

Getting started with the `glmmADMB` package

Ben Bolker, Hans Skaug, Arni Magnusson, Anders Nielsen

May 7, 2013

1 Introduction/quick start

`glmmADMB` is a package, built on the open source AD Model Builder nonlinear fitting engine, for fitting generalized linear mixed models and extensions.

- response distributions: Poisson, binomial, negative binomial (NB1 and NB2 parameterizations), Gamma, Beta, truncated Poisson and negative binomial; Gaussian; logistic
- link functions: log, logit, probit, complementary log-log ("`cloglog`"), inverse, identity
- zero-inflation (models with a constant zero-inflation value only); hurdle models via truncated Poisson/NB
- single or multiple (nested or crossed) random effects
- offsets
- post-fit MCMC chain for characterizing uncertainty

As of version 0.6.5, the package has been greatly revised to allow a wider range of response and link functions and to allow models with multiple random effects. For now, the resulting package is slower than the old (single-random-effect version), but we hope to increase its speed in the future.

In order to use `glmmADMB` effectively you should already be reasonably familiar with generalized linear mixed models (GLMMs), which in turn requires familiarity with (i) generalized linear models (e.g. the special cases of logistic, binomial, and Poisson regression) and (ii) ‘modern’ mixed models (those working via maximization of the marginal likelihood rather than by manipulating sums of squares).

In order to fit a model in `glmmADMB` you need to:

- specify a model for the fixed effects, in the standard R (Wilkinson-Rogers) formula notation (see `?formula` or Section 11.1 of the Introduction to R. Formulae can also include *offsets*).

- specify a model for the random effects, in the notation that is common to the `nlme` and `lme4` packages. Random effects are specified as `e|g`, where `e` is an effect and `g` is a grouping factor (which must be a factor variable, or a nesting of/interaction among factor variables). For example, the formula would be `1|block` for a random-intercept model or `time|block` for a model with random variation in slopes through time across groups specified by `block`. A model of nested random effects (block within site) would be `1|site/block`; a model of crossed random effects (block and year) would be `(1|block)+(1|year)`.

Random effects can be specified either in a separate `random` argument (as in `nlme`) or as part of the model formula (as in `lme4`).

- choose the error distribution by specifying the family (as a string: e.g. `"poisson"` or `"binomial"`)
- specify a link function (as a string: e.g. `"logit"` or `"log"`).
- optionally specify that zero-inflation is present `zeroInflation=TRUE`. In the current version, zero-inflation can only be specified as a single constant term across the entire model — i.e. it cannot vary across groups or with covariates.

This document was generated using R Under development (unstable) (2013-04-25 r62670) and package versions:

##	bbmle	coda	ggplot2	glmmADMB	lme4
##	1.0.5.2	0.16-1	0.9.3.1	0.7.5	0.99999911-2
##	MASS	scapeMCMC			
##	7.3-26	1.1-3			

```

citation("glmmADMB")

##
## To cite package 'glmmADMB' in publications use:
##
## Fournier DA, Skaug HJ, Ancheta J, Ianelli J, Magnusson A, Maunder
## M, Nielsen A and Sibert J (2012). "AD Model Builder: using
## automatic differentiation for statistical inference of highly
## parameterized complex nonlinear models." _Optim. Methods Softw._,
## *27*, pp. 233-249.
##
## Skaug H, Fournier D, Nielsen A, Magnusson A and Bolker B ($Date:
## 2013-02-20 16:11:03 -0500 (Wed, 20 Feb 2013) $). _Generalized
## Linear Mixed Models using AD Model Builder_. R package version
## 0.7.5.

```

2 Owls data

These data, taken from Zuur et al. (2009) and ultimately from Roulin and Bersier (2007), quantify the number of negotiations among owlets (owl chicks) in different nests *prior* to the arrival of a provisioning parent as a function of food treatment (deprived or satiated), the sex of the parent, and arrival time. The total number of calls from the nest is recorded, along with the total brood size, which is used as an offset to allow the use of a Poisson response.

Since the same nests are measured repeatedly, the nest is used as a random effect. The model can be expressed as a zero-inflated generalized linear mixed model (ZIGLMM).

First we draw some pictures (Figures ??, ??).

Load the `glmmADMB` package to get access to the `Owls` data set; load the `ggplot2` graphics package.

```
library("glmmADMB")
library("ggplot2")
```

Various small manipulations of the data set: (1) reorder nests by mean negotiations per chick, for plotting purposes; (2) add log brood size variable (for offset); (3) rename response variable.

```
Owls <- transform(Owls, Nest = reorder(Nest, NegPerChick), NCalls = SiblingNegotiation)
```

(If you were really using this data set you should start with `summary(Owls)` to see what columns are there and what their characteristics are.)

```
## theme_text is deprecated. Use 'element_text' instead. (Deprecated;  
last used in version 0.9.1)  
## theme_text is deprecated. Use 'element_text' instead. (Deprecated;  
last used in version 0.9.1)
```

```
## geom_smooth: method="auto" and size of largest group is <1000,  
so using loess. Use 'method = x' to change the smoothing method.
```

Now fit some models:

The basic `glmmadmb` fit — a zero-inflated Poisson model.

```
fit_zipoiss <- glmmadmb(NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +  
  offset(log(BroodSize)) + (1 | Nest), data = Owls, zeroInflation = TRUE,  
  family = "poisson")
```




Figure 1: Basic view of owl data (arrival time not shown). (FIXME: clip at zero?)




Figure 2: Basic view of owl data, #2 (nest identity not shown)

```
summary(fit_zipoiss)

##
## Call:
## glmmadmb(formula = NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +
##   offset(log(BroodSize)) + (1 | Nest), data = Owls, family = "poisson",
##   zeroInflation = TRUE)
##
## AIC: 3987
##
## Coefficients:
##                                     Estimate Std. Error z value Pr(>|z|)
## (Intercept)                        2.8564     0.3872   7.38 1.6e-13
## FoodTreatmentSatiated              -0.3314     0.0635  -5.22 1.8e-07
## ArrivalTime                       -0.0807     0.0156  -5.18 2.3e-07
## SexParentMale                      0.2885     0.3575   0.81  0.42
## FoodTreatmentSatiated:SexParentMale  0.0740     0.0761   0.97  0.33
## ArrivalTime:SexParentMale          -0.0150     0.0143  -1.05  0.29
##
## (Intercept)                        ***
## FoodTreatmentSatiated              ***
## ArrivalTime                       ***
## SexParentMale
## FoodTreatmentSatiated:SexParentMale
## ArrivalTime:SexParentMale
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Number of observations: total=599, Nest=27
## Random effect variance(s):
## Group=Nest
##               Variance StdDev
## (Intercept)   0.1399  0.374
##
## Zero-inflation: 0.2583 (std. err.: 0.01811 )
##
## Log-likelihood: -1985
```

The `coefplot2` package knows about `glmmadmb` fits:

```
library("coefplot2")
coefplot2(fit_zipoiss)
```

We can also try a standard zero-inflated negative binomial model; the default is the “NB2” parameterization (variance = $\mu(1 + \mu/k)$).

```
fit_zinbinom <- glmmadmb(NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +
  offset(log(BroodSize)) + (1 | Nest), data = Owls, zeroInflation = TRUE,
  family = "nbinom")
```

Alternatively, use an “NB1” fit (variance = $\phi\mu$).

```
fit_zinbinom1 <- glmmadmb(NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +
  offset(log(BroodSize)) + (1 | Nest), data = Owls, zeroInflation = TRUE,
  family = "nbinom1")
```

Relax the assumption that total number of calls is strictly proportional to brood size (i.e. using $\log(\text{brood size})$ as an offset):

```
fit_zinbinom1_bs <- glmmadmb(NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +
  BroodSize + (1 | Nest), data = Owls, zeroInflation = TRUE, family = "nbinom1")
```

Every change we have made so far improves the fit — changing distributions improves it enormously, while changing the role of brood size makes only a modest (-1 AIC unit) difference:

```
library("bbmle")
AICtab(fit_zipoiss, fit_zinbinom, fit_zinbinom1, fit_zinbinom1_bs)

##              dAIC  df
## fit_zinbinom1_bs    0.0 10
## fit_zinbinom1       1.2  9
## fit_zinbinom        68.7  9
## fit_zipoiss        637.0  8
```

Compare the parameter estimates:

```
vn <- c("food", "arrivaltime", "sex", "food:sex", "arrival:sex", "broodsize")
coefplot2(list(ZIP = fit_zipoiss, ZINB = fit_zinbinom, ZINB1 = fit_zinbinom1,
  ZINB1_brood = fit_zinbinom1_bs), varnames = vn, legend = TRUE)
```

2.1 Hurdle models

In contrast to zero-inflated models, hurdle models treat zero-count and non-zero outcomes as two completely separate categories, rather than treating the zero-count outcomes as a mixture of structural and sampling zeros.

As of version 0.6.7.1, `glmmADMB` includes truncated Poisson and negative binomial families and hence can fit hurdle models. The two parts of the model have to be fitted separately, however. First we fit a truncated distribution to the non-zero outcomes:

```
fit_hnbinom1 <- glmmadmb(NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +
  BroodSize + (1 | Nest), data = subset(Owls, NCalls > 0), family = "truncnbinom1")
```

Then we fit a model to the binary part of the data (zero vs. non-zero). In this case, I started by fitting a simple (intercept-only) model with intercept-level random effects only. This comes a bit closer to matching the previous (zero-inflation) models, which treated zero-inflation as a single constant level across the entire data set (in fact, leaving out the random effects and just using `glmmADMB(nz~1,data=Owls,family="binomial")`, or `glm(nz~1,data=Owls,family="binomial")`, would be an even closer match). I then fitted a more complex binary model — this is all a matter of judgment about how complex a model it's worth trying to fit to a given data set — but it does look as though the zero-inflation varies with arrival time and satiation.

```
Owls$nz <- as.numeric(Owls$NCalls > 0)
fit_count <- glmmadmb(nz ~ 1 + (1 | Nest), data = Owls, family = "binomial")
fit_ccount <- glmmadmb(nz ~ (FoodTreatment + ArrivalTime) * SexParent + (1 |
  Nest), data = Owls, family = "binomial")
AICtab(fit_count, fit_ccount)

##              dAIC df
## fit_ccount  0.0  7
## fit_count   84.1  2

summary(fit_ccount)

##
## Call:
## glmmadmb(formula = nz ~ (FoodTreatment + ArrivalTime) * SexParent +
##   (1 | Nest), data = Owls, family = "binomial")
##
## AIC: 580.2
##
## Coefficients:
##
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)         7.3109    2.1577   3.39  0.0007
## FoodTreatmentSatiated -1.8250    0.3479  -5.25 1.6e-07
## ArrivalTime         -0.2171    0.0846  -2.57  0.0102
## SexParentMale         2.7699    3.0248   0.92  0.3598
## FoodTreatmentSatiated:SexParentMale -0.3646    0.4740  -0.77  0.4418
## ArrivalTime:SexParentMale -0.0821    0.1175  -0.70  0.4849
##
## (Intercept)          ***
## FoodTreatmentSatiated ***
## ArrivalTime          *
## SexParentMale
```



```
## FoodTreatmentSatiated:SexParentMale
## ArrivalTime:SexParentMale
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Number of observations: total=599, Nest=27
## Random effect variance(s):
## Group=Nest
##           Variance StdDev
## (Intercept)  1.423  1.193
##
##
## Log-likelihood: -283.1
```

2.2 Testing and inference

(Sketchy: to be expanded.)

There are many challenging statistical issues surrounding tests of terms in GLMMs. Most often people use minor variations of existing approaches (Wald tests, likelihood ratio tests, etc.), either accounting for or sweeping under the rug some of the differences that should be accounted for when moving either from linear mixed models (LMMs) or generalized linear models (GLMs) to GLMMs. See <http://glmm.wikidot.com/faq>, or a good book on GLMMs (!!), for a discussion of these issues.

In the meantime, keeping the limitations in mind, you can

- use AIC to select models or generated weighted predictions (see example above).
- use `anova` to perform a likelihood ratio test:

```
anova(fit_zipoiss, fit_zinbinom)

## Analysis of Deviance Table
##
## Model 1: NCalls ~ (FoodTreatment + ArrivalTime) * SexParent + offset(log(BroodSize))
## Model 2: NCalls ~ (FoodTreatment + ArrivalTime) * SexParent + offset(log(BroodSize))
##   NoPar LogLik Df Deviance Pr(>Chi)
## 1      8 -1985
## 2      9 -1700  1      570  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

- use `Anova` from the `car` package to generate Wald tests:

```
car::Anova(fit_zinbinom)

## Analysis of Deviance Table (Type II tests)
##
## Response: NCalls
##
##           Df Chisq Pr(>Chisq)
## FoodTreatment      1  8.84    0.0030 **
## ArrivalTime        1  8.53    0.0035 **
## SexParent           1  0.14    0.7036
## FoodTreatment:SexParent  1  0.99    0.3186
## ArrivalTime:SexParent  1  0.11    0.7389
## Residuals          590
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

- coefplot2
- drop1??,
- simulate??,
- parametric bootstrapping ??

2.3 Integration with lme4

FIXME: commented out while lme4 bugs are being ironed out!

```
library("lme4")
gm1_lme4 <- glmer(cbind(incidence, size - incidence) ~ period + (1 | herd),
  data = cbpp, family = binomial)
gm1_glmmADMB <- glmmadmb(cbind(incidence, size - incidence) ~ period + (1 |
  herd), data = cbpp, family = "binomial")
## sessionInfo()
fixef(gm1_lme4)
fixef(gm1_glmmADMB) ## or coef()
unlist(ranef(gm1_lme4))
unlist(ranef(gm1_glmmADMB))
VarCorr(gm1_lme4)
detach("package:lme4") ## kluge!
VarCorr(gm1_glmmADMB)
```

2.3.1 Convert glmmADMB parameters to lme4

this section is UNDER CONSTRUCTION

```

library("lme4")

##
## Attaching package: 'lme4'
## The following object is masked from 'package:glmmADMB':
##
##      VarCorr

## seems to crash r-forge tests at the moment?  requires DEVELOPMENT
## version of lme4 ...
new_lme4 <- packageVersion("lme4") > "0.999375.42"
if (new_lme4) {
  lme4fun <- update(gm1_lme4, devFunOnly = TRUE)
  deviance(gm1_lme4)
  lme4fun(c(0, fixef(gm1_lme4))) ## variance set to zero
  v1 <- getME(gm1_lme4, "theta") ## log-Cholesky factor: equal in this case to standard c
  lme4fun(c(v1, fixef(gm1_lme4))) ## observed variance (*almost* identical)
  v2 <- sqrt(gm1_glmmADMB$S[[1]])
  lme4fun(c(v2, fixef(gm1_lme4))) ## ???
}

## FIXME: figure out why this crashes knitr in non-interactive mode!!
## (then turn evaluation back on)
glmmadmbfun <- function(pars, minval = exp(-9.5), verbose = FALSE) {
  if (pars[1] == 0) {
    warning(sprintf("variance parameter set to min val (%f)", minval))
    pars[1] <- minval
  }
  g0 <- glmmadmb(cbind(incidence, size - incidence) ~ period + (1 | herd),
    data = cbpp, family = "binomial", extra.args = c("-maxfn 1 -phase 6"),
    verbose = verbose, start = list(RE_sd = log(pars[1]), fixed = pars[-1]))
  -logLik(g0)
}
## glmmadmbfun(c(0,fixef(gm1_lme4))) ## fails
v2B <- glmmadmbfun(c(v2, fixef(gm1_lme4)))
v1B <- glmmadmbfun(c(v1, fixef(gm1_lme4)))
all.equal(v1B, v2B, -logLik(gm1_glmmADMB))

```

2.4 MCMC fitting

AD Model Builder can run a *post hoc* Markov chain to assess variability — that is, it uses the MLE as a starting point and the estimated sampling distribution (variance-covariance matrix) of the parameters as a candidate distribution, and

“jumps around” the parameter space in a consistent way (using Metropolis-Hastings updating with a multivariate normal candidate distribution) to generate a series of samples from a posterior distribution of the parameter distribution (by default assuming flat, usually improper, priors: please see the ADMB documentation, or Bolker (2008), for more details).

This is a useful way to find confidence intervals etc., but is often slow. In the example below, I ran a chain of 50,000 MCMC iterations — on examination, the default chain of 1000 iterations was much too short — which took several hours on a modern (2012) desktop.

```
OwlModel_nb1_bs_mcmc <- glmmadmb(NCalls ~ (FoodTreatment + ArrivalTime) * SexParent +
  BroodSize + (1 | Nest), data = Owls, zeroInflation = TRUE, family = "nbinom1",
  mcmc = TRUE, mcmc.opts = mcmcControl(mcmc = 50000))
```


Extract the MCMC chain from the fit; load the `coda` package for analyzing and plotting MCMC results.

```
library("coda")
m <- OwlModel_nb1_bs_mcmc$mcmc
```

Look at the trace plots. (Something a bit odd happens at the end of the chain, so we drop the last few values ...there may be a bug in the import-handling for MCMC for very long chains ...) For convenience and brevity, I will discard the random-effects columns, which start with `u.`:

```
tm <- window(m, 1, 320)
tm <- tm[, !grepl("^u\\.\\.", colnames(m))]
```

```
library("scapeMCMC")
plotTrace(tm)
```



The Geweke diagnostic gives Z scores for each variable for a comparison between (by default) the first 10% and last 50% of the chain

```
(gg <- geweke.diag(tm))

##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
##              pz              (Intercept)
##              0.9514             -0.2756
##      FoodTreatmentSatiated      ArrivalTime
##              1.2285              0.1011
##      SexParentMale      BroodSize
##             -0.4542              0.3225
## FoodTreatmentSatiated:SexParentMale      ArrivalTime:SexParentMale
##             -0.3902              0.4172
##              tmpL              alpha
##             0.8514             -0.9233

summary(2 * pnorm(abs(gg$z), lower.tail = FALSE))

##      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##      0.219  0.366   0.663   0.578   0.734   0.920
```

The most frequently used diagnostic, Gelman-Rubin (`gelman.diag`), requires multiple chains. The full set of diagnostic functions available in `coda` is:

```
## [1] autocorr.diag gelman.diag geweke.diag heidel.diag raftery.diag
```

`effectiveSize` gives the effective length of the chain for each variable, i.e. the number of samples corrected for autocorrelation:

```
range(effectiveSize(tm))

## [1] 233.1 417.1
```

`HPDinterval` gives the highest posterior density (credible interval):

```
try(detach("package:lme4")) ## kluge!!
HPDinterval(tm)

##               lower      upper
## pz              0.04298  0.14201
## (Intercept)      2.64801  5.94767
## FoodTreatmentSatiated -1.19929 -0.68662
## ArrivalTime      -0.18632 -0.06142
## SexParentMale     -1.71678  2.19173
## BroodSize         0.04601  0.32412
## FoodTreatmentSatiated:SexParentMale -0.01308  0.63548
## ArrivalTime:SexParentMale -0.09442  0.06758
## tmpL             0.22402  0.57588
## alpha            4.79370  6.94246
## attr(,"Probability")
## [1] 0.95
```

You might prefer inferences based on the quantiles instead:

```
t(apply(tm, 2, quantile, c(0.025, 0.975)))

##               2.5%      97.5%
## pz              0.04603  0.14935
## (Intercept)      2.57679  5.87872
## FoodTreatmentSatiated -1.19408 -0.66973
## ArrivalTime      -0.18398 -0.05809
## SexParentMale     -1.51902  2.52424
## BroodSize         0.05040  0.34938
## FoodTreatmentSatiated:SexParentMale -0.04280  0.61561
## ArrivalTime:SexParentMale -0.10927  0.05843
## tmpL             0.23783  0.60309
## alpha            4.83412  7.06562
```

You can also look at density plots or pairwise scatterplots (“splom” in `lattice` and `scapeMCMC`, for **Scatterplot matrices**) (not shown):

```
plotDens(tm)
plotSplom(tm, pch = ".")
```

The MCMC output in `glmmADMB` is not completely translated. It includes, in order:

pz zero-inflation parameter (raw)

fixed effect parameters Named in the same way as the results of `coef()` or `fixef()`.

tmpL variances (standard-deviation scale)

tmpL1 correlation/off-diagonal elements of variance-covariance matrices (“off-diagonal elements of the Cholesky factor of the correlation matrix”). (If you need to transform these to correlations, you will need to construct the relevant matrices with 1 on the diagonal and compute the cross-product, CC^T (see `tcrossprod`); if this makes no sense to you, contact the maintainers ...)

alpha overdispersion/scale parameter

u random effects (*unscaled*: these can be scaled using the estimated random-effects standard deviations from `VarCorr()`)

If you need to some insufficiently-processed component of the MCMC output and can’t figure out how, please contact the maintainers and encourage them to work on them some more (!)

3 Other information

The standard set of accessors is available:

coef extract (fixed-effect) coefficients

fixef a synonym for **coef**, for consistency with `nlme/lme4`

ranef extract random effect coefficients (“BLUPs” or “conditional modes”)

residuals extract (Pearson) residuals

fitted fitted values

predict predicted values (*based only on fixed effects, not on random effects*), possibly with standard errors (*based only on uncertainty of fixed effects*), possibly for new data

`logLik` extract log-likelihood

`AIC` extract AIC

`summary` print summary

`stdEr` extract standard errors of coefficients

`vcov` extract estimated variance-covariance matrix of coefficients

`VarCorr` extract variance-covariance matrices of random effects

`confint` extract confidence intervals of fixed-effect coefficients

In case this list is out of date, you can try `methods(class="glmmadmb")` to tell you what methods are currently available.

4 To do/road map

4.1 Vignette

- More examples
- Show how to specify starting values
- fix MCMC! Apply phi, std dev
- General troubleshooting (extra arguments, running outside R)
- basic intro to `R2admb`?
- (appendix?) document details of TPL file – robustness hacks, etc.

4.2 Code

- Speed improvement by identifying special cases?
- Spatial models?
- Additional flexibility:
 - Allow model specification for zero-inflation
 - Allow model specification for shape parameter
 - More complex variance models (see AS-REML/MCMCglmm for interface/syntax ideas)
- Improve `predict` method: allow prediction based on REs
- `simulate` method
- fix `step/stepAIC` (???)
- instructions for evaluating likelihood at a particular set of parameters/ translating parameters (and log-likelihoods) to scales/parameterizations used by other packages

References

- Bolker, B. M. (2008). *Ecological Models and Data in R*. Princeton, NJ: Princeton University Press.
- Roulin, A. and L. Bersier (2007). Nestling barn owls beg more intensely in the presence of their mother than in the presence of their father. *Animal Behaviour* 74, 1099–1106.
- Zuur, A. F., E. N. Ieno, N. J. Walker, A. A. Saveliev, and G. M. Smith (2009, March). *Mixed Effects Models and Extensions in Ecology with R* (1 ed.). Springer.