

An Introduction to Model-Fitting with the R package `glm`

Christina Knudson

April 5, 2022

Contents

1	Introduction	2
2	Formatting the Data	2
3	Fitting the Model	3
4	Reading the Model Summary	5
5	Isolating the Parameter Estimates	6
6	Calculating Confidence Intervals	7
7	Monte Carlo Standard Error	9
8	Estimating the Variance-Covariance Matrix	9
9	Accessing Additional Output	10
10	Optional Model-Fitting Arguments	11
10.1	Setting Variance Components Equal	11
10.2	Altering the Importance Sampling Distribution	12
10.3	Adjusting Optimization Arguments	12
10.4	Starting at a Specified Parameter Value	13

1 Introduction

The R package `glmm` approximates the entire likelihood function for generalized linear mixed models (GLMMs) with a canonical link. `glmm` calculates and maximizes the Monte Carlo likelihood approximation (MCLA) to find Monte Carlo maximum likelihood estimates (MCMLEs) for the fixed effects and variance components. Additionally, the value, gradient vector, and Hessian matrix of the MCLA are calculated at the MCMLEs. The Hessian of the MCLA is used to calculate the standard errors for the MCMLEs.

The basis of `glmm` is MCLA, which was first proposed by Geyer (1990) for approximating the likelihood of unnormalized densities. MCLA was used by Geyer and Thompson (1992) to approximate the likelihood for normalized densities for models with random effects. Gelfand and Carlin (1993) proposed applying MCLA to unnormalized densities for models with random effects. The theoretical foundation for MCLA was established by Geyer (1994). Sung and Geyer (2007) prepared an R package `bernor` that, when given model matrices, fits maximum likelihood estimates for the logit-normal model. Their importance sampling distribution is chosen independently of the data.

2 Formatting the Data

The following vectors can be used to fit a generalized linear mixed model using the `glmm` package. These vectors can be contained in a data frame, but they do not need to be.

1. A response vector. If your response is Poisson, then the entries in the response vector must be natural numbers. If your response is Bernoulli, then the entries in the response vector must be 0 and 1. If your response is binomial, then you will have two vectors: a vector of successes and a vector of failures. (For this version of `glmm`, these are the only response types possible. If you need to fit a model with a different response, contact me.)
2. At least one vector that will be used for defining the random effects' design matrix. For this version of `glmm`, the vector(s) should be class `factor`.
3. Vector(s) that will be used for defining the fixed effects' design matrix. The vector(s) can be of class `factor` or `numeric`.

The first two types of vectors described in the list are required. The last type is optional. That is, the minimum requirement to fit a `glmm` model is the response vector and one vector for defining the random effects' design matrix.

We use the `salamander` dataset as an example in this vignette. For your convenience, it is already included in the `glmm` package. The data arose from an experiment conducted at the University of Chicago in 1986 and were first presented by McCullagh and Nelder (1989, section 14.5). Scientists paired female and male salamanders of two types (Rough Butt and White Side) and collected data on whether or not they mated.

The variable `Mate` tells us whether the pair of salamanders mated. The value is 1 if they successfully mated and 0 if they did not. The variable `Cross` describes the type of female and

male salamander. For example, `Cross = W/R` indicates a White Side female was crossed with a Rough Butt male. The variable `Female` contains the identification number of the female salamander, and the variable `Male` contains the identification number of the male salamander.

The first R command shown below gives us access to the `glmm` package and all of its commands. The second line of code gives us access to the `salamander` data frame. The next three commands help us begin to understand the data. We have four variables: `Mate`, `Cross`, `Female`, and `Male`. The summary shows us `Mate` is numeric, `Cross` is a factor with four levels, `Female` is a factor, and `Male` is a factor.

```
library(glmm)
data(salamander)
names(salamander)

## [1] "Mate" "Cross" "Female" "Male"

head(salamander)

##   Mate Cross Female Male
## 1    1  R/R     10   10
## 2    1  R/R     11   14
## 3    1  R/R     12   11
## 4    1  R/R     13   13
## 5    1  R/R     14   12
## 6    1  R/W     15   28

summary(salamander)

##      Mate      Cross      Female      Male
## Min.   :0.000  R/R:90   10      : 6   10      : 6
## 1st Qu.:0.000  R/W:90   11      : 6   11      : 6
## Median :1.000  W/R:90   12      : 6   12      : 6
## Mean   :0.525  W/W:90   13      : 6   13      : 6
## 3rd Qu.:1.000          14      : 6   14      : 6
## Max.   :1.000          15      : 6   15      : 6
##                                     (Other):324  (Other):324
```

3 Fitting the Model

Following Model A from Karim and Zeger (1992), we set `Mate` as the response, `Cross` as the fixed effect variable, and `Female` and `Male` as the random effect variables. That is, we would like to fit a generalized linear mixed model with a logit link (because the response is Bernoulli). We will have four fixed effect parameters ($\beta_{R/R}, \beta_{R/W}, \beta_{W/R}, \beta_{W/W}$). There is likely to be variability among the females and variability among the males. That is, some females will be more likely to mate than other females, and we would like the model to reflect the tendencies of the individual salamanders. We incorporate this into the model by including a random effect for each female salamander and a random effect for each male salamander.

We believe the female salamanders' random effects are i.i.d. draws from $N(0, \nu_F)$, where ν_F is an unknown parameter to be estimated. Similarly, we believe the male salamanders' random effects are i.i.d. draws from $N(0, \nu_M)$, where ν_M is an unknown parameter to be estimated. Finally, we believe the female and male random effects are independent of one another.

In the following code, we fit the model using the `glmm` command and save the model under the name `sal`. Because `Mate` is our response, it is on the left of the \sim . We want to have a fixed effect for each of the four levels of `Cross`, so we type `Mate ~ 0 + Cross`. Because `Cross` is a factor, typing `Mate ~ Cross` would fit an equivalent model.

Next, the `random` list creates the design matrices for the random effects. Since we want two random effects for each cross (one from the female salamander and one from the male salamander), we type `list(~ 0 + Female, ~ 0 + Male)`. We include the 0 because we want our random effects to be centered at 0. Almost always, you will want your random effects to have mean 0.

Following the `random` list, the argument `varcomps.names` allows us to name the list of variance components. In the `random` list, we have placed the females first. Therefore, the order of the variance components names are first "F" and then "M."

Next, we specify the name of our data set. This is an optional argument. If the data set is not specified, `glmm` looks to the parent environment for the variables you have referenced.

After the name of the data set, we need to specify the type of the response. In the salamander mating example, the response is binary: the salamanders either mated or they did not. Therefore, the family is `bernoulli.glmm`. If your response is a count, then the family is `poisson.glmm`.

Next, we specify our Monte Carlo sample size `m`. The general rule is the larger the Monte Carlo sample size, the more accurate the Monte Carlo likelihood approximation (MCLA) will be, and the more accurate the resulting Monte Carlo maximum likelihood estimates (MCMLEs) will be. Ideally, you want the largest `m` that time allows. For this vignette, we have chosen a Monte Carlo sample size that allows for quick computation. If you are interested in accuracy in the resulting estimates for the salamander model, we suggest a larger Monte Carlo sample size.

Finally, we specify the cluster we would like to use for the value, gradient vector and Hessian matrix calculations, done in parallel, using `cluster` and the weighting scheme we would like to use for relevance weighting each observation, using `weights`. To read more about the `cluster` argument, please read "Developments in the R package `glmm`"

We put this all together in the following commands. Note that we set the seed so that we can have reproducible results. In other words, if you set your seed to the same number and type the exact command listed below, your results should be identical to those listed here.

```
set.seed(1234)
clust <- makeCluster(2)
sal <- glmm(Mate ~ 0 + Cross, random = list(~ 0 + Female,
```

```

~ 0 + Male), varcomps.names = c("F", "M"), data = salamander,
family.glmm = bernoulli.glmm, m = 10^4, debug = TRUE, cluster = clust)
stopCluster(clust)

```

4 Reading the Model Summary

The `summary` command displays

- the function call (to remind you of the model you fit).
- the link function.
- the fixed effect estimates, their standard errors (calculated using observed Fisher information), their `z value` test statistics (testing whether the coefficients are significantly different from zero), the test's p-values, and the R-standard significance stars (optional).
- the variance component estimates, their standard errors (calculated using observed Fisher information), their `z value` test statistics (testing whether the coefficients are significantly different from zero), the test's p-values, and the R-standard significance stars (optional).

Note that the p-value for the fixed effects is calculated using a two-sided alternative hypothesis ($H_A : \beta \neq 0$) while the p-value for the variance components is calculated using a one-sided alternative hypothesis ($H_A : \nu > 0$) because variance components must be nonnegative.

To view the model summary, we use the `summary` command.

```

summary(sal)

##
## Call:
## glmm(fixed = Mate ~ 0 + Cross, random = list(~0 + Female, ~0 +
##   Male), varcomps.names = c("F", "M"), data = salamander, family.glmm = bernoulli.glmm,
##   m = 10^4, debug = TRUE, cluster = clust)
##
##
## Link is: "logit (log odds)"
##
## Fixed Effects:
##           Estimate Std. Error z value Pr(>|z|)
## CrossR/R      1.230      0.300   4.045 5.24e-05 ***
## CrossR/W       0.320      0.267   1.198  0.23077
## CrossW/R     -2.000      0.330  -6.042 1.52e-09 ***
## CrossW/W       0.920      0.300   3.084  0.00204 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
##

```

```
## Variance Components for Random Effects (P-values are one-tailed):
## Estimate Std. Error z value Pr(>|z|)/2
## F      1.46      0.31  4.695  1.33e-06 ***
## M      1.64      0.33  4.918  4.36e-07 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Looking at our output, we can see that the type of cross significantly affects the salamanders' odds of mating. Additionally, both the variance components are significantly different from zero and should be retained in the model.

The summary provides the estimates needed to write our model. First, we establish a little notation. Let π_i represent the probability of successful mating for salamander pair i . Let $I()$ be an indicator function, so that $I(\text{Cross}=\text{R/R})$ is 1 when the variable **Cross** = **R/R** and 0 otherwise. Let u_i^F represent the random effect from the female salamander in the i th pair. Let u_i^M represent the random effect from the male salamander in the i th pair. Since the response is Bernoulli, the canonical link is the log odds of successful mating. Using this notation, we write the model as follows.

$$\begin{aligned} \log\left(\frac{\pi_i}{1-\pi_i}\right) &= 1.2328 * I(\text{Cross}=\text{R/R}) + 0.3199 * I(\text{Cross}=\text{R/W}) \\ &\quad + -1.9986 * I(\text{Cross}=\text{W/R}) + 0.9159 * I(\text{Cross}=\text{W/W}) \\ &\quad + u_i^F + u_i^M \\ u_i^F &\overset{i.i.d.}{\sim} N(0, 1.459) \\ u_i^M &\overset{i.i.d.}{\sim} N(0, 1.644) \end{aligned}$$

Recall that **m** in the above model was chosen for convenience to save time. The resulting parameter estimates have a little too much variability. If we increase **m**, the Monte Carlo standard error decreases.

5 Isolating the Parameter Estimates

If we wish to extract the estimates for the fixed effect coefficients or the variance components, we use the commands `coef` and `varcomps`, respectively. These commands isolate the estimates that are shown in the summary (as displayed in section 4).

To extract the fixed effect coefficients, the only argument needed is the model. The commands `coef` and `coefficients` are interchangeable. We can type either of the following:

```
coef(sal)
## CrossR/R CrossR/W CrossW/R CrossW/W
## 1.2327950 0.3199320 -1.9986114 0.9159357
```

```

coefficients(sal)

##   CrossR/R   CrossR/W   CrossW/R   CrossW/W
##  1.2327950   0.3199320  -1.9986114   0.9159357

```

To extract the variance components, the only argument needed is the model.

```

varcomps(sal)

##           F           M
## 1.458761  1.643627

```

To further isolate variance components or fixed effects, use indexing. The following demonstrates how to extract the last two fixed effects and the first variance component.

```

coef(sal)[c(3,4)]

##   CrossW/R   CrossW/W
## -1.9986114   0.9159357

varcomps(sal)[1]

##           F
## 1.458761

```

6 Calculating Confidence Intervals

We can calculate confidence intervals for parameters using the `confint` command. (Note that prediction is not yet possible in this version of the package). If we wish to calculate 95% confidence intervals for all of our parameters, the only argument is the model name.

```

confint(sal)

##           0.025           0.975
## CrossR/R  0.6354145  1.8301756
## CrossR/W -0.2033230  0.8431871
## CrossW/R -2.6469583 -1.3502645
## CrossW/W  0.3338634  1.4980081
## F         0.8498360  2.0676851
## M         0.9886561  2.2985979

```

The output is a matrix. Each row represents one parameter. The first column is the lower bound of the confidence interval, and the second column is the upper bound of the confidence interval.

If we wish to change the level of confidence from the default of 95%, we use the argument `level` and specify a number between 0 and 1. For example, to 90% confidence intervals and 99% confidence intervals, we type the following:

```

confint(sal,level=.9)

##           0.05      0.95
## CrossR/R  0.7314575  1.7341326
## CrossR/W -0.1191974  0.7590615
## CrossW/R -2.5427213 -1.4545015
## CrossW/W  0.4274452  1.4044263
## F         0.9477349  1.9697861
## M         1.0939581  2.1932959

```

```

confint(sal,level=.99)

##           0.005      0.995
## CrossR/R  0.4477039  2.017886
## CrossR/W -0.3677417  1.007606
## CrossW/R -2.8506837 -1.146539
## CrossW/W  0.1509630  1.680908
## F         0.6584980  2.259023
## M         0.7828493  2.504405

```

We can calculate 90% confidence intervals for the first and third fixed effects through indexing or by listing the names of the fixed effects:

```

confint(sal,level=.9,c(1,3))

##           0.05      0.95
## CrossR/R  0.7314575  1.734133
## CrossW/R -2.5427213 -1.454502

confint(sal,level=.9,c("CrossR/R","CrossW/R"))

##           0.05      0.95
## CrossR/R  0.7314575  1.734133
## CrossW/R -2.5427213 -1.454502

```

To calculate a 93 percent confidence interval for the variance component for the female salamanders, we can again either use indexing or list the name of the variable. Note that there are four fixed effects so ν_F is the fifth parameter in this model. (Similarly, ν_M is the sixth parameter in this model).

```

confint(sal,level=.93,c(5))

##           0.035      0.965
## F  0.8958334  2.021688

confint(sal,level=.93,c("F"))

##           0.035      0.965
## F  0.8958334  2.021688

```

Note that all confidence intervals are calculated using the observed Fisher information from the Monte Carlo likelihood approximation.

7 Monte Carlo Standard Error

A common question is “How big should the Monte Carlo sample size (m) be?” A larger m leads to higher accuracy in your likelihood-based inference, but the trade-off is a larger m takes more computing time. To assess whether your chosen Monte Carlo sample size m is large enough, you can use the Monte Carlo standard error.

The point estimates produced by `glm` (and by Monte Carlo likelihood approximation in general) have two sources of variability. First, there is variability from sample to sample. That is, if we conducted the experiment on another set of 120 salamanders, our point estimates would differ slightly. We measure this sample to sample variability with standard error. Second, there is the variability between calls of `glm` because different random numbers are used to calculate the Monte Carlo likelihood approximation. We measure this variability with Monte Carlo standard error.

If the Monte Carlo standard error is large (relative to the standard error), then you should use a large Monte Carlo sample size m . How much larger should m be? The Monte Carlo standard error decreases at a square root rate; if you want a Monte Carlo standard error that is 10% your current Monte Carlo standard error, then you should multiple your Monte Carlo sample size by 100.

You can find the Monte Carlo standard errors with the `mcse` command.

```
mcse(sal)
##      CrossR/R      CrossR/W      CrossW/R      CrossW/W      F      M
## 0.058272155 0.008764209 0.019981486 0.055950656 0.053502304 0.050063421
```

You can compare the Monte Carlo standard errors (shown above) to the standard errors of the point estimates (shown below).

```
se(sal)
##      CrossR/R      CrossR/W      CrossW/R      CrossW/W      F      M
## 0.3047916 0.2669718 0.3307953 0.2969811 0.3106815 0.3341750
```

8 Estimating the Variance-Covariance Matrix

The variance-covariance matrix for the parameter estimates can be found using the `vcov` function. The only input is the model name.

```
(myvcov <- vcov(sal)[1:4, 1:4])
```

```
##           CrossR/R      CrossR/W      CrossW/R      CrossW/W
## CrossR/R  0.092897916  0.002448403  0.0035382549 -0.0076306989
## CrossR/W  0.002448403  0.071273920  0.0014922529  0.0044663567
## CrossW/R  0.003538255  0.001492253  0.1094255383  0.0009020289
## CrossW/W -0.007630699  0.004466357  0.0009020289  0.0881978020
```

The variance-covariance matrix can be useful for some hypothesis testing. For example, suppose we want to test the hypotheses:

$$H_0 : \beta_{RR} - \beta_{WW} = 0$$
$$H_0 : \beta_{RR} - \beta_{WW} \neq 0.$$

The Wald test statistic is

$$\frac{\hat{\beta}_{RR} - \hat{\beta}_{WW} - 0}{\sqrt{\text{Var}(\hat{\beta}_{RR} - \hat{\beta}_{WW})}} \sim N(0, 1).$$

To calculate

$$\text{Var}(\hat{\beta}_{RR} - \hat{\beta}_{WW}) = \text{Var}(\hat{\beta}_{RR}) + \text{Var}(\hat{\beta}_{WW}) - 2 \text{Cov}(\hat{\beta}_{RR}, \hat{\beta}_{WW})$$

we use the variances and covariances from the variance-covariance matrix:

```
myvar <- myvcov[1,1] + myvcov[4,4] - 2* myvcov[1,4]
SE <- sqrt(myvar)
```

Then the test statistic and its associated p-value can be calculated:

```
test.stat <- (coef(sal)[1] - coef(sal)[4]) / SE
as.numeric(2 * pnorm(test.stat))

## [1] 1.525429
```

Therefore, we do not have evidence to reject $H_0 : \beta_{RR} = \beta_{WW}$. The probability of two White Side salamanders mating is not significantly different from the probability of two Rough Butt salamanders mating. This makes sense, considering how close β_{RR} and β_{WW} are.

Similarly, we could do a Wald-style hypothesis test to find the two variance components ν_F and ν_M are not significantly different.

9 Accessing Additional Output

The model produced by `glmm` has information that is not displayed by the `summary` command. The `names` command helps us see what we can access.

```
names(sal)

## [1] "beta"           "nu"             "loglike.value"  "loglike.gradient"
## [5] "loglike.hessian" "trust.converged" "mod.mcml"       "fixedcall"
## [9] "randcall"       "x"              "y"              "z"
## [13] "weights"        "family.glm"     "call"           "varcomps.names"
## [17] "varcomps.equal" "umat"           "pvec"           "beta.pql"
## [21] "nu.pql"         "u.pql"          "zeta"           "cluster"
## [25] "cores"          "debug"
```

The first two items are `beta` and `nu`. These are the MCMLEs for the fixed effects and variance components.

The third item is `likelihood.value`, the value of the MCLA evaluated at the MCMLEs. The fourth item is `likelihood.gradient`, the gradient vector of the MCLA evaluated at the MCMLEs. The fifth item is `likelihood.hessian`, the Hessian matrix of the MCLA evaluated at the MCMLEs.

Next is `trust.converged`, which tells us whether the `trust` function in the `trust` package was able to converge to the optimizer of the MCLA.

Items 7 through 16 relate to the original function call. `mod.mcml` contains the model matrix for the fixed effects, a list of model matrices for the random effects, and the response vector. These are also displayed in `x`, `z`, and `y`, respectively. Then, the call (the original formula representations of the fixed and random effects) are contained in `fixedcall`, `randcall`, and `call`.

The last argument is `debug`. If the model was fit with the default `debug = FALSE`, then this argument is just `FALSE`. If the model was fit with `debug = TRUE`, then `debug` contains a list of output for advanced users and programmers.

10 Optional Model-Fitting Arguments

Additional arguments may be added for more control over the model fit. If you're an introductory user, go ahead and ignore this section.

10.1 Setting Variance Components Equal

By default, `glmm` assumes each variance component should be distinct. Suppose we want to set $\nu_F = \nu_M$. Then we would add the argument `varcomps.equal` to indicate the equality. Since the list of random effects has two entries and we want those entries to share a variance component, we would set `varcomps.equal = c(1,1)`. In this scenario, we would only have one variance component, so we only need one entry in `varcomps.names`. Thus, the new command to fit this updated model with one variance component could be the following:

```
sal <- glmm(Mate ~ 0 + Cross, random = list(~ 0 + Female,
~ 0 + Male), varcomps.equal = c(1, 1), varcomps.names =
c("Only Varcomp"), data = salamander, family.glm =
```

```
bernoulli.glmm, m = 10^4, debug = TRUE, cluster = clust)
```

As another example, suppose the list `random` has three entries, indicating three variance components ν_1, ν_2, ν_3 . To set $\nu_1 = \nu_3$, we write `varcomps.equal = c(1,2,1)`. Thus, the shared variance component would be listed first in any output, and ν_2 would be listed second. Note that the entries in the `varcomps.equal` vector must start at 1, then continue through the integers. The order of the names of the variance components listed in `varcomps.names` must correspond to the integers in `varcomps.equal`. In this problem, the names could be `varcomps.names = c("shared", "two")`.

10.2 Altering the Importance Sampling Distribution

The following default arguments can be adapted to alter the importance sampling distribution: `doPQL`, `p1`, `p2`, `p3`, and `zeta`.

By default, penalized quasi-likelihood estimates are used to form the importance sampling distribution for the generated random effects. To skip PQL, add the argument `doPQL=FALSE`. If PQL is skipped, then the importance sampling distribution uses arbitrary estimates of 0 for the random effects, 0 for the fixed effects, and 1 for the variance components. Sometimes the examples in the `glmm` documentation skip the PQL step so that the package can load more quickly. Most of the time, the model will fit more accurately and efficiently if PQL estimates are used in the importance sampling distribution.

The importance sampling distribution is a mixture of three distributions. By default, the mixture is evenly weighted, with each component's contribution set at 1/3. If you wish to change the mixture, you can alter `p1`, `p2`, and `p3` from the default of `p1 = 1/3`, `p2 = 1/3`, and `p3 = 1/3`. The only restrictions are that the three probabilities must sum to 1 and `p1` must be positive.

The first component of the importance sampling distribution is a scaled multivariate t-distribution with `zeta` degrees of freedom. Therefore, another way to alter the importance sampling distribution is by changing `zeta` from its default of 5.

10.3 Adjusting Optimization Arguments

It may be useful to adjust the `trust` arguments `rmax` and `iterlim`. The argument `rmax` is the maximum allowed trust region radius. By `glmm` default, this is set to the arbitrary, somewhat large number of 1000. If this is set to a small number, then the optimization will move more slowly.

The argument `iterlim` must be a positive integer that limits the length of the optimization. If `iterlim` is too small, then the `trust` optimization will end before the MCMLA has been maximized.

If `iterlim` is reached, then `trust` has not converged to the MCMLE. When the `summary` command is called, a warning will be printed telling the user that the parameter values are

not MCMLEs, but `glmm` can be rerun starting at these outputted parameter values. To do this, use the `par.init` argument in section 10.4.

10.4 Starting at a Specified Parameter Value

Rather than using the PQL estimates, you can provide parameter values to `glmm` using the argument `par.init`. The `glmm` argument `par.init` is a vector that specifies the user-supplied values of the fixed effects and variance components. The parameters must be inputted in the order that `summary` outputs them, with fixed effects followed by variance components.

If `par.init` is provided, then PQL estimates will not be computed. The `par.init` estimates will be used instead to form the importance sampling distribution. Then, `trust` will use `par.init` as the starting point for the optimization. This argument may be useful for very hard problems that require iteration.

References

- Gelfand, A. and Carlin, B. (1993). Maximum-likelihood estimation for constrained- or missing-data models. *Canadian Journal of Statistics*, 21:303–311.
- Geyer, C. (1990). *Likelihood and Exponential Families*. PhD thesis, University of Washington.
- Geyer, C. J. (1994). On the convergence of Monte Carlo maximum likelihood calculations. *Journal of the Royal Statistical Society, Series B*, 61:261–274.
- Geyer, C. J. and Thompson, E. (1992). Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society, Series B*, 54:657–699.
- Karim, M. and Zeger, S. (1992). generalized linear models with random effects; salamander mating revisited. *Biometrics*, 48:631–644.
- McCullagh, P. and Nelder, J. (1989). *Generalized Linear Models*. Chapman and Hall/CRC.
- Sung, Y. J. and Geyer, C. J. (2007). Monte Carlo likelihood inference for missing data models. *Annals of Statistics*, 35:990–1011.

Model-Fitting in Parallel with the R package `glmm`

Sydney Benson

August 17, 2018

Contents

1	Introduction	2
2	Additional Model-Fitting Arguments	2
3	Additional Output	3

1 Introduction

The R package `glmm` approximates the likelihood function for generalized linear mixed models (GLMMs) with a canonical link. `glmm` calculates and maximizes the Monte Carlo likelihood approximation (MCLA) to find Monte Carlo maximum likelihood estimates (MCMLEs) for the fixed effects and variance components. The value, gradient vector, and Hessian matrix of the MCLA are calculated at the MCMLEs. The Hessian of the MCLA is used to calculate the standard errors for the MCMLEs.

In version 1.2.4, the R package `glmm` has been revised to calculate the value, gradient vector and Hessian matrix in parallel. This addition has added an optional argument to the `glmm` command, additional output, and has decreased the time it takes to fit the model.

2 Additional Model-Fitting Arguments

In the following code, we fit the model using the `glmm` command and save the model under the name `sal`. The final argument for the `glmm` command is `cores`. The `cores` argument is optional and, by default, R will calculate the optimal number of cores to use for the parallel calculations if `cores` is not specified. The optimal number of cores is one less than the number of cores available in the device. However, if you choose to specify a different number of cores, you may do that using the `cores` argument. In this example, we will use two cores.

Note: you may only specify the number of cores you want to use if it is less than or equal to the optimal number of cores that R defines. If you specify a number greater than the optimal number defined by R, the number of cores used will default to R's optimal number.

Using multiple cores is useful for reducing the time the `glmm` command takes to run, thus allowing for an increased `m` without additional computational expense. We can see the time reduction made by running the calculations in parallel using the `proc.time` command. The times shown here are from fitting a model on a MacBook Air that cost 1000 USD in 2015.

```
> library(glmm)
> data(salamander)
```

```

> set.seed(1234)
> start <- proc.time()
> sal <- glmm(Mate ~ 0 + Cross, random = list(~ 0 + Female,
+ ~ 0 + Male), varcomps.names = c("F", "M"), data = salamander,
+ family.glmm = bernoulli.glmm, m = 10^4, debug = TRUE, cores = 1)
> proc.time() - start

      user  system elapsed
101.479    7.791  341.730

> set.seed(1234)
> start <- proc.time()
> sal <- glmm(Mate ~ 0 + Cross, random = list(~ 0 + Female,
+ ~ 0 + Male), varcomps.names = c("F", "M"), data = salamander,
+ family.glmm = bernoulli.glmm, m = 10^4, debug = TRUE, cores = 2)
> proc.time() - start

      user  system elapsed
 90.983    5.009  237.323

```

To read about the other arguments in the `glmm` command, please read "An Introduction to Model-Fitting with the R package `glmm`".

3 Additional Output

Using the `names` command allows us to see the additional output that we can access, beyond the model summary. With the addition of parallel calculations, we also have the addition of the `cores` output. `cores` will tell you how many cores were used during the value, gradient vector and Hessian matrix calculations.

```

> names(sal)

 [1] "beta"           "nu"             "likelihood.value"
 [4] "likelihood.gradient" "likelihood.hessian" "trust.converged"
 [7] "mod.mcml"       "fixedcall"      "randcall"
[10] "x"              "y"              "z"
[13] "family.glmm"    "call"           "varcomps.names"
[16] "varcomps.equal" "umat"           "pvec"
[19] "beta.pql"       "nu.pql"         "u.pql"
[22] "zeta"           "cores"          "debug"

```

```
> sal$scores
```

```
[1] 2
```

To read about the model summary or the other additional output provided when using the `glm` command, please read "An Introduction to Model-Fitting with the R package `glm`"