

Supplementary File 5: *Tutorial for MCMCglmm version*

Tutorial 1 (MCMCglmm) - Estimating the heritability of birth weight

This tutorial will demonstrate how to run a univariate animal model using the R package `MCMCglmm` and example data files provided. Additionally, this module provides some information that applies to `MCMCglmm`-based analyses in general, but that will not be included in other tutorials. Most importantly, this applies to some of the simplest ways of determining the performance of a run using `MCMCglmm`, i.e., verification of the validity of the posterior distribution. This tutorial is not a substitute for working through the `MCMCglmm` tutorial, which is available from CRAN (the Comprehensive R Archive Network, <http://cran.r-project.org/>, or can be accessed in R using the command `vignette("Tutorial", "MCMCglmm")`). These tutorials (1-3) do not introduce one of the main advantages of using `MCMCglmm` for analyses of data from natural populations - the ability to properly model non-normal responses. These capabilities are introduced in the tutorial that is distributed with `MCMCglmm`, and available from CRAN.

Scenario

In a population of gryphons there is strong positive selection on birth weight with heavier born individuals having, on average higher fitness. To find out whether increased birth weight will evolve in response to the selection, and if so how quickly, we want to estimate the heritability of birth weight.

Data files

Open `gryphonped.txt` and `gryphon.txt` in your text editor. The structure and contents of these files is fairly self-explanatory. The pedigree file `gryphonped.txt` contains three columns containing unique identifiers that correspond to each animal, its father, and its mother. Note that this is a multigenerational pedigree, with the earliest generation (for which parentage information is necessarily missing) at the beginning of the file. For later born individuals maternal identities are all known but paternity information is incomplete (a common situation in real world applications).

The phenotype data, as well as additional factors and covariates that we may wish to include in our model are contained in `gryphon.txt`. Columns correspond to individual identity (`ANIMAL`), maternal identity (`MOTHER`), year of birth (`BYEAR`), sex (`SEX`, 1=female, 2=male), birth weight (`BWT`), and tarsus length (`TARSUS`). Each row of the data file contains a record for a different individual.

`gryphonped.txt` and `gryphon.txt` were prepared specifically for use with `ASREML`. However, tab delimited files such as these are easily read into R. The following code reads the data in the two files into data frames called `Ped` and `Data` (the names are arbitrary, almost any character string will do for names). This code also modifies some of the structure of the data in some simple ways. First, we rename the column originally called `ANIMAL` to lower case, as the lower case "animal" is a reserved word for `MCMCglmm` that

we will use when fitting animal models. This code also ensures that R recognizes those data that we will use as factors and numeric data as such. The `head()` commands show us just the first few lines of the phenotypic and pedigree data frames, just so that we can reassure ourselves that we've read in and handled the data properly. We could view the entire data frames by typing their names and pressing return.

```
> Data<-as.data.frame(read.table(file=
+                           "~/Desktop/JAE_MCMCglimm/gryphon.txt",header=TRUE))
> names(Data)[1]<-"animal"
> Data$animal<-as.factor(Data$animal)
> Data$MOTHER<-as.factor(Data$MOTHER)
> Data$BYEAR<-as.factor(Data$BYEAR)
> Data$SEX<-as.factor(Data$SEX)
> Data$BWT<-as.numeric(Data$BWT)
> Data$TARSUS<-as.numeric(Data$TARSUS)
> head(Data)
```

	animal	MOTHER	BYEAR	SEX	BWT	TARSUS
1	1029	1145	968	1	10.77	24.77
2	1299	811	968	1	9.30	22.46
3	643	642	970	2	3.98	12.89
4	1183	1186	970	1	5.39	20.47
5	1238	1237	970	2	12.12	NA
6	891	895	970	1	NA	NA

```
> Ped<-as.data.frame(read.table(file=
+                           "~/Desktop/JAE_MCMCglimm/gryphonped.txt",header=TRUE))
> for(x in 1:3) Ped[,x]<-as.factor(Ped[,x])
> head(Ped)
```

	ID	FATHER	MOTHER
1	1306	<NA>	<NA>
2	1304	<NA>	<NA>
3	1298	<NA>	<NA>
4	1293	<NA>	<NA>
5	1290	<NA>	<NA>
6	1288	<NA>	<NA>

This code is intended to run in a Linux/Unix environment. The only difference on a PC is that the file location would have to be specified slightly differently. For example if the files were in a folder called JAE_MCMCglimm on the root C: drive, the two lines above that read the data might be

```
Data<-as.data.frame(read.table(file=
+                           "C:\JAE_MCMCglimm\gryphon.txt",header=TRUE))
Ped<-as.data.frame(read.table(file=
+                           "C:\JAE_MCMCglimm\gryphonped.txt",header=TRUE))
```

Similarly on a Mac (OSX) if the files were in a folder called JAE_MCMCglmm in the documents area, the two lines above that read the data might be

```
Data<-as.data.frame(read.table(file=
  "//users//name/documents/JAE_MCMCglmm/gryphon.txt",header=TRUE))
Ped<-as.data.frame(read.table(file=
  "//users//name/documents/JAE_MCMCglmm/gryphonped.txt",header=TRUE))
```

Running the model

First load MCMCglmm:

```
> library(MCMCglmm)
```

```
pscl 1.03          2008-11-24
```

You should check that you have the most current version of MCMCglmm with the command `help(package=MCMCglmm)`. You can check the number of the current version on CRAN. If you need to update (or install) MCMCglmm, use `install.packages()` and follow the prompted instructions.

The first model we will fit is a simple animal model with no fixed effects, and only an “animal” random effect relating individuals to their additive genetic values through the pedigree. First we are going to define priors. In a way we might want to avoid using priors, because we would like all of the information in our analysis to come from our data. By default MCMCglmm uses improper priors, but this can cause inferential and numerical problems. We will specify a prior, but we will tell MCMCglmm to pay very little heed to its specific value. We will specify priors for the animal effect and the residual variance using the following code:

```
> p.var<-var(Data$BWT,na.rm=TRUE)
> prior1.1<-list(G=list(
+   G1=list(V=matrix(p.var/2),n=1)),
+   R=list(V=matrix(p.var/2),n=1))
```

We have provided a guess for the animal (the G1 structure) and residual effects (the R structure) by dividing the observed phenotypic variation in birth weight in half (the V terms). We have made this assumption for the purposes of this tutorial, but in reality careful consideration of priors is necessary. We have told MCMCglmm to pay little heed to these numbers by specifying degree of belief parameters (`n`) of 1. Since this is a univariate analysis, the priors are matrices of order 1 and thus $n = 1$ is the smallest degree of belief that provides what is known as a “proper” prior, avoiding numerical problems. In fact, there is a lot of information in the data regarding the marginal distributions of the parameters, and MCMCglmm will run most of the models that we suggest in these tutorials without priors. However, this is poor practice, and we will therefore use priors throughout these tutorials.

We can now fit an animal model. The model to decompose variation in birth weight into genetic and residual effects is as follows:

```
> model1.1<-MCMCglmm(BWT~1,random=~animal,  
+                   pedigree=Ped,data=Data,prior=prior1.1)
```

```
MCMC iteration = 0
```

```
MCMC iteration = 1000
```

```
MCMC iteration = 2000
```

```
MCMC iteration = 3000
```

```
MCMC iteration = 4000
```

```
MCMC iteration = 5000
```

```
MCMC iteration = 6000
```

```
MCMC iteration = 7000
```

```
MCMC iteration = 8000
```

```
MCMC iteration = 9000
```

```
MCMC iteration = 10000
```

```
MCMC iteration = 11000
```

```
MCMC iteration = 12000
```

```
MCMC iteration = 13000
```

In addition: Warning messages:

```
1: In MCMCglmm(BWT ~ 1, random = ~animal, pedigree = Ped, data = Data, :  
  some combinations in animal do not exist and 225 missing records have  
  been generated
```

After typing this code, `MCMCglmm` will run, taking about 20 seconds on a modern desktop computer. The R console may temporarily appear to crash under windows. Unless an error message appears, just let it run, the jitters result from poor communication between R and the C routines in `MCMCglmm`. The progress of the run will be printed to the screen. Also, note the warning message will be printed at the end of the run. This is natural too. In order for the MCMC algorithm to work, `MCMCglmm` must keep track of effects associated with unmeasured individuals appearing in the pedigree. This will not affect the answers, but when many unmeasured individuals exist, it can hinder the ability of the algorithm to explore the parameter space (more on this, and a solution, later).

Lets have a look at the `MCMCglmm` outputs. First we will evaluate how confident we can be that `MCMCglmm` found good answers. By entering




Figure 1: The posterior distribution of the fixed effect (the intercept, or mean) in model 1.1.

```
> plot(model1.1$Sol)
```

in the console, we get Figure 1 (p. 5). The plot on the left shows a time series of the values of 1000 samples of the posterior distribution of the the model intercept (mean birthweight). The plot on the right shows the same data as a distribution. Complicated statistical methods for estimating population means are of course of little interest; rather, we are examining these outputs to check that MCMCglmm's algorithms worked well for our data and for this model. The important point here is that a consistent amount of variation around a largely unchanging mean value of the intercept was obtained, and the posterior distribution of the intercept appears to be valid. More rigorous means of evaluation the independence of the samples in the posterior distribution (evaluating autocorrelation) are discussed in the MCMCglmm tutorial, available from CRAN. Note that your output for model 1.1 may not be identical to this due to Monte Carlo (random number) error.

The posterior distributions of the the variance components are generally of more interest to animal model users. We can view plots of the posterior distribution for the variance components for model 1.1 by

```
> plot(model1.1$VCV)
```

which generates Figure 2 (p. 6). Here we see distributions of the estimates of the additive genetic (**animal**) and residual (**units**) effects. These samples contain some autocorrelation, i.e., trends are apparent in the left-hand plot. We can deal with this easily. We will simply re-run the model for a longer number of iterations, and sample the chain

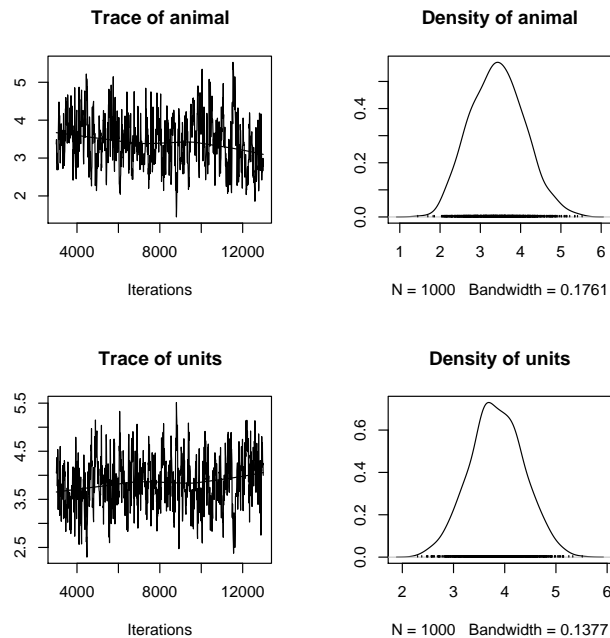


Figure 2: The posterior distributions of the variance components of model 1.1, based on an analysis with the default values for `nitt`, `burnin`, and `thin` in `MCMCglmm`.

less frequently. So far we have been running `MCMCglmm` with its default values. These defaults are a total run length of 13000 iterations, the first 3000 of which are discarded as a “burn-in” period to make sure that model the converges to the part of the parameter space where the maximum likelihood exists. The remaining 10000 iterations are sampled (estimates retained) every 10 iterations (the thinning interval). Because the values in the left-hand plots in figure 2 to appear to have different values at the beginning of the run, we might suspect that a longer burn-in period might be required. We can reduce the autocorrelation by lengthening the rest of the run and sampling the chain less frequently. The following code runs the same model 1.1, but is likely to produce better samples of the posterior distributions. This model should take about two minutes to analyze.

```
> model1.1<-MCMCglmm(BWT~1,random=~animal,
+                   pedigree=Ped,data=Data,
+                   nitt=65000,thin=50,burnin=15000,
+                   prior=prior1.1,verbose=FALSE)
```

Notice that we have now included the command `verbose=FALSE` in the `MCMCglmm` call. We will continue this throughout the tutorial so that more complete screen outputs can be included in this document without using too much space.

Now produce the plots of the samples of the fixed and random effects (they have not been included in this document). Note that the autocorrelation is much reduced. A more compact way to evaluate the validity of the posterior distributions is to calculate autocorrelation among samples, as follows:

```
> autocorr(model1.1$VCV)
, , animal
      animal      units
Lag 0  1.000000000 -0.796134942
Lag 50  0.230652982 -0.205946426
Lag 250 -0.009709827  0.001097322
Lag 500 -0.012504144  0.008050763
Lag 2500 -0.026189620  0.028163506
, , units
      animal      units
Lag 0  -0.796134942  1.000000000
Lag 50  -0.192836753  0.16383254
Lag 250  0.015065923 -0.01274703
Lag 500 -0.006800833 -0.01250495
Lag 2500  0.011141997 -0.01341721
```

We will consider these levels of autocorrelation acceptable, at least for the purposes of this tutorial. Ideally, all samples of the posterior distribution should be independent, and the autocorrelation for all lag values greater than zero should be near zero. However, in practice this will not strictly be achievable for all analytical scenarios. Certainly the levels of autocorrelation observed here should not be tolerated in any formal analysis. Note that the validity of posterior distributions of any analysis should always be checked; however, for brevity we will not continue to be so consistently diligent throughout the rest of these tutorials. We can now proceed with confidence to recover some more information from these samples. We can obtain estimates of the additive genetic and residual variance by calculating the modes of the posterior distributions:

```
> posterior.mode(model1.1$VCV)
      animal      units
3.380911 3.732298
```

We can obtain the Bayesian equivalent of confidence intervals by calculating the the values of the estimates that bound 95% (or any other proportion) of the posterior distributions:

```
> HPDinterval(model1.1$VCV)
      lower      upper
animal 2.079416 4.492744
units  2.958496 4.880523
attr(,"Probability")
[1] 0.95
```

We specified weak priors in this analyses. Now we will check whether or not these priors influenced the results that we obtained. The simplest way to do this is to rerun the model with different priors. Rather than splitting the observed phenotypic variance evenly between the priors for the genetic and residual effects, we will construct priors with the same degree of belief parameters, but we will specify a much larger proportion of genetic control:

```
> prior1.1.2<-list(G=list(
+       G1=list(V=matrix(p.var*0.95),n=1)),
+       R=list(V=matrix(p.var*0.05),n=1))
> model1.1.2<-MCMCglmm(BWT~1,random=~animal,
+       pedigree=Ped,data=Data,prior=prior1.1.2,
+       nitt=65000,thin=50,burnin=15000,verbose=FALSE)
> posterior.mode(model1.1$VCV)

  animal    units
3.380911 3.732298

> posterior.mode(model1.1.2$VCV)

  animal    units
3.375395 3.756474
```

and we can therefore conclude that the prior has little effect on the outcome of the analysis.

Estimating heritability

A useful property of Bayesian posterior distributions is that we can apply almost any transformation to these distributions and they will remain valid. This applies to the calculation of heritabilities. We can obtain an estimate of the heritability by applying the basic formula $h^2=V_A/V_P$ to each sample of the posterior distribution:

```
> posterior.heritability1.1<-model1.1$VCV[,"animal"]/
+       (model1.1$VCV[,"animal"]+model1.1$VCV[,"units"])
> HPDinterval(posterior.heritability1.1,0.95)

      lower    upper
var1 0.3108154 0.6043704
attr(,"Probability")
[1] 0.95

> posterior.mode(posterior.heritability1.1)

  var1
0.4580992
```

Generate a plot of the posterior distribution of this heritability estimate (Figure 3, p. 9):

```
> plot(posterior.heritability1.1)
```





Figure 3: The posterior distributions the heritability from model 1.1.

Adding fixed effects

To add effects to a univariate model we simply modify the fixed effect portion of the the model specification:

```
> model1.2<-MCMCglmm(BWT~SEX,random=~animal,
+                   pedigree=Ped,data=Data,prior=prior1.1,
+                   nitt=65000,thin=50,burnin=15000,verbose=FALSE)
```

We can assess the significance of sex as a fixed effect by examining its posterior distribution.

```
> posterior.mode(model1.2$Sol[, "SEX2"])

      var1
2.209543

> HPDinterval(model1.2$Sol[, "SEX2"], 0.95)

      lower      upper
var1 1.912343 2.531451
attr(,"Probability")
[1] 0.95
```

The posterior distribution of the SEX2 term does not overlap zero. Thus the we can infer that sex has a statistical effect on birthweight in this model and is a useful addition

to the model, for most purposes. `MCMCglmm` has designated `SEX2` as the contrast between the two factor levels (male and female).

It is also worth noting that the variance components have changed slightly:

```
> posterior.mode(model1.2$VCV)

  animal    units
2.941984 3.009324
```

In fact since `SEX` effects were previously contributing to the residual variance of the model our estimate of V_R (denoted 'units' in the output) is now slightly lower than before. This has an important consequence for estimating heritability since if we calculate V_P as $V_A + V_R$ then as we include fixed effects we will soak up more residual variance driving V_P . Assuming that V_A is more or less unaffected by the fixed effects fitted then as V_P goes down we expect our estimate of h^2 will go up.

```
> posterior.heritability1.2<-model1.2$VCV[,"animal"]/
+   (model1.2$VCV[,"animal"]+model1.2$VCV[,"units"])
> posterior.mode(posterior.heritability1.2)

  var1
0.5005594

> HPDinterval(posterior.heritability1.2,0.95)

      lower      upper
var1 0.3714756 0.6529849
attr(,"Probability")
[1] 0.95
```

Here h^2 has increased slightly from 0.4581 to 0.5006 (again, your values may differ slightly due to Monte Carlo error). Which is the better estimate? It depends on what your question is. The first is an estimate of the proportion of variance in birth weight explained by additive effects, the latter is an estimate of the proportion of variance in birth weight *after conditioning on sex* that is explained by additive effects.

Adding random effects

This is done by simply modifying the model statement in the same way, but requires addition of a prior for the new random effect. For instance, we can fit an effect of birth year:

```
> prior1.3<-list(G=list(G1=list(V=matrix(p.var/3),n=1),
+   G2=list(V=matrix(p.var/3),n=1)),
+   R=list(V=matrix(p.var/3),n=1))
> model1.3<-MCMCglmm(BWT~SEX,random=~animal+BYEAR,
+   pedigree=Ped,data=Data,
+   nitt=65000,thin=50,burnin=15000,
+   prior=prior1.3,verbose=FALSE)
> posterior.mode(model1.3$VCV)
```

```

  animal    BYEAR    units
2.5703970  0.8418539  2.3566518

```

Here the variance in birth weight explained by birth year is 0.8419. Note that although V_A has changed somewhat, most of what is now partitioned as a birth year effect was previously partitioned as V_R . Thus what we have really done here is to partition environmental effects into those arising from year to year differences versus everything else, and we do not really expect much change in h^2 (since now $h^2 = V_A/(V_A + V_{BY} + V_R)$).

However, we get a somewhat different result if we also add a random effect of MOTHER to test for maternal effects:

```

> prior1.4<-list(G=list(G1=list(V=matrix(p.var/4),n=1),
+                       G2=list(V=matrix(p.var/4),n=1),
+                       G3=list(V=matrix(p.var/4),n=1)),
+               R=list(V=matrix(p.var/4),n=1))
> model1.4<-MCMCglmm(BWT~SEX,random=~animal+BYEAR+MOTHER,
+                   pedigree=Ped,data=Data,
+                   nitt=65000,thin=50,burnin=15000,
+                   prior=prior1.4,verbose=FALSE)
> posterior.mode(model1.4$VCV)

  animal    BYEAR    MOTHER    units
2.2902304  0.8380849  1.1363555  1.7368250

```

Here partitioning of significant maternal variance has resulted in a further decrease in V_R but also a decrease in V_A . The latter is because maternal effects of the sort we simulated (fixed differences between mothers) will have the consequence of increasing similarity among maternal siblings. Consequently they can look very much like additive genetic effects and if present, but unmodelled, represent a type of “common environment effect” that can - and will- cause upward bias in V_A and so h^2 . Let's compare the estimates of heritability from each of models 1.2, 1.3 and 1.4:

```

> posterior.heritability1.3<-model1.3$VCV[,"animal"]/
+   (model1.3$VCV[,"animal"]+model1.3$VCV[,"BYEAR"]+model1.3$VCV[,"units"])
> posterior.heritability1.4<-model1.4$VCV[,"animal"]/
+   (model1.4$VCV[,"animal"]+model1.4$VCV[,"BYEAR"]+
+     model1.4$VCV[,"MOTHER"]+model1.4$VCV[,"units"])
> posterior.mode(posterior.heritability1.2)

  var1
0.5005594

> posterior.mode(posterior.heritability1.3)

  var1
0.4241203

> posterior.mode(posterior.heritability1.4)

  var1
0.3690422

```

Testing significance of variance components

While testing the significance of fixed effects by evaluating whether or not their posterior distributions overlap zero was simple and valid, this approach does not work for variance components. Variance components are bound to be positive (given a proper prior), and thus even when a random effect is not meaningful, its posterior distribution will never overlap zero. We can test the statistical significance of variance components using the deviance information criterion (DIC). It should be noted that the properties of DIC are not as well understood. The implementation of DIC in `MCMCglmm` is further described in the reference manual. DIC values are calculated by `MCMCglmm` by default. Briefly, DIC balances model fit and parameter number simultaneously, and small values of DIC are preferred. We can test the statistical significance of the maternal effect by comparing models 1.4 and 1.3, i.e., models with and without the mother term:

```
> model1.3$DIC
```

```
[1] 3551.41
```

```
> model1.4$DIC
```

```
[1] 3326.87
```

model 1.4 has a much lower DIC value. Since the maternal effect term is the only difference between the models, we can consider the inclusion of this term statistically justifiable. We should note however that DIC has a large sampling variance and should probably only be calculated based on much longer MCMC runs.

Tutorial 2 (MCMCglmm) – A bivariate animal model

This tutorial will demonstrate how to run a multivariate animal model using the R package `MCMCglmm` and example data files provided.

Scenario

Since natural selection rarely acts on single traits, to understand how birth weight might evolve in our population of gryphons, we may also want to think about possible covariance with other traits. If tarsus length at fledging is also under positive selection what implications does this have for birth weight and vice versa? If the two traits are positively genetically correlated then this will facilitate evolution of larger size (since response of one trait will induce a positively correlated response in the other). If there is negative genetic covariance then this could act an evolutionary constraint.

Using multivariate models allows the estimation of parameters relating to each trait alone (i.e. V_A , h^2 , etc.), but also yields estimates of covariance components between traits. These include the (additive) genetic covariance COV_A which is often rescaled to give the genetic correlation r_G . However, covariance can also arise through other random effects (e.g. maternal covariance) and these sources can be explicitly modelled in a bivariate analysis.

Data

Pedigree and phenotypic data files are the same as those used in tutorial 1 (i.e, `gryphonped.txt` and `gryphon.txt` respectively). See tutorial 1, section 2 for information on loading these data into R. Note that if you are continuing from tutorial number 1 that these data will still be loaded and no further processing of the raw data is required at this point.

Fitting the model

Fitting a multivariate model in `MCMCglmm` involves several new consideration above those for fitting univariate models. First, we have to fit multivariate priors; second, we have to specify the ways in which effects on different traits may covary, including the nature of residual (co)variation; and third, we will have to be a little more specific when specifying to `MCMCglmm` what type of distributions from which we assume our data are drawn.

Our most basic model can be specified as:

```
> phen.var<-matrix(c(var(Data$BWT,na.rm=TRUE),0,0,
+                    var(Data$TARSUS,na.rm=TRUE)),2,2)
> prior2.1<-list(G=list(G1=list(V=phen.var/2,n=2)),
+               R=list(V=phen.var/2,n=2))
> model2.1<-MCMCglmm(cbind(BWT,TARSUS)~trait-1,
+                    random=~us(trait):animal,
+                    rcov=~us(trait):units,
+                    family=c("gaussian","gaussian"),
+                    pedigree=Ped,data=Data,
+                    prior=prior2.1,verbose=FALSE)

> plot(model2.1$VCV[, "animal.trait.TARSUS.TARSUS"])
```

We have constructed the prior similarly to the those in the univariate models in tutorial 1, only we are specifying a matrix based on the the phenotypic variances (and arbitrarily ignoring covariance). In order to provide proper priors, we have set the degree of belief parameter to the order of the matrices describing the priors, i.e., 2. We have used the R command `cbind`, “column-bind”, to specify the two response variables, body weight and tarsus length. The nature of genetic and residual (co)variance (between and) of the traits is specified as unstructured matrices, the most general structure available for this type of data. Finally, we have specified that we wish to treat both traits as gaussian responses.

In tutorial 1, we used full autocorrelation tables to evaluate the validity of the posterior distribution. Note that we have not done this here. For a bivariate model this table can become very complex. Nonetheless, it is worth evaluating, rather it is simply too large to include here. It can be viewed in the console as before. Here we have displayed only the autocorrelation for estimates of additive genetic effects for tarsus length with a lag of 50 iterations (five samples given this `MCMCglmm` run with default values). This lag of 0.3697 is clearly unacceptable. The posterior distribution of the additive genetic effect on tarsus length is shown in Figure 4 (p. 14), note the autocorrelation evident in the left-hand plot. We will opt to run the analysis for longer. This longer run could be run using the following code:




Figure 4: The posterior distribution of the additive genetic effect for tarsus length in a MCMCglmm run with default values.

```
model2.1<-MCMCglmm(cbind(BWT,TARSUS)~trait-1,
  random=~us(trait):animal,
  rcov=~us(trait):units,
  family=c("gaussian","gaussian"),
  pedigree=Ped,data=Data,
  nitt=130000,thin=100,burnin=30000,
  prior=prior2.1,verbose=FALSE)
```

```
# the autocorrelation of the genetic variance of TARSUS at Lag 5
autocorr(model2.2$VCV)[,,"animal.trait.TARSUS.TARSUS"][3,4]
```

However, this run would might take as long as an hour. For the purpose of this tutorial we have provided an output for such a run. It can be obtained and manipulated as follows, assuming that the file `~/Desktop/JAE_MCMCglmm/model1point1LongRun.Rdat` is available at the specified location:

```
> model2.1<-dget(file="~/Desktop/JAE_MCMCglmm/model2point1LongRun.Rdat")
> # the autocorrelation of the genetic variance of TARSUS at Lag 5
> autocorr(model2.1$VCV)[,,"animal.trait.TARSUS.TARSUS"][3,4]

[1] -0.03403771
```

This level of autocorrelation is more acceptable, at least for the purpose of demonstration in this tutorial.

We can recover variance components, heritabilities, and genetic correlations from the posterior distribution of this model:

```
> posterior.mode(model2.1$VCV)

      animal.trait.BWT.BWT      animal.trait.TARSUS.BWT
      3.424151                2.253418
animal.trait.BWT.TARSUS animal.trait.TARSUS.TARSUS
      2.253418                12.574780
      units.trait.BWT.BWT      units.trait.TARSUS.BWT
      3.674433                3.222885
units.trait.BWT.TARSUS units.trait.TARSUS.TARSUS
      3.222885                17.909409

> heritability.BWT2.1<-model2.1$VCV[,"animal.trait.BWT.BWT"]/
+      (model2.1$VCV[,"animal.trait.BWT.BWT"]+
+      model2.1$VCV[,"units.trait.BWT.BWT"])
> posterior.mode(heritability.BWT2.1)

      var1
0.465334

> heritability.TARSUS2.1<-model2.1$VCV[,"animal.trait.TARSUS.TARSUS"]/
+      (model2.1$VCV[,"animal.trait.TARSUS.TARSUS"]+
+      model2.1$VCV[,"units.trait.TARSUS.TARSUS"])
> posterior.mode(heritability.TARSUS2.1)

      var1
0.4142762

> genetic.correlation2.1<-model2.1$VCV[,"animal.trait.BWT.TARSUS"]/
+      sqrt(model2.1$VCV[,"animal.trait.BWT.BWT"]*
+      model2.1$VCV[,"animal.trait.TARSUS.TARSUS"])
> posterior.mode(genetic.correlation2.1)

      var1
0.3758
```

Adding fixed and random effects

Fixed and random effects can be added just as for the univariate case. Given that our full model of BWT from tutorial 1 had SEX as a fixed effect as well as random effects of BYEAR and MOTHER we could specify a bivariate formulation of this using the following code:

```
prior2.2<-list(G=list(G1=list(V=phen.var/4,n=2),
                        G2=list(V=phen.var/4,n=2),
                        G3=list(V=phen.var/4,n=2)),
              R=list(V=phen.var/4,n=2))
```

```

model2.2<-MCMCgllmm(cbind(BWT,TARSUS)~trait-1+trait:SEX,
  random=~us(trait):animal+us(trait):BYEAR+us(trait):MOTHER,
  rcov=~us(trait):units,
  family=c("gaussian","gaussian"),
  pedigree=Ped,data=Data,
  nitt=130000,thin=100,burnin=30000,
  prior=prior2.2,verbose=FALSE)

```

Again we have provided the data from one such run. It can be accessed using the code:

```

> model2.2<-dget(file=~/Desktop/JAE_MCMCgllmm/model2point2LongRun.Rdat")
> # the autocorrelation of the genetic variance of TARSUS at Lag 5
> autocorr(model2.2$VCV)[,,"animal.trait.TARSUS.TARSUS"][3,4]

[1] 0.01044468

```

As before we can obtain the raw variance component estimates and genetic correlations for the random effects:

```

> posterior.mode(model2.2$VCV)

  animal.trait.BWT.BWT    animal.trait.TARSUS.BWT
                1.8451734                3.3756017
 animal.trait.BWT.TARSUS animal.trait.TARSUS.TARSUS
                3.3756017                9.4761196
  BYEAR.trait.BWT.BWT    BYEAR.trait.TARSUS.BWT
                1.1043305                0.1952849
 BYEAR.trait.BWT.TARSUS BYEAR.trait.TARSUS.TARSUS
                0.1952849                3.8908216
  MOTHER.trait.BWT.BWT    MOTHER.trait.TARSUS.BWT
                1.1826161                -1.2795865
 MOTHER.trait.BWT.TARSUS MOTHER.trait.TARSUS.TARSUS
                -1.2795865                5.5880575
  units.trait.BWT.BWT    units.trait.TARSUS.BWT
                1.7685080                4.0626904
 units.trait.BWT.TARSUS units.trait.TARSUS.TARSUS
                4.0626904                12.5483773

> genetic.correlation2.2<-model2.2$VCV[,"animal.trait.BWT.TARSUS"]/
+   sqrt(model2.2$VCV[,"animal.trait.BWT.BWT"]*
+   model2.2$VCV[,"animal.trait.TARSUS.TARSUS"])
> maternal.correlation2.2<-model2.2$VCV[,"MOTHER.trait.BWT.TARSUS"]/
+   sqrt(model2.2$VCV[,"MOTHER.trait.BWT.BWT"]*
+   model2.2$VCV[,"MOTHER.trait.TARSUS.TARSUS"])
> posterior.mode(genetic.correlation2.2)

  var1
0.7735107

```



```
> posterior.mode(maternal.correlation2.2)
      var1
-0.4911696
```

Evaluation of the statistical support for these genetic and maternal correlations is straightforward. Because we imposed no constraint on their estimation, we can evaluate the extent to which the posterior distributions overlap zero:

```
> HPDinterval(genetic.correlation2.2,0.95)
      lower      upper
var1 0.5699206 0.8782866
attr(,"Probability")
[1] 0.95
```

```
> HPDinterval(maternal.correlation2.2,0.95)
      lower      upper
var1 -0.7359405 -0.1739729
attr(,"Probability")
[1] 0.95
```

Neither of these posterior distributions overlaps zero, so we can consider them both statistically supported. Alternatively, we could constrain one or both covariances to zero using `idh()` in place of `us()` in the `MCMCglmm` model specifications, and proceed with significance testing using DIC.

Tutorial 3 (MCMCglmm) – A repeated measures animal model

This tutorial will demonstrate how to run a univariate animal model for a trait with repeated observations using the software `MCMCglmm` and example data files provided.

Scenario

Since gryphons are iteroparous, multiple observations of reproductive traits are available for some individuals. Here we have repeated measures of lay date (measured in days after Jan 1) for individual females of varying age from 2 (age of maturation) up until age 6. Not all females lay every year so the number of observations per female is variable. We want to know how repeatable the trait is, and (assuming it is repeatable) how heritable it is.

Data

We use for these examples the same pedigree as was used for tutorials 1 and 2, and we assume that it is still loaded. If this is not the case, revisit tutorial 1, section 2 for instructions for loading the pedigree. We are however going to use different phenotypic data collected from the same population of gryphons, and we need to load these data. We

can load the data from the file `gryphonRM.txt` into an R data frame using the following code

```
> DataRM<-as.data.frame(read.table(file="~/Desktop/JAE_MCMCglmm/gryphonRM.txt",header=
> names(DataRM)[1]<-"animal"
> DataRM$animal<-as.factor(DataRM$animal)
> DataRM$BYEAR<-as.factor(DataRM$BYEAR)
> DataRM$AGE<-as.factor(DataRM$AGE)
> DataRM$YEAR<-as.factor(DataRM$YEAR)
> DataRM$LAYDATE<-as.numeric(DataRM$LAYDATE)
> head(DataRM)
```

	animal	BYEAR	AGE	YEAR	LAYDATE
1	1	990	2	992	19
2	1	990	3	993	23
3	1	990	4	994	24
4	1	990	5	995	23
5	1	990	6	996	29
6	2	990	2	992	21

```
> DataRM$ID<-DataRM$animal
> head(DataRM)
```

	animal	BYEAR	AGE	YEAR	LAYDATE	ID
1	1	990	2	992	19	1
2	1	990	3	993	23	1
3	1	990	4	994	24	1
4	1	990	5	995	23	1
5	1	990	6	996	29	1
6	2	990	2	992	21	2

Note that we have created a new factor called ID that is the same as the 'animal' factor. MCMCglmm will recognize the term animal and use it to relate individuals to their records in a pedigree. The factor ID allows us to disassociate individual records from the pedigree, which is an important part of analyzing repeated measures data, as we will see.

Estimating repeatability

With repeated measures on individuals it is often of interest, prior to fitting a genetic model, to see how repeatable a trait is. We can estimate the repeatability of a trait as the proportion of phenotypic variance explained by individual identity using the commands below

```
> p.var<-var(DataRM$LAYDATE,na.rm=TRUE)
> prior3.1<-list(G=list(G1=list(V=matrix(p.var/2),n=1)),
+               R=list(V=matrix(p.var/2),n=1))
> model3.1<-MCMCglmm(LAYDATE~1,random=~ID,
```

```
+           data=DataRM,prior=prior3.1,verbose=FALSE)
> posterior.mode(model3.1$VCV)
```

```
      ID      units
10.75490 21.44650
```

Note the use of the new ID factor that we created in the previous section.

Between-individual variance is given by the ID component, while the residual component (Variance) therefore represents within-individual variance. Here then the repeatability of the trait can be determined by as 0.333 (i.e., $10.7549/(10.7549+21.4465)$). Given that we set up the simulation such that mean lay date changes with age (initially increasing to age 5 before a late life decline) we might ask what the repeatability of lay date is after conditioning on age effect. This would be done by adding age into the model as a fixed effect.

```
> model3.2<-MCMCglmm(LAYDATE~AGE,random=~ID,
+           data=DataRM,prior=prior3.1,verbose=FALSE)
> plot(model3.2$Sol)
> plot(model3.2$VCV)
> posterior.mode(model3.2$VCV)
```

```
      ID      units
12.42367 16.19491
```

Note that the random effect structure has remained unchanged, and so we have not modified the prior between models 3.1 and 3.2.

So the repeatability of laydate, after accounting for age effects, is now estimated as 0.432 (i.e., $12.4237/(12.4237+16.1949)$). So, just as we saw when estimating h^2 in tutorial 1, the inclusion of fixed effects will alter the estimated effect size if we determine total phenotypic variance as the sum of the variance components. Thus, proper interpretation is vital.

Here age is modelled as a 5 level factor (see the conversion of age to a factor in section 3.2). We could equally have fitted it as a continuous variable instead in which case, given the late life decline, we would probably also include a quadratic term.

Partitioning additive and permanent environment effects

Generally we expect that the repeatability will set the upper limit for heritability since, while additive genetic effects will cause among-individual variation, so will other types of effect. Non-additive contributions to fixed among-individual differences are normally referred to as “permanent environment effects”, although “non-heritable effects” that are consistent within individuals may be a better way to think of modelling this effect. If a trait has repeated measures then it is necessary to model permanent environment effects in an animal model to prevent upward bias in V_A . To illustrate this fit the animal model

```
> model3.3<-MCMCglmm(LAYDATE~1+AGE,random=~animal,
+           pedigree=Ped,data=DataRM,prior=prior3.1,verbose=FALSE)
> posterior.mode(model3.3$VCV)
```

```

  animal    units
13.51516 16.77824

```

This suggests that all of the among-individual variance is –rightly or wrongly – being partitioned as V_A here. In fact here the partition is wrong since the simulation included both additive genetic effects and additional fixed heterogeneity that was not associated with the pedigree structure (i.e. permanent environment effects). An more appropriate estimate of V_A is given by the model:

```

> p.var<-var(DataRM$LAYDATE,na.rm=TRUE)
> prior3.4<-list(G=list(G1=list(V=matrix(p.var/3),n=1),
+                       G2=list(V=matrix(p.var/3),n=1)),
+               R=list(V=matrix(p.var/3),n=1))
> model3.4<-MCMCglmm(LAYDATE~1+AGE,random=~animal+ID,
+                   pedigree=Ped,data=DataRM,prior=prior3.4,verbose=FALSE)
> posterior.mode(model3.4$VCV)

```

```

  animal      ID    units
4.764842 6.883453 16.385166

```

The estimate of V_A is now much lower (reduced from 13.5152 to 4.7648) since the additive and permanent environment effects are being properly separated. We could obtain estimates of h^2 and of the repeatability from this model using the following commands:

```

> # calculate phenotypic variance conditional on age
> model3.4.VP<-model3.4$VCV[,"animal"]+model3.4$VCV[,"ID"]+model3.4$VCV[,"units"]
> # calculate between individual variance
> model3.4.IDplusVA<-model3.4$VCV[,"animal"]+model3.4$VCV[,"ID"]
> # repeatability
> posterior.mode(model3.4.IDplusVA/model3.4.VP)

```

```

  var1
0.4278259

```

```

> # heritability
> posterior.mode(model3.4$VCV[,"animal"]/model3.4.VP)

```

```

  var1
0.1827556

```

Adding additional effects and testing significance

Models of repeated measures can be extended to include other fixed or random effects. For example try including year of measurement (YEAR).

```

> p.var<-var(DataRM$LAYDATE,na.rm=TRUE)
> prior3.5<-list(G=list(G1=list(V=matrix(p.var/5),n=1),
+                       G2=list(V=matrix(p.var/5),n=1),
+                       G3=list(V=matrix(p.var/5),n=1)),

```

```

+           G4=list(V=matrix(p.var/5),n=1)),
+           R=list(V=matrix(p.var/5),n=1))
> model3.5<-MCMCglmm(LAYDATE~1+AGE,
+                   random=~animal+ID+YEAR+BYEAR,
+                   pedigree=Ped,data=DataRM,
+                   prior=prior3.5,verbose=FALSE)
> posterior.mode(model3.5$VCV)

      animal      ID      YEAR      BYEAR      units
4.1747773 8.3037069 7.5496368 0.6672928 7.8099636

```

This model will return additional variance components corresponding to year of measurement effects and birth year (of the female effects). The latter were not simulated as should be apparent from the apparent parameter estimate (or better evaluated from DIC-based comparison of model3.5 and a model from which the birth year term had been eliminated, see tutorial 1). However, `YEAR` effects were simulated as should be apparent from the modal estimate and from the support interval (try this yourself using `HPDinterval()`) and this could be formally confirmed by comparison of DIC.

`YEAR` effects could alternatively be included as fixed effects (try this, you should be able to handle the new prior specification at this point). Since we simulated large year of measurement effects this treatment will reduce V_R and increase the estimates of heritability and repeatability which must now be interpreted as proportions of phenotypic variance after conditioning on both age and year of measurement effects.