

1 An Ecological Modeler's Primer on JAGS

2 N. Thompson Hobbs

3 April 1, 2013

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	7
12	4.1 Windows	8
13	4.2 LINUX	8
14	5 Running JAGS	9
15	5.1 The JAGS model	9
16	5.2 Technical notes	10
17	5.2.1 The model statement	10
18	5.2.2 for loops	10
19	5.2.3 Specifying priors	13
20	5.2.4 The <- operator	14
21	5.2.5 Vector operations	14
22	5.2.6 Keeping variables out of trouble.	14
23	5.3 Running JAGS from R	15
24	6 Output from JAGS	20
25	6.1 coda objects	20
26	6.1.1 Summarizing coda objects	20
27	6.1.2 The structure of coda objects (MCMC lists)	21
28	6.1.3 Manipulating coda objects	23
29	6.2 JAGS objects	24
30	6.2.1 Why another object?	24

31	6.2.2	Summarizing the JAGS object	25
32	6.2.3	The structure of JAGS objects (MCMC arrays)	26
33	6.2.4	Manipulating JAGS objects	27
34	6.2.5	Converting JAGS objects to coda objects	29
35	7	Which object to use?	29
36	8	Checking convergence using the coda package	29
37	8.1	Trace and density plots	30
38	8.2	Gelman and Rubin diagnostics	30
39	8.3	Heidelberger and Welch diagnostics	31
40	8.4	Raftery diagnostic	31
41	9	Monitoring deviance and calculating DIC	32
42	10	Differences between JAGS and WinBUGS / OpenBUGS	33
43	11	Troubleshooting	33
44	12	Answers to exercises	35
45		Literature Cited	39

46 1 Aim

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

54 2 Introducing MCMC Samplers

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

62 The BUGS language is currently implemented in three flavors of software: OpenBUGS,
63 WinBUGS, and JAGS. OpenBUGS and WinBUGS run on Windows operating systems, while
64 JAGS was specifically constructed to run multiple platforms, including Mac OS and Unix.
65 Although all three programs use essentially the same syntax, OpenBUGS and WinBUGS
66 run in an elaborate graphical user interface, while JAGS only runs from the command line
67 of a Unix shell or from R. However, all three can be easily called from R, and this is the
68 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.

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

70 3 Introducing JAGS

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

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

89 which, of course, is a linear model with intercept r and slope $\frac{r}{K}$. Note that these quantities
 90 enjoy a sturdy biological interpretation; r is the intrinsic rate of increase, $\frac{r}{K}$ is the strength of
 91 the feedback from population size to population growth rate, and K is the carrying capacity,
 92 that is, the population size (o.k., o.k., the gm N per gm soil) at which $\frac{dN}{dt} = 0$. Presume
 93 we have some data consisting of observations of per capita rate of growth of N paired with
 94 observations of N . The vector \mathbf{y} contains values for the rate and the vector \mathbf{x} contains
 95 aligned data on N , i.e., $y_i = \frac{1}{N_i} \frac{dN_i}{dt}$, $x_i = N_i$. A simple Bayesian model specifies the joint
 96 distribution of the parameters and data as

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

97 where the priors are uninformative. Now, I have full, abiding confidence that with a couple
 98 of hours worth of work, perhaps less, you could knock out a Gibbs sampler to estimate r , K ,
 99 and τ . However, I am all for doing things nimbly in 15 minutes that might otherwise take a
 100 sweaty hour of hard labor, so, consider the code in algorithm 1, below.

101 This code illustrates the purpose of JAGS (and other BUGS software): to translate the
 102 numerator of Bayes theorem (a.k.a., the joint, e.g., equation 2) into a specification of an
 103 MCMC sampler. JAGS parses this code, sets up proposal distributions and steps in the
 104 Gibbs sampler and returns the MCMC chain for each parameter. These chains form the
 105 basis for estimating posterior distributions and associated statistics, i.e., means, medians,
 106 standard deviations, and quantiles. As we will soon learn, it easy to derive chains for other
 107 quantities of interest and their posterior distributions, for example, $K/2$ (What is $K/2$?),
 108 N as a function of time or dN/dt as a function of N . It is easy to construct comparisons
 109 between of the growth parameters of two populations or among ten of them. If this seems

110 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
```

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

119 4 Installing JAGS

120 Update your version of R to the most recent one. Go to the package installer under Packages
121 and Data on the toolbar and check the box in the lower right corner for install dependencies.
122 Install the `rjags` package from a CRAN mirror of your choice. Now go to <http://www-ice.>
123 [iarc.fr/~martyn/software/jags/](http://www-ice.iarc.fr/~martyn/software/jags/) and look in the section under downloads. Click on the
124 files page link and then click on Download JAGSdist-.dmg (4.7 MB) where _____ is the

125 number of the latest version to get the disk mounting image. Install as you would any other
126 Mac software.

127 **4.1 Windows**

128 Update your version of R to the most recent one. Go to the package installer under Packages
129 and Data on the toolbar and check the box in the lower right corner for install dependencies.
130 Install the `rjags` package from a CRAN mirror of your choice. Check the version number
131 of `rjags`. Go to <http://sourceforge.net/projects/mcmc-jags/files/JAGS/>. Click on
132 3x then `JAGS-3.3.0.exe`.

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

139 **4.2 LINUX**

140 There is a link to the path for binaries found at <http://mcmc-jags.sourceforge.net/>
141 . If you want to compile from source code, there are detailed instructions at [http://](http://yusung.blogspot.com/2009/01/install-jags-and-rjags-in-fedora.html)
142 yusung.blogspot.com/2009/01/install-jags-and-rjags-in-fedora.html There are tar
143 files found at <http://sourceforge.net/projects/mcmc-jags/files/JAGS/3.x/Source/>.
144 You want `JAGS-3.0.0.tar.gz`. My guess is that you will need to download the `rjags`
145 package in R before installing JAGS.

146 5 Running JAGS

147 5.1 The JAGS model

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

- 151 1. code for priors,
- 152 2. code for the deterministic model,
- 153 3. code for the likelihood(s).

154 The similarity between the code and equation 2 should be pretty obvious, but there are a few
155 things to point out. Priors and likelihoods are specified using the \sim notation that we have
156 seen in class. For example, remember that

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

157 is the same as

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

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

163 **Exercise: always plot your priors** Plot priors for each parameter, scaling the x axis
164 appropriately for each value— r should be about .2, K about 1200, and τ should be about
165 2500. Discuss with you lab mates if $\text{gamma}(\theta \mid .001, .001)$ is vague for all parameters, i.e.,

166 $\theta = r, K, \tau$. Be sure to include lots of x points in your plots to get a good interpolation, at
167 least 1000.

168 5.2 Technical notes

169 5.2.1 The model statement

170 Your entire model must be enclosed in

```
171     model{  
172     .  
173     .  
174     .  
175     .  
176     } #end of model
```

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

181 5.2.2 for loops

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

```

189     mu[1] <- r - r/K * x[1]
190     y[1] ~ dnorm(mu[1],tau)
191     mu[2] <- r - r/K * x[2]
192     y[2] ~ dnorm(mu[3],tau)
193     mu[3] <- r - r/K * x[3]
194     y[3] ~ dnorm(mu[3],tau)
195     .
196     .
197     .
198     mu[n] <- r - r/K * x[n]
199     y[n] ~ dnorm(mu[n],tau)

```

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

```

202     for(i in 1:n){
203         mu[i] <- r - r/K * x[i]
204         y[i] ~ dnorm(mu[i],tau)
205     }

```

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

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

```

211     #WRONG!!!
212     for(i in 1:n){
213         mu <- r - r/K * x[i]

```

```

214     y[i] ~ dnorm(mu,tau)
215   }

```

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

```

219   #Wrong
220   mu <- r - r/K * x[1]
221   y[1] ~ dnorm(mu,tau)
222   mu <- r - r/K * x[2]
223   y[2] ~ dnorm(mu,tau)
224   mu <- r - r/K * x[3]
225   y[3] ~ dnorm(mu,tau)
226   .
227   .
228   .
229   mu <- r - r/K * x[n]
230   y[n] ~ dnorm(mu,tau),

```

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

```

238   for(i in 1:n){

```

```

239         for(j in 1:m){
240             expression[i,j]
241         } #end of j loop
242     } #end of i loop

```

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

```

245     for(1 in 1:length(x[])){
246         mu[i] <- r - r/K * x[i]
247         y[i] ~ dnorm(mu[i],tau)
248     }

```

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

251 5.2.3 Specifying priors

252 We specify priors in JAGS as `parameter ~ distribution(shape1, shape2)`. See the sup-
253plementary material for available distributions. Note that in the code (algorithm 1), the
254second argument to the normal density function is `tau`, which is the precision, defined as
255the reciprocal of the variance. This means that we must calculate `sigma` from `tau` if we
256want a posterior distribution on `sigma`. Be very careful about this—it is easy to forget that
257you must use the precision rather than the standard deviation as an argument to `dnorm` or
258`dlnorm`. For the lognormal, it is the precision on the log scale. If you would like, you can
259express priors on σ rather than τ using code like this:

```

260
261     sigma~dunif(0,100) #presuming this more than brackets the posterior of sigma
262     tau <- 1/sigma^2

```

263 There are times when this seems to work better than the gamma prior for

264 **5.2.4 The <- operator**

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

267 **5.2.5 Vector operations**

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

273 **5.2.6 Keeping variables out of trouble.**

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

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

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

285 5.3 Running JAGS from R

286 We implement our model using R (algorithm 2.) We will go through the R code step by
287 step. We start by bringing the growth rate data into R as a data frame. Next, we specify
288 the initial conditions for the MCMC chain in the statement `inits = . . .`. This is exactly the
289 same thing as you did when you wrote your MCMC code and assigned a guess to the first
290 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.real(pop.data$"Population Size"),
  y=as.real(pop.data$"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, n.thin=1)
```

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

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

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

294 you will get an error message when you execute the `jags.model` statement and your code
295 will not run. Second, this statement allows you to set up multiple chains³, which are needed
296 for some tests of convergence and to calculate DIC (more about these tasks later). For
297 example, if you want three chains, you would use something like:

```
298     inits=list(  
299     list(K=1500, r=.5, tau=1500), #chain 1  
300     list(K=1000, r=.1, tau=1000), #chain 2  
301     list(K=900, r=.3, tau=900) #chain 3  
302     ) #end of inits list
```

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

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

314 The next couple of statements,

```
315     n.xy = nrow(pop.data)  
316     data=list(n=n.xy,
```

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


```
317     x=as.real(pop.data$"Population Size"),
318     y=as.real(pop.data$"Growth Rate"))
```

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

```
321     x=as.real(pop.data$"Population Size")
```

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

```
324     y=as.real(pop.data$"Growth Rate")
```

325 creates a y vector required by the JAGS program from the column in your data frame called
326 `Growth Rate` (pretty cool, I think). Notice that if I had named the variable `Growth.Rate`
327 instead of `Growth Rate`, the quotes would not be needed. **It is important for you to un-**

328 derstand that the left hand side of the = corresponds to variable name for the data in the

329 JAGS program and the right hand side of the = is what they are called in R. Also, note

330 that because `pop.data` is a data frame I used `as.real()` to be sure that JAGS received
331 real numbers instead of characters or factors, as can happen with data frames. This can
332 be particularly important if you have missing data in the data. The `n` is required in the
333 JAGS program to index the `for` structure (algorithm 2) and it must be read as data in
334 this statement⁴. By the way, you don't need to call this list "data"—it could be anything
335 ("apples", "bookshelves", "xy" etc.)

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

```
338     library(rjags)
339     ##call to JAGS
```

⁴You could hard code the `for` index in the JAGS code, but this is bad practice.

```

340     n.adapt=5000
341     n.update = 10000
342     n.iter = 25000
343     jm=jags.model("Logistic example JAGS.R",data=data,init=init,n.chains=length(init),
344                 n.adapt = n.adapt)
345     #Burnin the chain.
346     update(jm, n.iter=n.update)
347     #generate coda object
348     zm=coda.samples(jm,variable.names=c("K", "r", "sigma"), n.iter=n.iter, n.thin=1)

```

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

359 The `jm=jags.model....` statement sets up the MCMC chain. Its first argument is the
360 name of the file containing the BUGS code. Note that in this case, the file resided in the
361 current working directory, which I specified at the top of the code (algorithm 2). Otherwise,
362 you would need to specify the full path name. (It is also possible to embed the BUGS code
363 within your R script, see Algorithm 3,). The next two expressions specify where the data
364 come from, where to get the initial values, and how many chains to create (i.e., the length
365 of the list `init`). Finally, it specifies the “burn-in” how many samples to throw away before

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()
```

366 beginning to save values in the chain. Thus, in this case, we will throw away the first 10,000
367 values.

368 The second statement (`zm=coda.samples...`) creates the chains and stores them as
369 an MCMC list (more about that soon). The first argument (`jm`) is the name of the jags
370 model you created in the `jags.model` function. The second argument (`variable.names`)
371 tells JAGS which variables to “monitor.” These are the variables for which you want poste-
372 rior distributions. Finally, `n.iter=n.iter` says we want 25000 elements in each chain and
373 `n.thin` specifies how many of these to keep. For example, if `n.thin = 10`, we would store
374 every 10th element. Sometimes setting `n.thin > 1` is a good idea to reduce the size of the
375 data files that you will analyze.

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

379 6 Output from JAGS

380 6.1 coda objects

381 6.1.1 Summarizing coda objects

382 The `zm` object produced by the statement

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

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

```
389   Iterations = 15001:25000
390   Thinning interval = 1
391   Number of chains = 3
392   Sample size per chain = 10000
393   1. Empirical mean and standard deviation for each variable,
394      plus standard error of the mean:
395
396           Mean          SD Naive SE Time-series SE
397   K      1.313e+03 1.180e+02 6.811e-01      1.244e+00
398   r      1.998e-01 1.101e-02 6.359e-05      1.113e-04
399   sigma 2.538e-02 5.204e-03 3.004e-05      3.604e-05
400
401   2. Quantiles for each variable:
402
403           2.5%       25%       50%       75%       97.5%
404   K      1.125e+03 1.235e+03 1.300e+03 1.374e+03 1.583e+03
405   r      1.776e-01 1.928e-01 1.999e-01 2.070e-01 2.213e-01
406   sigma 1.759e-02 2.169e-02 2.460e-02 2.814e-02 3.773e-02
```

404 Each of the two tables above has the properties of a matrix⁵. You can output the cells of
405 these tables using syntax as follows. To get the mean and standard deviation of r,

```
406 > summary(zm)$stat[2,1:2]
407           Mean           SD
408 0.19980128 0.01101439
```

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

```
410 > summary(zm)$quantile[1,c(1,5)]
411      2.5%      97.5%
412 1124.539 1582.647
```

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

417 6.1.2 The structure of coda objects (MCMC lists)

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

```
419 [[1]]
420 Markov Chain Monte Carlo (MCMC) output:
421 Start = 60001
422 End = 60010
423 Thinning interval = 1
424           K           r           sigma
```

⁵Consider `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.

```

425 [1,] 1096.756 0.1914722 0.02889710
426 [2,] 1196.326 0.2088859 0.03155777
427 [3,] 1401.511 0.1804327 0.02553913
428 [4,] 1471.539 0.1754886 0.03589013
429 [5,] 1245.909 0.1567580 0.04248644
430 [6,] 1134.738 0.2114307 0.04151478
431 [7,] 1105.661 0.2303630 0.03141035
432 [8,] 1108.569 0.2169765 0.03708956
433 [9,] 1134.755 0.1964426 0.02660658
434 [10,] 1161.750 0.2152418 0.03700475
435 .
436 .
437 .

```

438 as many rows as you have thinned iterations

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

```

444 [[1]]
445 Markov Chain Monte Carlo (MCMC) output:
446 Start = 10001
447 End = 10005
448 Thinning interval = 1
449           K           r           sigma
450 [1,] 1070.013 0.2126878 0.02652204
451 [2,] 1085.438 0.2279789 0.02488036

```

```

452     [3,] 1170.086 0.2259743 0.02331958
453     [4,] 1094.564 0.2228788 0.02137309
454     [5,] 1053.495 0.2368199 0.03209893
455     [[2]]
456     Markov Chain Monte Carlo (MCMC) output:
457     Start = 10001
458     End = 10005
459     Thinning interval = 1
460           K           r           sigma
461     [1,] 1137.501 0.2657460 0.04093364
462     [2,] 1257.340 0.1332901 0.04397191
463     [3,] 1073.023 0.2043738 0.03355776
464     [4,] 1159.732 0.2339060 0.02857740
465     [5,] 1368.568 0.2021042 0.05954259
466     attr(,"class")
467     [1] "mcmc.list"

```

468 **Exercise: Understanding coda objects:** Modify your code to produce a coda object with
469 3 chains called `zm.short`, setting `n.adapt = 500`, `n.update=500`, and `n.iter = 20`.

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

473 6.1.3 Manipulating coda objects

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

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

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

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

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

488 6.2 JAGS objects

489 6.2.1 Why another object?

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

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


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

502 6.2.2 Summarizing the JAGS object

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

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

506 If you run this and enter `zj` at the console, R will return the means of all the monitored
507 variables⁶. Try it. If you want other statistics, you would use syntax like:

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

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

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

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

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

⁶There 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.

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

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

522 6.2.3 The structure of JAGS objects (MCMC arrays)

523 Like coda objects, JAGS objects have a list structure, but instead of each element of the
524 list holding an array (i.e., matrix) for each chain, the JAGS objects holds an array for each
525 quantity estimated. This is easier illustrated than explained. The JAGS object below⁷ below
526 contains 5 iterations and two chains. Look at the object and think about how it is structured.

527 Note how the vector structure is preserved for the 16 estimates of μ :

```
528 > zj
529 $K
530 , , 1
531      [,1] [,2] [,3] [,4] [,5]
532 [1,] 1424.628 1411.863 1307.185 1338.801 1351.346
533 , , 2
534      [,1] [,2] [,3] [,4] [,5]
535 [1,] 1279.262 1326.353 1345.851 1243.561 1157.157
536 attr(,"class")
537 [1] "mccarray"
538 $mu
539 , , 1
540      [,1] [,2] [,3] [,4] [,5]
541 [1,] 0.17072948 0.19509308 0.19127273 0.19714752 0.19323022
542 [2,] 0.16631829 0.19000444 0.18586162 0.19170919 0.18795213
543 [3,] 0.16568811 0.18927749 0.18508861 0.19093228 0.18719812
544 [4,] 0.16442777 0.18782360 0.18354257 0.18937848 0.18569010
545 [5,] 0.15951244 0.18215340 0.17751305 0.18331862 0.17980879
546 [6,] 0.15888227 0.18142645 0.17674003 0.18254172 0.17905478
547 [7,] 0.14388420 0.16412508 0.15834225 0.16405139 0.16110928
548 [8,] 0.13770852 0.15700098 0.15076670 0.15643772 0.15371995
549 [9,] 0.12170217 0.13853649 0.13113209 0.13670435 0.13456802
550 [10,] 0.11628270 0.13228473 0.12448416 0.13002297 0.12808351
551 [11,] 0.09410068 0.10669615 0.09727399 0.10267593 0.10154226
```

⁷Actually, 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.

```

552      [12,] 0.09258827 0.10495147 0.09541876 0.10081136 0.09973263
553      [13,] 0.07822037 0.08837704 0.07779399 0.08309794 0.08254113
554      [14,] 0.06322230 0.07107567 0.05939621 0.06460761 0.06459562
555      [15,] 0.05288749 0.05915372 0.04671875 0.05186637 0.05222981
556      [16,] 0.03839356 0.04243390 0.02893938 0.03399757 0.03488752
557      , , 2
558          [,1]      [,2]      [,3]      [,4]      [,5]
559      [1,] 0.19328215 0.18103879 0.18031947 0.18834429 0.187960699
560      [2,] 0.18768794 0.17599534 0.17537282 0.18272716 0.181909482
561      [3,] 0.18688876 0.17527484 0.17466616 0.18192471 0.181045022
562      [4,] 0.18529042 0.17383386 0.17325283 0.18031982 0.179316103
563      [5,] 0.17905686 0.16821401 0.16774086 0.17406073 0.172573319
564      [6,] 0.17825769 0.16749352 0.16703420 0.17325828 0.171708860
565      [7,] 0.15923735 0.15034577 0.15021561 0.15416003 0.151134723
566      [8,] 0.15140544 0.14328494 0.14329031 0.14629604 0.142663020
567      [9,] 0.13110643 0.12498440 0.12534106 0.12591388 0.120705748
568     [10,] 0.12423353 0.11878816 0.11926375 0.11901283 0.113271397
569     [11,] 0.09610261 0.09342679 0.09438920 0.09076667 0.082842422
570     [12,] 0.09418460 0.09169760 0.09269321 0.08884080 0.080767719
571     [13,] 0.07596343 0.07527035 0.07658128 0.07054500 0.061058042
572     [14,] 0.05694309 0.05812261 0.05976269 0.05144675 0.040483906
573     [15,] 0.04383664 0.04630652 0.04817341 0.03828661 0.026306770
574     [16,] 0.02545564 0.02973517 0.03192015 0.01983031 0.006424201
575     attr(,"class")
576     [1] "marray"
577     $r
578     , , 1
579         [,1]      [,2]      [,3]      [,4]      [,5]
580     [1,] 0.1795519 0.2052704 0.2020950 0.2080242 0.2037864
581     , , 2
582         [,1]      [,2]      [,3]      [,4]      [,5]
583     [1,] 0.2044706 0.1911257 0.1902128 0.1995786 0.2000631
584     attr(,"class")
585     [1] "marray"
586     $sigma
587     , , 1
588         [,1]      [,2]      [,3]      [,4]      [,5]
589     [1,] 0.03038826 0.02973461 0.03196986 0.02771297 0.02342979
590     , , 2
591         [,1]      [,2]      [,3]      [,4]      [,5]
592     [1,] 0.02939191 0.02266891 0.01886645 0.01684712 0.02437535
593     attr(,"class")
594     [1] "marray"

```

595 6.2.4 Manipulating JAGS objects

596 To understand how you can extract elements of the JAGS object you need to know its
597 dimensions. For mcmc arrays that include scalars and vectors, each element in the list has
598 three dimensions. For the scalars in the list, the first dimension⁸ is always = 1, the second

⁸This gives the the length. A scalar is a vector with length = 1.

599 dimension = number of iterations and the third dimension = the number of the chain.
600 For vectors, the first dimension of the JAGS object is the length of the vector, the second
601 dimension is the number of iterations, and the third dimension is the number of the chain.
602 An easy way to remember this is simply to enter `dim(jags.object)` at the console. Because
603 the dimensions are named, there is no ambiguity about the structure of the object. So for
604 example,

```
605     #dimensions of mu in the zj jags object:
606     dim(zj$mu)
607     #a vector containing all iterations of the second chain for K:
608     zj$K[1,2]
609     #a matrix for sigma with 2 rows, one for each chain, containing
610     #iterations 1 to 1000:
611     zj$sigma[1,1:1000,]
612     #a matrix containing 16 rows, one for each element of mu
613     #containing elements from the third chain:
614     zj$mu[,3]
```

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

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

618 **Exercise: Manipulating JAGS objects**

- 619 1. Calculate the median of the second chain for K .
- 620 2. Calculate the upper and lower 95% quantiles for the 16th estimate of μ without using
621 the `summary` function.
- 622 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

645 conditions and that the estimates of parameters of the posterior distribution will not change
646 appreciably with additional iterations. The coda package (?) contains a tremendous set of
647 tools for evaluating and manipulating MCMC chains produced in coda objects (i.e., MCMC
648 lists). I urge you to look at the package documentation in R Help, because we will use only
649 a few of the tools it offers.

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

654 8.1 Trace and density plots

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

```
657 plot(coda.object)
658 xyplot(coda.object)
659 densityplot(coda.object)
```

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

661 8.2 Gelman and Rubin diagnostics

662 The standard method for assuring convergence is the Gelman and Rubin diagnostic (Gelman
663 and Rubin, 1992), which “determines when the chains have ‘forgotten’ their initial values,
664 and the output from all chains is indistinguishable”(?). It requires at least 2 chains to work.
665 For a complete treatment of how this works, enter ?gelman.diag at the console and read
666 the section on Theory. We can be sure of convergence if all values for point estimates and
667 97.5% quantiles approach 1. More iterations should be run if the 95% quantile > 1.05 .

668 The syntax is

669 `gelman.diag(coda.object)`

670 **8.3 Heidelberg and Welch diagnostics**

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

680 The syntax is

681 `heidel.diag(coda.object)`

682 **8.4 Raftery diagnostic**

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

687 `raftery.diag(coda.object)`

688 **Exercise:** Using the `zm.short` object your created above, increase `n.iter` in increments of
689 500 until you get convergence. For each increment:

- 690 1. Plot the chain and the posterior distributions of parameters using `xyplot` and `densityplot`.
- 691 2. Do Gelman-Rubin, Heidelberg and Welch, and Raftery diagnostics.

692 Discuss with you labmates how the plotting reveals convergence.

693 9 Monitoring deviance and calculating DIC

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

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

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

698 above your `jags.samples` statement and/or your `coda.samples` statement. In the list of
699 variables to be monitored, you add “deviance” i.e.,

```
700 zm=coda.samples(jm,variable.names=c("K", "r",
```

```
701 "sigma", "deviance"), n.iter=25000, n.thin=1)
```

702 Later in the course we will learn about the Bayesian model selection statistic, the deviance
703 information criterion (DIC). DIC values are generated using syntax like this:

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

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

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

707 If you enter `dic.j` at the console (or run it as a line of code in your script) R will respond
708 with something like:

```
709 Mean deviance: -46.54
```

```
710 penalty 1.852
```

```
711 Penalized deviance: -44.69
```


712 **10 Differences between JAGS and WinBUGS / Open-** 713 **BUGS**

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

719 **11 Troubleshooting**

720 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

722 12 Answers to exercises

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

```
725     for(i in 1:5){  
726         B[i] ~ dnorm(0,.000001)  
727     }
```

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

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

732 **Exercise: Manipulating coda summaries**

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

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

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

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

```
746 #exercises on manipulating coda objects converted to data frames
747 df=as.data.frame(rbind(zm[[1]],zm[[2]],zm[[3]]))
748 max(df$sigma) #problem 1
749 mean(df$K[1:1000]) #problem 2, first part
750 nr=length(df$K)
751 mean(df$K[(nr-1000):nr]) #problem 2, second part
752 plot(density(df$K),main="",xlim=c(800,2000),xlab="K") #problem 3
753 1-ecdf(df$K)(1600) #problem 4, first part
754 ecdf(df$K)(1200)-ecdf(df$K)(800) #problem 4, second part.
```

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

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

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

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

```

768     n.iter=50000, n.thin=1)
769     b=summary(zj$K,mean)$stat b=summary(zj$mu,quantile,
770     c(.025,.5,.975))$stat
771     plot(pop.data$"Population Size", pop.data$"Growth Rate", xlab="N",
772     ylab="Per capita growth rate")
773     lines(pop.data$"Population Size",b[2,])
774     lines(pop.data$"Population Size",b[1,],lty="dashed")
775     lines(pop.data$"Population Size",b[3,],lty="dashed")
776     plot(density(zj$K),xlab="K", main="", xlim=c(800,2500))

```




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

777 **Exercise: Manipulating JAGS objects**

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

```
782 > median(zj$K[1,,2])
783 [1] 1275.208
784 > quantile(zj$mu[16,,],c(.025,.975))

785 2.5% 97.5%

786 -0.01539839 0.05925297
787 > ecdf(zj$mu[16,,])(0)
788 [1] 0.1096533
789 >
```

Literature Cited

- 790
791 Gelman, A. and D. B. Rubin, 1992. Inference from iterative simulation using multiple
792 sequences. *Statistical Science* **7**:457–511.
- 793 Heidelberger, P. and P. Welch, 1983. Simulation run length control in the presence of an
794 initial transient. *Operations Research* **31**:1109–1044.
- 795 Knops, J. M. H. and D. Tilman, 2000. Dynamics of soil nitrogen and carbon accumulation
796 for 61 years after agricultural abandonment. *Ecology* **81**:88–98.
- 797 McCarthy, M. A., 2007. Bayesian methods for ecology. Cambridge University Press, Cam-
798 bridge, UK.
- 799 Plummer, M., 2011. JAGS version 3.0.0 user manual. [http://sourceforge.net/
800 projects/mcmc-jags/files/Manuals/3.x/jags_user_manual.pdf](http://sourceforge.net/projects/mcmc-jags/files/Manuals/3.x/jags_user_manual.pdf) .
- 801 Raftery, A. and S. Lewis, 1995. The number of iterations, convergence diagnostics and
802 generic Metropolis algorithms. Chapman and Hall, London, UK.

Index

803 **A**

804 `as.mcmc.list`, 29

805 `as.real`, 17

806 **C**

807 `c()`, 14

808 `coda.samples`, 19

809 **D**

810 `densityplot(coda.object)`, 30

811 deviance information criterion, 32

812 DIC, 32

813 `dic.samples`, 32

814 `dim(jags.object)`, 28

815 **F**

816 `for loops`, 10

817 **G**

818 Gelman and Rubin diagnostic, 30

819 `gelman.diag`, 31

820 **H**

821 Heidelberger and Welch diagnostic, 31

822 `heidel.diag`, 31

823 **J**

824 `jags.model`, 18

825 **L**

826 `length()`, 13

827 list of lists, 16

828 **M**

829 `max`, 14

830 MCMC arrays, 26

831 MCMC lists, 21

832 model statement, 10

833 **N**

834 nested for loops, 12

835 **P**

836 `plot(coda.object)`, 30

837 `precision`, 13

838 product likelihood, 12

839 **R**

840 Raftery diagnostic, 31

841 `raftery.diag`, 31

842 **U**

843 undefined values, 14

844 **X**

845 `xyplot(coda.object)`, 30