify your code to produce a coda object with
484 3 chains called `zm.short`, setting `n.adapt = 500`, `n.update=500`, and `n.iter = 20`.

- 485 1. Output the estimate of σ for the third iteration from the second chain.
- 486 2. Output all of the estimates of r from the first chain.
- 487 3. Verify your answers by printing the entire chain, i.e. enter `zm.short` at the console.

488 6.1.3 Manipulating coda objects

489 Any coda object can be converted to a data frame using syntax like

```
490     df = as.data.frame(rbind(co[[1]], co[[2]], ....co[[n]]))
```


491 where `df` is the data frame, `co` is the coda object and `n` is the number of chains in the coda
492 object, that is, the number of elements in the list. Once the coda object has been converted to
493 a dataframe, you can use any of the R tricks you have learned for manipulating data frames.
494 The thing to notice here is the double brackets, which is how we refer to the elements of a
495 list. Think about what this statement is doing.

496 **Exercise:** Convert the `zm` object to a data frame. Using the elements of data frame (not
497 `zm`) as input to functions:

- 498 1. Find the maximum value of σ .
- 499 2. Estimate the mean of r for the first 1000 and last 1000 iterations in the chain.
- 500 3. Produce a publication quality plot of the posterior density of K .
- 501 4. Estimate the probability that the parameter K exceeds 1600. (Hint: Look into using
502 the `ecdf()` function.) Estimate the probability that K falls between 1000 and 1300.

503 6.2 JAGS objects

504 6.2.1 Why another object?

505 The coda object is strictly tabular—it is a list of matrices where each element of the list an
506 MCMC chain with rows holding iterations and columns holding values to be estimated. This
507 is fine when the parameters you are estimating are entirely scalar, but sometimes you want
508 posterior distributions for all of the elements of vectors or for matrices and in this case, the
509 coda object can be quite cumbersome. For example, presume you would like to get posterior
510 distributions on the *predictions* of your regression model. To do this, you would simply ask
511 JAGS to monitor the values of `mu` by changing your `coda.samples` statement to read:

```
512   zm=coda.samples(jm,variable.names=c("K", "r", "sigma", 'mu'),  
513   n.iter=n.iter, n.thin=1)
```

514 **Exercise: vectors in coda objects:** Modify your code to include estimates of μ and
515 summarize the coda object. What if you wanted to plot the model predictions with 95%
516 credible intervals against the data. How would you do that?

517 6.2.2 Summarizing the JAGS object

518 As an alternative, replace `coda.samples` function with

```
519     zj=jags.samples(jm,variable.names=c("K", "r", "sigma","mu"),  
520     n.iter=n.iter, n.thin=1)
```

521 If you run this and enter `zj` at the console, R will return the means of all the monitored
522 variables⁷. Try it. If you want other statistics, you would use syntax like:

```
523     summary(zj$variable.name,FUN)$stat
```

524 that will summarize the variable using the function, `FUN`. The most useful of these is illus-
525 trated here:

```
526     hat=summary(zj$mu,quantile,c(.025,.5,.975))$stat
```

527 which produces the median and upper and lower .025% quantiles for μ , preserving its vector
528 structure. You can also give JAGS objects as arguments to other functions, a very handy
529 one being the empirical cumulative distribution function, `ecdf()`. For example the following
530 would estimate the probability that the parameter K is less than 900:

```
531     pK.lt.900 = ecdf(zj$K)(900)
```

⁷There is a *very important* caveat here. If the `rjags` library is not loaded when you enter an `jags` object name, R will not know to summarize it, and you will get the raw iterations. There can be a lot of these, leaving you bewildered as they fly by on the console. If you simply load the library, you will get more well behaved output.

532 **Exercise: making plots with JAGS objects** For the logistic example:

- 533 1. Plot the observations of growth rate as a function of observed population size.
- 534 2. Overlay the median of the model predictions as a solid line
- 535 3. Overlay the 95% credible intervals as dashed lines.
- 536 4. Prepare a separate plot of the posterior density of K .

537 6.2.3 The structure of JAGS objects (MCMC arrays)

538 Like coda objects, JAGS objects have a list structure, but instead of each element of the
539 list holding an array (i.e., matrix) for each chain, the JAGS objects holds an array for each
540 quantity estimated. This is easier illustrated than explained. The JAGS object below⁸ below
541 contains 5 iterations and two chains. Look at the object and think about how it is structured.
542 Note how the vector structure is preserved for the 16 estimates of μ :

```
543 > zj
544 $K
545 , , 1
546      [,1] [,2] [,3] [,4] [,5]
547 [1,] 1424.628 1411.863 1307.185 1338.801 1351.346
548 , , 2
549      [,1] [,2] [,3] [,4] [,5]
550 [1,] 1279.262 1326.353 1345.851 1243.561 1157.157
551 attr(,"class")
552 [1] "mccarray"
553 $mu
554 , , 1
555      [,1] [,2] [,3] [,4] [,5]
556 [1,] 0.17072948 0.19509308 0.19127273 0.19714752 0.19323022
557 [2,] 0.16631829 0.19000444 0.18586162 0.19170919 0.18795213
558 [3,] 0.16568811 0.18927749 0.18508861 0.19093228 0.18719812
559 [4,] 0.16442777 0.18782360 0.18354257 0.18937848 0.18569010
560 [5,] 0.15951244 0.18215340 0.17751305 0.18331862 0.17980879
561 [6,] 0.15888227 0.18142645 0.17674003 0.18254172 0.17905478
562 [7,] 0.14388420 0.16412508 0.15834225 0.16405139 0.16110928
563 [8,] 0.13770852 0.15700098 0.15076670 0.15643772 0.15371995
564 [9,] 0.12170217 0.13853649 0.13113209 0.13670435 0.13456802
565 [10,] 0.11628270 0.13228473 0.12448416 0.13002297 0.12808351
566 [11,] 0.09410068 0.10669615 0.09727399 0.10267593 0.10154226
```

⁸Actually, `rjags` makes it hard to “see” the object. If `rjags` is loaded, it presumes you want summaries. If you want to look at a complete listing of a JAGS object you save it, quit R, and restart it, load the JAGS object without loading `rjags`. The JAGS object then has the structure shown in the example.

```

567      [12,] 0.09258827 0.10495147 0.09541876 0.10081136 0.09973263
568      [13,] 0.07822037 0.08837704 0.07779399 0.08309794 0.08254113
569      [14,] 0.06322230 0.07107567 0.05939621 0.06460761 0.06459562
570      [15,] 0.05288749 0.05915372 0.04671875 0.05186637 0.05222981
571      [16,] 0.03839356 0.04243390 0.02893938 0.03399757 0.03488752
572      , , 2
573           [,1]      [,2]      [,3]      [,4]      [,5]
574      [1,] 0.19328215 0.18103879 0.18031947 0.18834429 0.187960699
575      [2,] 0.18768794 0.17599534 0.17537282 0.18272716 0.181909482
576      [3,] 0.18688876 0.17527484 0.17466616 0.18192471 0.181045022
577      [4,] 0.18529042 0.17383386 0.17325283 0.18031982 0.179316103
578      [5,] 0.17905686 0.16821401 0.16774086 0.17406073 0.172573319
579      [6,] 0.17825769 0.16749352 0.16703420 0.17325828 0.171708860
580      [7,] 0.15923735 0.15034577 0.15021561 0.15416003 0.151134723
581      [8,] 0.15140544 0.14328494 0.14329031 0.14629604 0.142663020
582      [9,] 0.13110643 0.12498440 0.12534106 0.12591388 0.120705748
583     [10,] 0.12423353 0.11878816 0.11926375 0.11901283 0.113271397
584     [11,] 0.09610261 0.09342679 0.09438920 0.09076667 0.082842422
585     [12,] 0.09418460 0.09169760 0.09269321 0.08884080 0.080767719
586     [13,] 0.07596343 0.07527035 0.07658128 0.07054500 0.061058042
587     [14,] 0.05694309 0.05812261 0.05976269 0.05144675 0.040483906
588     [15,] 0.04383664 0.04630652 0.04817341 0.03828661 0.026306770
589     [16,] 0.02545564 0.02973517 0.03192015 0.01983031 0.006424201
590     attr(,"class")
591     [1] "mccarray"
592     $r
593     , , 1
594           [,1]      [,2]      [,3]      [,4]      [,5]
595     [1,] 0.1795519 0.2052704 0.2020950 0.2080242 0.2037864
596     , , 2
597           [,1]      [,2]      [,3]      [,4]      [,5]
598     [1,] 0.2044706 0.1911257 0.1902128 0.1995786 0.2000631
599     attr(,"class")
600     [1] "mccarray"
601     $sigma
602     , , 1
603           [,1]      [,2]      [,3]      [,4]      [,5]
604     [1,] 0.03038826 0.02973461 0.03196986 0.02771297 0.02342979
605     , , 2
606           [,1]      [,2]      [,3]      [,4]      [,5]
607     [1,] 0.02939191 0.02266891 0.01886645 0.01684712 0.02437535
608     attr(,"class")
609     [1] "mccarray"

```

6.2.4 Manipulating JAGS objects

To understand how you can extract elements of the JAGS object you need to know its dimensions. For mcmc arrays that include scalars and vectors, each element in the list has three dimensions. For the scalars in the list, the first dimension⁹ is always = 1, the second

⁹This gives the the length. A scalar is a vector with length = 1.

614 dimension = number of iterations and the third dimension = the number of the chain.
615 For vectors, the first dimension of the JAGS object is the length of the vector, the second
616 dimension is the number of iterations, and the third dimension is the number of the chain.
617 An easy way to remember this is simply to enter `dim(jags.object)` at the console. Because
618 the dimensions are named, there is no ambiguity about the structure of the object. So for
619 example,

```
620     #dimensions of mu in the zj jags object:
621     dim(zj$mu)
622     #a vector containing all iterations of the second chain for K:
623     zj$K[1,2]
624     #a matrix for sigma with 2 rows, one for each chain, containing
625     #iterations 1 to 1000:
626     zj$sigma[1,1:1000,]
627     #a matrix containing 16 rows, one for each element of mu
628     #containing elements from the third chain:
629     zj$mu[,3]
```

630 So, if you wanted to find the mean of the third prediction of `mu` across all iterations and all
631 chains, you would use

```
632     mean(zj$mu[3,,])
```

633 **Exercise: Manipulating JAGS objects**

- 634 1. Calculate the median of the second chain for K .
- 635 2. Calculate the upper and lower 95% quantiles for the 16th estimate of μ without using
636 the `summary` function.
- 637 3. Calculate the probability that the 16th estimate of $\mu < 0$.

6.2.5 Converting JAGS objects to coda objects

It is possible to convert individual elements of the JAGS object to coda objects, which can be helpful for using convergence diagnostics (as described in the next section) if you haven't created a coda object directly using the `coda.samples` function. The syntax is

```
coda.object=as.mcmc.list(object.name$element.name).
```

So, for example, if you want to create a coda object for K , you would use

```
K.coda = as.mcmc.list(zj$K)
```

It is not possible to convert all of the elements of a JAGS object into coda objects in a single statement, i.e., the following will not work:

```
#wrong
```

```
jm = as.mcmc.list(zj)
```

7 Which object to use?

Coda and JAGS objects are both useful, and for most of my work I eventually create both types. Coda objects are somewhat better for producing tabular summaries of estimates and are required for checking convergence, but JAGS objects are somewhat better for plotting. Coda objects are also produced by WinBUGS and OpenBUGS, so if you ever need to use them, everything you learned about coda objects will apply. I generally start development of models using coda objects alone, and when I reach the final output stage, I produce both types of objects with multiple chains.

8 Checking convergence using the coda package

Remember from lecture that the MCMC chain will provide a reliable estimate of the posterior distribution only after it has converged, which means that it is no longer sensitive to initial

660 conditions and that the estimates of parameters of the posterior distribution will not change
661 appreciably with additional iterations. The coda package (Plummer et al., 2010) contains
662 a tremendous set of tools for evaluating and manipulating MCMC chains produced in coda
663 objects (i.e., MCMC lists). I urge you to look at the package documentation in R Help,
664 because we will use only a few of the tools it offers.

665 There are several ways to check convergence, but we will use four here: 1) visual inspection
666 of density and trace plots 2) Gelman and Rubin diagnostics, 3) Heidelberger and Welch
667 diagnostics, and 4) Raftery diagnostics. For all of these to work, the coda library must be
668 loaded.

669 **8.1 Trace and density plots**

670 There are three useful ways to plot the chains and the posterior densities. I am particularly
671 fond of the latter two because they show more detail.

```
672 plot(coda.object)
673 xyplot(coda.object)
674 densityplot(coda.object)
```

675 You will examine how to use these for diagnosing convergence in the subsequent exercise.

676 **8.2 Gelman and Rubin diagnostics**

677 The standard method for assuring convergence is the Gelman and Rubin diagnostic (Gelman
678 and Rubin, 1992), which “determines when the chains have ‘forgotten’ their initial values,
679 and the output from all chains is indistinguishable”(R Core Team, 2012). It requires at least
680 2 chains to work. For a complete treatment of how this works, enter `?gelman.diag` at the
681 console and read the section on Theory. We can be sure of convergence if all values for point
682 estimates and 97.5% quantiles approach 1. More iterations should be run if the 95% quantile
683 > 1.05 .

684 The syntax is

```
685 gelman.diag(coda.object)
```

686 8.3 Heidelberger and Welch diagnostics

687 The Heidelberger and Welch diagnostic (Heidelberger and Welch, 1983) works for a single
688 chain, which can be useful during early stages of model development before you have initial-
689 ized multiple chains. The diagnostic tests for stationary in the distribution and also tests if
690 the mean of the distribution is accurately estimated. For details do `?heidel.diag` and read
691 the part on Details. We can be confident of convergence if out all chains and all parameters
692 pass the test for stationarity and half width mean. We can be sure that the chain converged
693 from the first iteration (i.e, burn in was sufficiently long) if the `start.iteration = 1`. If it is
694 greater than 1, the burn in should be longer, or `1:start.iteration` should be discarded
695 from the chain.

696 The syntax is

```
697 heidel.diag(coda.object)
```

698 8.4 Raftery diagnostic

699 The Raftery diagnostic Raftery and Lewis (1995) is useful for planning how many iterations
700 to run for each chain. It is used early in the analysis with a relatively short chain, say 10000
701 iterations. It returns and estimate of the number of iterations required for convergence for
702 each of the parameters being estimated. Syntax is

```
703 raftery.diag(coda.object)
```

704 **Exercise:** Using the `zm.short` object your created above, increase `n.iter` in increments of
705 500 until you get convergence. For each increment:

- 706 1. Plot the chain and the posterior distributions of parameters using `xyplot` and `densityplot`.
- 707 2. Do Gelman-Rubin, Heidelberger and Welch, and Raftery diagnostics.
- 708 Discuss with you labmates how the plotting reveals convergence.

709 9 Monitoring deviance and calculating DIC

710 It is often a good idea to report the deviance of a model which is defined as $-2\log [P(y|\theta)]$.

711 To obtain the deviance of a JAGS model you need to do two things. First, you need to add

712 the statement

```
713     load.module("dic")
```

714 above your `jags.samples` statement and/or your `coda.samples` statement. In the list of

715 variables to be monitored, you add “deviance” i.e.,

```
716     zm=coda.samples(jm,variable.names=c("K", "r",  
717     "sigma", "deviance"), n.iter=25000, n.thin=1)
```

718 Later in the course we will learn about the Bayesian model selection statistic, the deviance

719 information criterion (DIC). DIC values are generated using syntax like this:

```
720     dic.object.name = dic.samples(jags.model, n.iter, type='pD')
```

721 So, to use your regression example, you would write something like:

```
722     dic.j = dic.samples(jm,n.iter=2500, type="pD")
```

723 If you enter `dic.j` at the console (or run it as a line of code in your script) R will respond

724 with something like:

```
725     Mean deviance:  -46.54  
726     penalty 1.852  
727     Penalized deviance:  -44.69
```

728 **10 Differences between JAGS and WinBUGS / Open-** 729 **BUGS**

730 The JAGS implementation of the BUGS language closely resembles the implementation
731 in WinBUGS and OpenBUGS, but there are some important structural differences that are
732 described in Chapter 8 of the JAGS manual (?). There are also some functions (for example,
733 matrix multiplication and the \wedge symbol for exponentiation) that are available in JAGS has
734 but that are not found in the other programs.

735 **11 Troubleshooting**

736 Some common error messages and their interpretation are found in Table 1.

Message	Interpretation
Unable to resolve parameter O[38,1:2] (one of its ancestors may be undefined)	May be due to NA in data or illegal value in variable on rhs of <- or ~.
Error parsing model file: syntax error on line 9 near "="	You used an = instead of <- for assignment
Error: Error in node Failure to calculate log density	You will get this with a Poisson density if you give it continuous numbers as data. It will also occur if variables take on undefined values like log of negative.
Warning message: In readLines(file) : incomplete final line found on 'SS2.R'	Will occur when you don't have a hard return after the last } for the model
syntax error, unexpected '}', expecting \$end	Occurs when there are mismatched parens
Error in jags.model("beta", data = data, n.chain = 1, n.adapt = 1000) : Error in node y[7] Invalid parent values	Occurs when there is an illegal mathematical operation or argument on the rhs. For example, negative values for argument to beta distribution or Poisson, divide by 0, log of negative, etc.
Error in setParameters(init.values[[i]], i) : Error in node sigma.s[1] Attempt to set value of non-variable node	You get this error when you have a variable in your init list that is not a stochastic node in the model, i.e., it is constant

738 12 Answers to exercises

739 **Exercise: using for loops** Write a code fragment to set vague normal priors [`dnorm(0,10e-6)`]
740 for 5 regression coefficients stored in the vector B.

```
741     for(i in 1:5){  
742         B[i] ~ dnorm(0,.000001)  
743     }
```

744 **Exercise: Understanding coda objects** Modify your code to produce a coda object
745 with 3 chains with 5 iterations each. Output

- 746 1. The estimate of σ for the third iteration from the second chain, `zm[[2]][2,3]`
- 747 2. All of the estimates of r from the first chain. `zm[[1]][,2]`

748 **Exercise: Manipulating coda summaries**

```
749     m=summary(zm)  
750     mu_sd=m$stat[,1:2] #make columns for mean and sd  
751     q=m$quantile[,c(3,1,5)] #make columns for median and CI  
752     table=cbind(mu_sd,q) #make table  
753     write.csv(file="/Users/Tom/Documents/Ecological Modeling Course/JAGS Primer/table_e
```

754 **Exercise:** Convert the `zm` object to a data frame. Using the elements of data frame (not
755 `zm`) as input to functions:

- 756 1. Find the maximum value of σ .
- 757 2. Estimate the mean of r for the first 1000 and last 1000 iterations in the chain.
- 758 3. Plot the density of K . (This is very handy for producing publication quality graphs of
759 posterior distributions.)

760 4. Estimate the probability that the parameter K exceeds 1600. (Hint: Look into using
761 the `ecdf()` function.) Estimate the probability that it falls between 800 and 1200.

```
762 #exercises on manipulating coda objects converted to data frames
763 df=as.data.frame(rbind(zm[[1]],zm[[2]],zm[[3]]))
764 max(df$sigma) #problem 1
765 mean(df$K[1:1000]) #problem 2, first part
766 nr=length(df$K)
767 mean(df$K[(nr-1000):nr]) #problem 2, second part
768 plot(density(df$K),main="",xlim=c(800,2000),xlab="K") #problem 3
769 1-ecdf(df$K)(1600) #problem 4, first part
770 ecdf(df$K)(1200)-ecdf(df$K)(800) #problem 4, second part.
```

771 **Exercise: vectors in coda objects:** Modify your code as described above and summarize
772 the coda object. What if you wanted to plot the model predictions with 95% credible intervals
773 against the data. How would you do that? There are several ways this can be done, but
774 the general idea is that you need to extract the rows of the coda object that contain the
775 quantiles for μ , which can be tedious and error prone. For example, if you use rows in the
776 summary table and add or subtract parameters to be estimated, then your row counts will
777 be off. There are ways to use rownames, but a far better way to plot vectors is described in
778 the section on JAGS objects.

779 **Exercise: using JAGS objects to plot vectors** For the logistic example:

- 780 1. Plot the data as points,
- 781 2. Overlay the median of the model predictions as a solid line
- 782 3. Overlay the 95% credible intervals as dashed lines.

```
783 zj=jags.samples(jm,variable.names=c("K", "r", "sigma", "mu"),
```

```

784     n.iter=50000, n.thin=1)
785     b=summary(zj$K,mean)$stat b=summary(zj$mu,quantile,
786     c(.025,.5,.975))$stat
787     plot(pop.data$"Population Size", pop.data$"Growth Rate", xlab="N",
788     ylab="Per capita growth rate")
789     lines(pop.data$"Population Size",b[2,])
790     lines(pop.data$"Population Size",b[1,],lty="dashed")
791     lines(pop.data$"Population Size",b[3,],lty="dashed")
792     plot(density(zj$K),xlab="K", main="", xlim=c(800,2500))

```




Figure 1: Median and 95% credible intervals for predicted growth rate and posterior density of K .

793 **Exercise: Manipulating JAGS objects**

- 794 1. Calculate the median of the second chain for K .
- 795 2. Calculate the upper and lower 95% quantiles for the 16th estimate of μ without using
796 the `summary` function.
- 797 3. Calculate the probability that the 16th estimate of $\mu < 0$.

```
798 > median(zj$K[1,,2])
799 [1] 1275.208
800 > quantile(zj$mu[16,,],c(.025,.975))

801 2.5% 97.5%

802 -0.01539839 0.05925297
803 > ecdf(zj$mu[16,,])(0)
804 [1] 0.1096533
805 >
```

Literature Cited

Gelman, A. and D. B. Rubin, 1992. Inference from iterative simulation using multiple sequences. *Statistical Science* **7**:457–511.

Heidelberger, P. and P. Welch, 1983. Simulation run length control in the presence of an initial transient. *Operations Research* **31**:1109–1044.

Knops, J. M. H. and D. Tilman, 2000. Dynamics of soil nitrogen and carbon accumulation for 61 years after agricultural abandonment. *Ecology* **81**:88–98.

McCarthy, M. A., 2007. Bayesian methods for ecology. Cambridge University Press, Cambridge, UK.

Plummer, M., 2011. JAGS version 3.0.0 user manual. http://sourceforge.net/projects/mcmc-jags/files/Manuals/3.x/jags_user_manual.pdf .

Plummer, M., N. Best, K. Cowles, and K. Vines, 2010. coda: Output analysis and diagnostics for MCMC. R package version 0.14-4. <http://CRAN.R-project.org/package=coda> .

R Core Team, 2012. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Raftery, A. and S. Lewis, 1995. The number of iterations, convergence diagnostics and generic Metropolis algorithms. Chapman and Hall, London, UK.

Index

823 **A**

- 824 `as.double`, 17
- 825 `as.mcmc.list`, 30

826 **C**

- 827 `c()`, 14
- 828 `coda.samples`, 19

829 **D**

- 830 `densityplot(coda.object)`, 31
- 831 deviance information criterion, 33
- 832 DIC, 33
- 833 `dic.samples`, 33
- 834 `dim(jags.object)`, 29

835 **F**

- 836 `for` loops, 11

837 **G**

- 838 Gelman and Rubin diagnostic, 31
- 839 `gelman.diag`, 32

840 **H**

- 841 Heidelberger and Welch diagnostic, 32
- 842 `heidel.diag`, 32

843 **J**

- 844 `jags.model`, 19

845 **L**

- 846 `length()`, 13
- 847 list of lists, 17

848 **M**

- 849 `max`, 15
- 850 MCMC arrays, 27
- 851 MCMC lists, 22
- 852 model statement, 10

853 **N**

- 854 nested `for` loops, 13

855 **P**

- 856 `plot(coda.object)`, 31
- 857 `precision`, 14
- 858 product likelihood, 13

859 **R**

- 860 Raftery diagnostic, 32
- 861 `raftery.diag`, 32

862 **U**

- 863 undefined values, 15

864 **X**

- 865 `xyplot(coda.object)`, 31