

Sources

Packages

"BayesTraits" {Command-line driven executable} (Pagel & Meade 2006)
"Mesquite" (Maddison & Maddison 2009)

Data

<code>primates.txt:</code>	tab separated text file, data on the presence or absence of oestrus advertisement, and the presence or absence of multi-male groups for 60 Old World Monkey and Great Ape species
<code>primates.trees:</code>	a sample of evolutionary trees in <i>nexus</i> format showing the phylogenetic relationships between these 60 species created using genetic data.

Tutorials and exercises

Introduction

In this practical we're going to be using the program BayesTraits for performing phylogenetic comparative analyses of discrete characters in a Bayesian MCMC framework. This program is free to download from <http://www.evolution.reading.ac.uk/SoftwareMain.html>. The aim of the practical is to familiarize yourself with the program and introduce some of the features and types of analyses that it can perform. The practical is designed to get you up-and-running in using BayesTraits and is by no means meant as a full introduction to Bayesian statistics. Please refer to the chapters in the book for information about the theoretical issues involved in Phylogenetic Comparative Methods. The manual that accompanies the download of BayesTraits also provides more information about the capabilities of BayesTraits and should be consulted if you want to use the program for your own analyses. This practical will introduce the concept of model-based approaches via the Maximum Likelihood method of estimation (which should be used as a first-step in these kinds of analyses in any case), before going on to demonstrate how full Bayesian analyses can be performed. At the end of the practical we will introduce a check-list of the steps you need to take, and the things you need to look out for and check in order to achieve a reliable Bayesian analysis.

Data

The data we will be using were introduced in the main book chapter and are taken from Pagel and Meade's 2006 paper, and are also included with the download of BayesTraits. The manual that comes with BayesTraits uses these data in some of its examples, and represents a valuable source of information on how to use the program. The data are discretely distributed, binary (0 or 1) data that describe the presence or absence of 1) Oestrus Advertisement, and 2) Multimale Groups in a sample of 60 Old World Monkey and Ape species. The evolutionary relationships between these species are represented by a sample of 500 phylogenetic trees created using genetic data (the variation between trees in this sample represents our uncertainty about how exactly these species are related, and how much divergence between

species has occurred). These data can be visualized by mapping them onto one of the trees from the tree sample (Figure 1). Visualizing your data is an important step in any statistical analysis. This can be done in various R packages and other programs such as Mesquite (Maddison & Maddison 2009). From this figure we can see that closely related species tend to resemble each other in both of these traits, and that most species have both oestrus advertisement and multimale groups, or they lack both of these features. Looking at the data in this way reveals that the association between these features tends to be clustered in clades, e.g. Macaques and Baboons tend to have both oestrus advertisement and multimale groups, while Great Apes (with the exception of Chimps) tend to have neither. As we discussed in the main chapter this pattern of association between these traits could be the result of a functional link, or could be the result of both traits emerging independently during evolutionary history such that the relevant species possess both traits due to common inheritance.




Figure 1. Data on Oestrus advertisement (left) and Multimale systems (right) mapped onto a phylogeny of 60 Old World monkey and ape species (this is just one tree from a sample of 500). White represents absence, and black represents presence of these traits (grey indicates that data are missing for that species). The probability of each character state at internal nodes has been reconstructed using Maximum Likelihood in the program Mesquite.

Basic information about BayesTraits

BayesTraits is command line driven (i.e. there is no graphical user interface) so we have to enter all commands using the keyboard. To run analyses in BayesTraits we need two types of file: 1) a tree file containing a phylogenetic tree, or sample of trees, 2) a data file containing the trait data to be analyzed. BayesTraits requires input phylogenetic trees to be in NEXUS format (i.e. it should have a #NEXUS tag at the top of the file). For this practical the trees are in the file `Primates.trees`. The trait data are in the file `Primates.txt`. The species

names need to be listed in the first column (these need to be **exactly the same** as in your tree file) with trait data in subsequent columns separated by a tab like this:

Nasalis_larvatus	0	0
Cercocebus_torquatus	1	1
Cercopithecus_aethiops	0	1
Cercopithecus_mona	0	0
Cercopithecus_nictitans	0	0

There are no column headings and missing data needs to be represented by a hyphen (-).

Example 1 Running BayesTraits from the command prompt

BayesTraits is run from the command prompt (if you double-click on it trying to open it like other programs then nothing much will happen).

In Windows the command prompt can be found in the Accessories folder, or you can call it by searching for “command” in Windows Explorer or the Start Menu (on a Mac the program that brings up the command line is called Terminal). When you open up the command line you will be in your home directory. However, to make entering commands more straightforward it is preferable to be in the folder that contains the BayesTraits program. In Windows Explorer (or Finder etc.) navigate to the folder containing BayesTraits and copy the file path (i.e. something like C:/Username/Documents/BayesTraits , NB: the precise file path will depend on the folder structure of your computer). In the command prompt type `cd` (change directory) and paste the file path, e.g.:

```
> cd C:/Username/Documents/BayesTraits
```

You should now be in the folder that contains BayesTraits. You need to tell the command line to run BayesTraits and need to tell it which treefile and datafile you want to analyse. So in the command line type:

```
BayesTraits.exe Primates.trees Primates.txt
```

and press return. A menu should appear with several different options, like this:

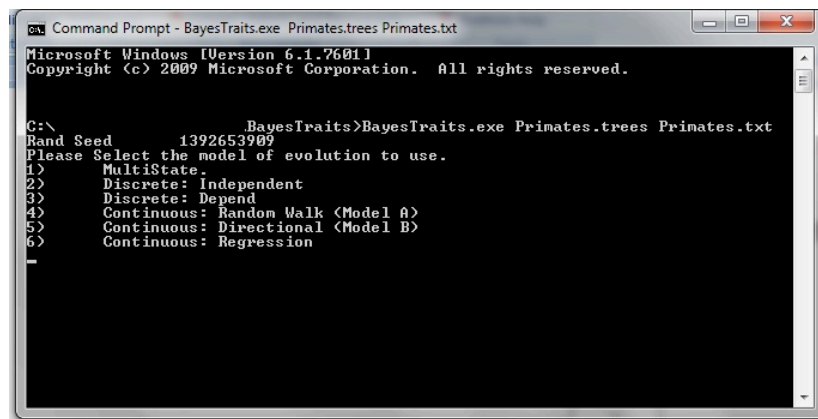


Figure 2. First option menu in BayesTraits

The top three options (1-3) are for categorical variables, while the bottom three are for continuous traits (4-6) (NB: the precise options you see in this window may vary as BayesTraits develops and new options and capabilities are added). The MULTISTATE option is for analyzing a single trait that can take more than one state. The two DISCRETE options are for comparing the dependent and independent models of trait co-evolution using two binary traits.

Binary Trait Co-evolution in BayesTraits

In the next section we will see how we can run co-evolutionary analyses in BayesTraits. As you remember from the book chapter, under the independent model two binary traits evolve without influencing each other, whereas under the dependent model the rate of change of one trait can be influenced by the state of the other trait. We will first see an example of the commands used for the Independent model and the output it generates. We then move onto the Dependent model and see how models of evolution can be compared and assessed.

Example 2 Independent model

In the BayesTraits menu select option 2 *DISCRETE: Independent model* by pressing the number 2 key. You will then be asked whether to run a Maximum Likelihood or a Bayesian MCMC analysis. We'll be starting with Maximum Likelihood analyses first, so press number 1 key. Information about the default settings and parameter states will appear. This list tells you information about the analysis setup (it can be called up by typing `info` at any time before running an analysis) (see Figure 3). You will see that there are four rate parameters to be estimated α_1 , β_1 , α_2 , and β_2 . Alpha designates a change from 0 to 1, and beta from 1 to 0 (see Figure 3). The numbers after alpha and beta refer to the different traits in the same order as they are entered in columns in the data file, i.e. 1 is Oestrus, 2 is Mating System.

```

C:\ Command Prompt - BayesTraits.exe Primates.trees Primates.txt
Options:
Model: Discete Independant
Tree File Name: Primates.trees
Data File Name: Primates.txt
Log File Name: Primates.txt.log.txt
Summary: False
Analysis Type: Maximum Likelihood
ML attempt per tree: 10
No of Rates: 4
Base frequency <PI's> None
Character Symbols 00,01,10,11
Using a covarion model: False
Restrictions:
alpha1 None
beta1 None
alpha2 None
beta2 None
Tree Information
Trees: 500
Taxa: 60
Sites: 1
States: 4
>

```

Figure 3. Settings and command line for the independent model

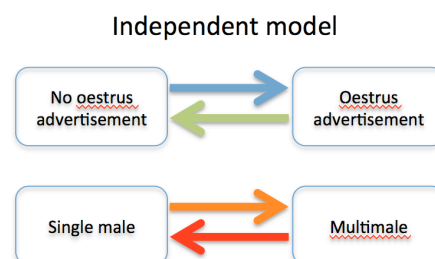


Figure 4. Representation of the Independent model of evolution. In the example used in this practical the blue arrow is parameter α_1 , green is β_1 , orange is α_2 , red is β_2 .

First, let's just run an analysis using these defaults. Type

```
run
```

The program will perform a Maximum Likelihood analysis on each tree in the treefile. The program will produce an output in the command line window, and will also write the results to a log-file. By default this file is named after the datafile with ".log.txt" added to the end (e.g. primates.txt.log.txt). You will see that the output gives the Maximum likelihood estimate (lh), the ML estimate of the four rate parameters, and the probability of the states at the root (two for each trait). It is often best to view these results by opening up the log-file in a spreadsheet program such as Excel. We can see that the likelihood is around -40 log units. The likelihood, the parameter estimates and the probability at the root all show some variation from tree to tree, indicating that this may be an important factor to be aware when assessing our results.

So that we don't overwrite our previous results, let's tell BayesTraits to store the next output in a new log-file (lf) (note that the first character is an "L" not an "I"), so type:

```
lf primates_ML_dependent.txt
```

If we run the analysis using the default settings we obtain the output for each tree in the sample (you may notice the analysis takes a little longer, this is because of the extra parameters that have to be estimated under the dependent model).

Dependent model rate parameters										Probabilities of ancestral state at root of tree				
Tree No	Lh	q12	q13	q21	q24	q31	q34	q42	q43	Root - P(0,0)	Root - P(0,1)	Root - P(1,0)	Root - P(1,1)	
28	1	-34.9232	2.123784	1.55736	5.749772	0.33565	47.03867	112.5859	2.04045	0	0.5347	0.295479	0.141713	0.028108
29	2	-34.3585	3.737294	0	4.417328	6.628631	15.49355	0	0	0	0.408657	0.066458	0.524885	0
30	3	-35.6457	1.683824	1.797969	3.482494	0	43.35764	68.09574	1.905618	0	0.72878	0.123207	0.145984	0.00203
31	4	-34.3301	1.450907	1.794687	6.077645	0	50.08663	144.8653	2.909977	0	0.505711	0.349103	0.106413	0.038772
32	5	-36.3476	1.502279	2.30143	4.920616	0.143111	75.56404	83.65309	2.858958	0	0.605783	0.187309	0.196396	0.010513
33	6	-35.0151	1.579096	1.227233	5.490283	1.484488	55.5164	77.0724	2.461743	0	0.452412	0.247906	0.263981	0.035701
34	7	-33.8409	3.676165	0	6.870779	3.768068	24.72041	0	1.82524	0	0.365901	0.226361	0.384413	0.023325
35	8	-35.1885	1.971772	2.108792	5.647048	0	95.35246	149.6636	2.123764	0	0.58673	0.219332	0.182053	0.011885
36	9	-34.9591	1.159434	1.228314	7.520798	0	39.61688	82.59376	3.453712	0	0.260472	0.083191	0.460317	0.19602
37	10	-36.3958	2.162058	1.569816	7.033134	1.014626	86.16544	115.4437	2.112017	0.00006	0.426394	0.339508	0.187164	0.046934

Figure 7. Results of analysis using a dependent model of evolution for first 10 trees in the treefile, displayed using Excel. The 3rd to 10th columns are the estimated rates of change. The 11th column (Root - P(0,0)) denotes the probability that Trait 1 and Trait 2 were both zero at the root of the tree. Columns 12-14 do the same for the other combinations of trait states (i.e. 0,1; 1,0; 1,1).

Again we see variability from tree to tree in likelihood scores and parameter values (Figure 7). This time the likelihood scores are around -34 log-units. In order to assess whether the dependent model is better than the independent model we can compare the likelihoods of the different for each tree using a likelihood ratio test. The likelihood ratio (LR) statistic is calculated as follows:

$$LR = -2(\log\text{-likelihood}(H_0 \text{ model}) - \log\text{-likelihood}(H_1 \text{ model})) \quad (1)$$

As we noted in the chapter the LR is often taken to approximate a chi-squared distribution with degrees of freedom equal to the difference in the number of parameters between the two models. In this case the degrees of freedom would be 4 (the dependent model has 8 rate parameters and the independent model has 4). The critical value at which we begin to obtain p-values less than 0.05 is 9.49. In these analyses the difference between the models is around 6-9 log-units, and gives LR statistics that range between around 12 and 18. These values are well above the critical value meaning that the dependent model is consistently significantly better than the independent. Unfortunately there is no way to combine maximum likelihood results from a sample of trees like this to come up with an overall assessment of the support for one model over the other, this can only be done in the Bayesian analyses.

Using command files in BayesTraits

To save yourself unnecessary typing you can use a command file to automatically run analyses (this is particularly good if you have a number of commands to enter each time).

Example 4 creating and using command files

Open up your text editor. The first line will be the number of the analysis we want to run from the first menu in BayesTraits (1-6). We want to run DISCRETE so type 2 or 3 on the first line. The second line refers to the second menu of BayesTraits (Maximum Likelihood or MCMC analysis). Type 1 for a ML analysis. In the following lines we can add any other commands such as `lf...` (see Figure 8). Finally at the end type `run`. Save this file (something like `command_file.txt`), and then to run the commands from the command file instead of the keyboard when you run BayesTraits again type:

```
BayesTraits.exe Primates.trees Primates.txt < command_file.txt
```




Figure 8. Example of a command file that can be used in BayesTraits

An introduction to Bayesian MCMC analyses using a sample of phylogenetic trees

In BayesTraits the same analyses described above can also be done within a Bayesian framework. Importantly, as discussed in the book chapter, this approach allows us to integrate the analyses over a sample of phylogenetic trees.

Example 5 Running a dependent model in MCMC mode

Go back to your command prompt and start BayesTraits again (Don't use the input command file at this stage). Let's run a dependent model of evolution but this time using the Bayesian MCMC framework. On the second BayesTraits menu choose MCMC (option 2)(not Maximum Likelihood). Notice the different options and parameters


```

C:\> Command Prompt - BayesTraits.exe Primates.trees Primates.txt
Model: Discete Dependent
Tree File Name: Primates.trees
Data File Name: Primates.txt
Log File Name: Primates.txt.log.txt
Summary: False
Analysis Type: MCMC
Sample Period: 100
Iterations: 5050000
Burn in: 50000
Rate Dev: 2.000000
No of Rates: 8
Base frequency (PI's) None
Character Symbols 00,01,10,11
Using a covarion model: False
Restrictions:
  q12 None
  q13 None
  q21 None
  q24 None
  q31 None
  q34 None
  q42 None
  q43 None
Prior Information:
  Prior Categories: 100
  q12 uniform 0.00 100.00
  q13 uniform 0.00 100.00
  q21 uniform 0.00 100.00
  q24 uniform 0.00 100.00
  q31 uniform 0.00 100.00
  q34 uniform 0.00 100.00
  q42 uniform 0.00 100.00
  q43 uniform 0.00 100.00
Tree Information
  Trees: 500
  Taxa: 60
  Sites: 1
  States: 4

```

Figure 9. Options and settings for a dependent model Discrete analysis in BayesTraits

As the chapter discussed, in a Bayesian analysis, rather than getting a single set of parameter values we get a sample of values that have been drawn from the posterior distribution. This distribution is sampled using the Markov Chain Monte Carlo (MCMC) method. This is a chain that visits parameter values in proportion to their likelihood (i.e. those values that return higher likelihoods are visited more often). Most of the options in this menu relate to properties of the chain so that it ends up sampling from the appropriate region of parameter space as efficiently as possible. The important options are explained in Table 1, and relate to issues discussed in the chapter.

Option	Explanation
Sample	Sample Period: The number of steps in the chain between samples being taken (sampling at a sensible frequency prevents autocorrelation)
Iterations	The number of steps that the chain runs for (this needs to be long enough for the chain to converge on the posterior distribution, to produce an independent sample of parameter estimates, and for the Harmonic mean to stabilize – see below)
Burnin	The number of iterations before samples are taken (this avoids taking samples before the chain has reached convergence, i.e. the burn-in period)
Number of rates	The number of rate parameters in the model to be estimated (in the default dependent model this is 8, in the independent model this is 4; restricting rates to 0 or to be equal to other rates reduces the number

	of rates estimated – see below)
Using a covarion model	This option allows the rates of change to switch between fast and slow over the tree, and therefore relaxes the assumption of a constant rate of change. To turn on the covarion, type <code>cv</code>
Restrictions	Tells you if rate parameters have been set to take a value of 0, or have been set to be equal to each other
Prior information	Prior distributions to be used for each rate parameter

Table 1. Different options in the MCMC version of analyses of Discrete character in *BayesTraits*

In practical terms, one important difference that you need to be aware of is the use of priors. As was discussed in the main chapter we need to specify the prior distribution of the parameters we are estimating. We need to be careful as sometimes the choice of prior can determine the result. How restrictive your prior distribution needs to be depends on the strength of the signal in your data. Ideally you want to use an uninformative prior (this is known as uniform prior and is the default in *BayesTraits*). You can use the information from the Maximum Likelihood analyses to help set your prior distribution. The *BayesTraits* manual that accompanies the software deals with this issue in more detail. At the moment let's just start using the default option of uniform distributions on the range 0-100.

We need to sample from the MCMC at appropriate intervals. If samples are too close together then there will be autocorrelation between your samples, meaning that samples in the chain are not independent (i.e. you can predict the next value in the chain based on the previous value). You can check this using a spreadsheet or another statistical package. You can change this using the sample command (see below).

Let's start by using the default values, and type `run` to start the analysis. After a short pause you should see lots of output appearing quickly. (NB: at any point you can stop the analysis by pressing the `Ctrl` and `C` keys together. Because they produce lots of output it becomes particularly useful to view the output of MCMC analyses in a spreadsheet program (e.g. Excel), or a specialized program such as Tracer that accompanies the BEAST phylogenetic inference software (Drummond & Rambaut 2007).

In a real analysis you will have to spend time adjusting these parameters and it can be quite tricky. However, to get you started here are some sensible values. Open up your input command file. Change the number on the second line to 2 so that the program will perform a MCMC analysis. Write the following lines:

```
3
2
Sample 5000
priorall exp 5
run
```

The command `priorall exp 5` is telling *BayesTraits* to run an analysis using the same prior for each rate parameter. A prior distribution has specific distribution parameters (e.g. an exponential distribution has one such parameter, a gamma distribution has two). Here we are telling *BayesTraits* to use an exponential prior distribution with a mean parameter taking a

value of 5. It is also perfectly possible to set different priors for different parameters using commands such as `prior q12 exp 3`, etc.). Another option is to use a hyper prior, which is less restrictive in that the distribution parameters of the prior themselves can vary (they are seeded from a uniform distribution between specified values). For example, you can give the command `hp q12 exp 0 30` (this command is telling the analysis to select parameter values from exponential prior distributions that can have a mean values anywhere between 0 and 30).

After typing in the above commands allow the program to run and then examine the output in your spreadsheet program. The various output columns are explained in Table 2 below. Notice the variation in likelihood and parameter estimates from each sample of the MCMC. To see if the chain is sampling well from the posterior distribution we can create a line plot of the likelihood from each sample of the chain (Figure 10). If the chain is sampling from the posterior distribution well, it should show no trend towards increasingly higher likelihoods but should instead be fluctuating between higher and lower values. You may not see the initial hill-climbing phase of the MCMC because the `burnin` command tells BayesTraits to return samples from the chain only after a certain number of iterations. To confirm that the chain really has converged on the posterior distribution you should run the analysis several times and see if the different chains produce similar distributions of likelihoods and parameter estimates.

Column	Explanation
Iteration	The point in the chain at which the sample was taken
Lh	Log likelihood of the model
Harmonic Mean	The running Harmonic mean of the likelihoods. This is used as an approximation of the marginal mean and the final value is used in hypothesis testing (see below)
Tree No	The tree used in the calculation for each sample. The MCMC should sample all trees with approximately equal probability.
q12-q43	Values of rate parameters
Root	Estimated probability of each state combination at the root of the tree (other nodes can be estimated using the <code>addNode</code> or <code>addMRCA</code> commands – see the BayesTraits manual)
q12-q43 means	NB: extra variables such as these are only shown when using hyperpriors. These are the values used to seed the prior distribution (again these should well distributed and not pushing up against an upper or lower boundary)

Table 2. Variables from the results of Bayesian MCMC analyses and their meaning




Figure 10. Line plot of likelihoods from samples of the MCMC. The likelihoods fluctuate between lower and higher values indicating that the chain has converged on the posterior distribution (the initial “burn-in” or “hill-climbing” phase of the chain is not seen here because the chain only begins sampling after the number of iterations indicated by the `burnin` command).

Comparing models using the Harmonic Mean

In a Bayesian framework we compare alternative models of evolution using the marginal likelihood of the different models. This is difficult to estimate but can be approximated using the Harmonic mean of the likelihoods of the samples from your chain. You take the value of the harmonic mean from the final iteration of your chain. It therefore needs to be stable (i.e. the harmonic mean should not change much after a period of time). You can check this by plotting out the harmonic mean in your spreadsheet program (see below). Try running an independent model using the MCMC method and compare the likelihoods and Harmonic means with that of the dependent model. Figure 11 shows the plot of the harmonic mean for the dependent model (blue) and the independent model (red). The harmonic mean sometimes takes large drops when the MCMC chains samples a lower likelihood value. Here it can be seen that the Harmonic mean in both runs is reasonably stable and not changed much in the second half of the analyses (although ideally we might want to check this does not change too much more by running longer analyses). The difference between independent and dependent models is around 5 log-units. As we discussed in the main chapter this value is doubled to calculate the Bayes Factor. A value of 10 on this scale relates to strong evidence in favour of the dependent model.

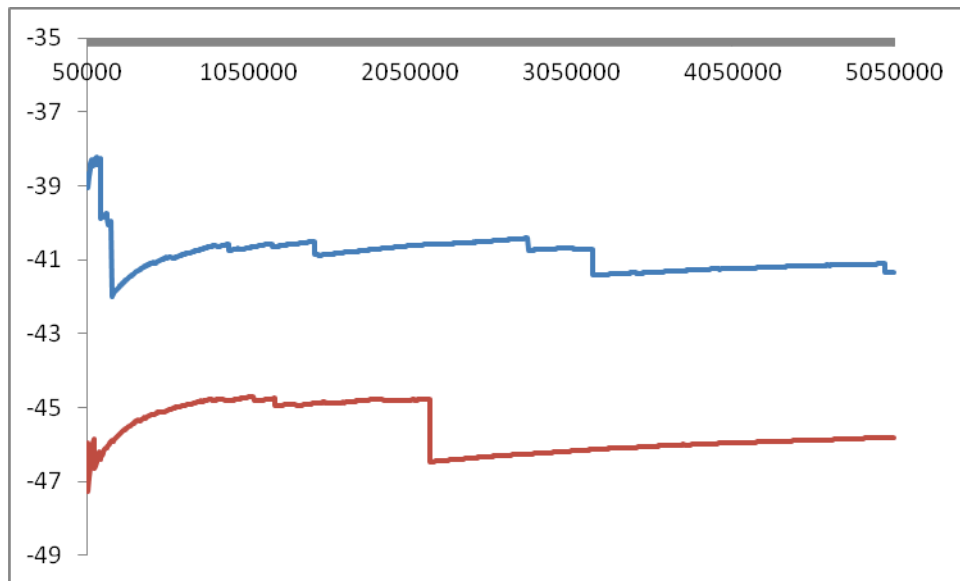


Figure 11. Line plot of the Harmonic means of the Independent model (Red) and the Dependent model (blue).

Examining rates of change

Having discerned that the dependent model is a better fit than the independent model we can discern more information about the evolutionary process by examining the rates of change. It is often useful to visualize these rates in the form of a flow diagram. It should be noted that the value of these rates depends on the branch lengths of the phylogenetic trees being used and should only be interpreted relative to each other (NB: however, values of 0 always indicate that this particular change is actually zero, i.e. it has not occurred). The figure below shows the dependent model flow diagram with the mean rates of change associated with each arrow. It can be seen that the highest rates of change are those going away from the state 1,0 (i.e. oestrus advertisement, non-multimale systems). This suggests that this state is unstable, and indeed none of the species in this dataset have this combination of traits. It is also possible to examine the order in which a change from a state of the absence of oestrus advertisement and a lack of multimale systems (0,0) to the presence of both of these traits (1,1) occurs. The value of q_{12} is greater than q_{13} suggesting that a change from a single male system to a multimale system occurs at a higher rate than a change to oestrus advertisement. This could be tested explicitly by setting q_{13} to zero, or by restricting q_{12} to be equal to q_{13} (see (Pagel 1994) for more information). In BayesTraits this is done by using the `restrict` command, e.g.:

```
restrict q12 q13
res q13 0
```

In effect these restriction commands create alternative evolutionary hypotheses that can be assessed using Bayes Factors in the same way as with the full unrestricted dependent and the independent models. This can give some indication about possible evolutionary pathways and the order in which traits change. It is important to note that to make strong claims about the order in which traits have changed it is advisable to explicitly infer the states of traits at

ancestral nodes in the tree. This enables us to see where and when traits changed over the evolutionary history of the species being analysed. This issue is beyond the scope of this introductory practical except to say that it involves giving BayesTraits commands to either return the proportional probabilities of different trait states at a defined node. Alternatively, analyses can be run with that node constrained to be first in one of the states, then the other, and then comparing the likelihoods of these two alternative scenarios via Bayes Factors.




Figure 12. Flow diagram of the dependent model of evolution with mean values of the estimated rate parameters.

Tips for running MCMC analyses in BayesTraits

More information about how to use BayesTraits can be gained from the manual that accompanies the program, including how to reconstruct ancestral states at nodes in the trees, and the Reversible-Jump (RJ) implementation of Discrete that was mentioned in the chapter (which samples the models of evolution themselves)(see (Pagel & Meade 2006).

To help with future analyses the following checklist should serve as a guide to conducting reliable analyses in a Bayesian framework:

- 1) examine the effects of different priors. Uniform priors have the fewest assumptions, and hyperpriors are less restrictive than regular priors.
- 2) check the posterior distribution of rate values (and prior distribution parameters where relevant) to make sure they are not bumping up against a boundary imposed by your prior
- 3) set the `sample` option so that there is no significant autocorrelation between samples
- 4) the analysis should visit all trees in the sample with roughly equal probability
- 5) check that the burnin phase of the MCMC has finished, and that the MCMC has converged, which leads to...
- 6) run several analyses to check stability between runs in terms of likelihoods, and parameter values
- 7) plot the harmonic mean is make sure it is stable as it can be hard to estimate
- 8) The program was designed to be flexible and will allow you to create a variety of models and perform a range of analyses. It is best to be guided by theory wherever possible as it is possible to perform analyses or reach conclusions that may not be meaningful biologically.

9) The model parameters are estimated from the data, so it is important to make sure you have sufficient data to estimate the number of parameters required in your model.

Good luck with your analyses!

References

- Drummond, A. J. & Rambaut, A. 2007 BEAST: Bayesian evolutionary analysis by sampling trees. *BMC Evolutionary Biology* **7**, 214.
- Maddison, W. P. & Maddison, D. R. 2009 Mesquite: a modular system for evolutionary analysis.
- Pagel, M. 1994 Detecting Correlated Evolution on Phylogenies - A General-Method for the Comparative-Analysis of Discrete Characters. *Proceedings of the Royal Society of London Series B-Biological Sciences* **255**, 37-45.
- Pagel, M. & Meade, A. 2006 Bayesian analysis of correlated evolution of discrete characters by reversible-jump Markov chain Monte Carlo. *American Naturalist* **167**, 808-825.