

Posterior comparison of regression coefficients

March 8, 2014

The idea here is pretty simple: The credible intervals associated with each coefficient in the regressions identifies associations for which we have good evidence, i.e., those for which the credible interval doesn't overlap zero. But we can't compare the regression coefficients for *Protea* and *Pelargonium* simply by comparing which coefficients are detectably different from zero, except in the simple case when coefficients don't overlap zero and have opposite signes.

It is, however, easy to construct statistics to compare them using posterior comparisons [?]. We already have a posterior sample for each set of regression coefficients, so we simply select one coefficient from the posterior in the *Protea* model, one from the posterior in the *Pelargonium* model, take the difference, repeat a large number of times, calculate the credible interval of the difference, and determine whether it overlaps zero. If it does, we don't have evidence that the coefficients are different. If it doesn't, we do.

```
load("results-2014-03-06-09-01-02.Rsave")
protea <- fit$BUGSoutput$sims.list
load("results-2014-03-08-13-25-09.Rsave")
pellie <- fit$BUGSoutput$sims.list
rm(fit)
```

`protea` and `pellie` are now simple objects that contain the posterior sample for each coefficient in the model. We're only interested in comparing the `beta.*` coefficients. Since the two chains were run separately from one another and are of the same size we can construct our posterior sample by simply running through each sample fromt the posterior sequentially. First, we set up the posterior comparison function. That's really easy (so long as

we're using symmetric credible intervals) because all we have to do is determine the lower and upper bounds, take the difference of the input vectors and return the quantiles.

```
compare <- function(x, y, prob = 0.95) {
  lo <- (1 - prob)/2
  hi <- 1 - lo
  diff <- x - y
  quantile(diff, probs = c(lo, hi))
}
```

Now we simply iterate through each of the **beta.*** and report the results. To make the coefficients that are detectably different from one another a bit easier to spot, we add a '*' after the credible interval.

```
## get the indices for beta coefficients in the
## objects they'll be the same for protea and
## pellie
idx <- grep("beta", names(pellie))
for (i in idx) {
  for (j in 1:4) {
    cred <- compare(protea[[i]][, j], pellie[[i]][,
      j])
    cat(names(protea)[[i]], "[", j, "]", " (",
      cred[1], ",", cred[2], ")", sep = "")
    ## we know that the credible interval doesn't
    ## overlap zero if either the lower bound is
    ## positive or the upper bound is negative
    if ((cred[1] > 0) || (cred[2] < 0)) {
      cat("*", sep = "")
    }
    cat("\n")
  }
}

## beta.cdd[1] (-0.03624,0.3088)
## beta.cdd[2] (-0.785,-0.3749)*
## beta.cdd[3] (-0.04717,0.182)
```

```
## beta.cdd[4] (-0.0724,0.2932)
## beta.elev[1] (-0.4811,0.06748)
## beta.elev[2] (-0.4061,0.2245)
## beta.elev[3] (-0.1862,0.2201)
## beta.elev[4] (-0.1839,0.3939)
## beta.inso[1] (-0.3677,-0.07219)*
## beta.inso[2] (-0.07133,0.2912)
## beta.inso[3] (-0.1027,0.1064)
## beta.inso[4] (-0.02798,0.2842)
## beta.map[1] (-0.1653,0.2921)
## beta.map[2] (-0.4748,0.07021)
## beta.map[3] (0.1035,0.424)*
## beta.map[4] (-0.271,0.1902)
## beta.mat[1] (0.2681,0.7574)*
## beta.mat[2] (-0.5638,0.05674)
## beta.mat[3] (-0.182,0.1915)
## beta.mat[4] (-0.2947,0.2045)
## beta.ratio[1] (0.01782,0.4088)*
## beta.ratio[2] (-0.4621,0.02279)
## beta.ratio[3] (-0.3998,-0.1067)*
## beta.ratio[4] (-0.2784,0.1458)
```

It's probably helpful to visualize the differences by plotting the posteriors.

```
source("posterior-comparisons.R")

## Loading required package: ggplot2
```







