Linear Mixed Effects Regression

<2014.2.10>

AC's attempts at learning to do this in R.

Added package **lme4** to R. http://cran.r-project.org/web/packages/lme4/lme4.pdf

Also added **plyr** package, because it has a rename() function for renaming columns in a data frame: (http://www.cookbook-r.com/Manipulating_data/Renaming_columns_in_a_data_frame/)

However, I found it easier to do this:

```
colnames(x) <- c("name1","name2"...)
```

based on guidance from:

https://stat.ethz.ch/R-manual/R-devel/library/base/html/colnames.html

and from

http://cran.r-project.org/doc/manuals/r-devel/R-data.html

> Column names can be given explicitly via the col.names; explicit names override the header line (if present).

Added **arm** package: http://cran.r-project.org/web/packages/arm/index.html

> `arm: Data Analysis Using Regression and Multilevel/Hierarchical Models`
>
> R functions for processing lm, glm, svy.glm, merMod and polr outputs.

arm package provides the `coefplot` function for plotting the regression coefficient values and std. errors.

Need to specify that subjects is a nominal variable, using the `factor` command:

http://www.statmethods.net/input/datatypes.html

> Tell R that a variable is nominal by making it a factor. The factor stores the nominal values as a vector of integers in the range [1... k] (where k is the number of unique values in the nominal variable), and an internal vector of character strings (the original values) mapped to these integers.

Interpretation of glmer output

A nice description from:

https://stat.ethz.ch/pipermail/r-sig-mixed-models/2011q1/015590.html

> In this case the coefficient corresponds to one of the
> terms in the model and I would advocate performing a likelihood ratio
> test comparing the two models
>
> ```
> >fm <- glmer(SameSite~BreedSuc1+Sex+(1|Bird), family="binomial")
> >fm0 <- glmer(SameSite~Sex+(1|Bird), family="binomial")  # the null
> >hypothesis model
> >anova(fm0, fm)
> >
> ```
>
> Even though the function is called anova it will, in this case,
> perform a likelihood ratio test (LRT). It also prints the values of
> AIC and BIC if you prefer to compare models according to one of those
> criteria but I prefer using the likelihood ratio for nested models.
> However, before doing that comparison you should ask yourself whether
> you want to compare models that have the, apparently unnecessary term
> for Sex in them. The way I would approach the model building is first
> to reduce the model to
>
> ```
> >fm1 <- lmer(SameSite~BreedSuc1+(1|Bird), family="binomial")
> >
> ```
>
> You could then compare
>
> ```
> >anova(fm1, fm)
> >
> ```
>
> which I presume will give a large p-value for the LRT, so we prefer
> the simpler model, fm1. After that, I would compare
>
> ```
> >fm2 <- lmer(SameSite ~ 1 + (1|Bird), family="binomial")
> >
> ```

From:
https://www.wiki.anthonycate.org/ - **Visual Cognitive Neuroscience**

Permanent link:
https://www.wiki.anthonycate.org/doku.php?id=resources:lme&rev=1435159889

Last update: **2019/05/22 20:08**