

Additional **multcomp** Examples

Torsten Hothorn

January 7, 2009

It is assumed that the reader is familiar with the theory and applications described in the `generalsiminf` vignette.

1 Simple Examples

Example: Simple Linear Model. Consider a simple univariate linear model regressing the distance to stop on speed for 50 cars:

```
R> lm.cars <- lm(dist ~ speed, data = cars)
R> summary(lm.cars)
```

Call:

```
lm(formula = dist ~ speed, data = cars)
```

Residuals:

Min	1Q	Median	3Q	Max
-29.07	-9.53	-2.27	9.21	43.20

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-17.579	6.758	-2.60	0.012 *
speed	3.932	0.416	9.46	1.5e-12 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 15.4 on 48 degrees of freedom

Multiple R-squared: 0.651, Adjusted R-squared: 0.644

F-statistic: 89.6 on 1 and 48 DF, p-value: 1.49e-12

The estimates of the regression coefficients β and their covariance matrix can be extracted from the fitted model via:

```
R> betahat <- coef(lm.cars)
R> Vbetahat <- vcov(lm.cars)
```

At first, we are interested in the hypothesis $\beta_1 = 0$ and $\beta_2 = 0$. This is equivalent to the linear hypothesis $\mathbf{K}\beta = 0$ where $\mathbf{K} = \text{diag}(2)$, i.e.,

```
R> K <- diag(2)
R> Sigma <- diag(1/sqrt(diag(K %%% Vbetahat %%% t(K))))
R> z <- Sigma %%% K %%% betahat
R> Cor <- Sigma %%% (K %%% Vbetahat %%% t(K)) %%% t(Sigma)
```

Note that $\mathbf{z} = (-2.6011, 9.464)$ is equal to the t statistics. The multiplicity-adjusted p values can now be computed by means of the multivariate t distribution utilizing the `pmvt` function available in package **mvtnorm**:

```
R> library("mvtnorm")
R> df.cars <- nrow(cars) - length(betahat)
R> sapply(abs(z), function(x) 1 - pmvt(-rep(x, 2), rep(x,
+      2), corr = Cor, df = df.cars))
```

```
[1] 1.661e-02 2.458e-12
```

Note that the p value of the global test is the minimum p value of the partial tests.

The computations above can be performed much more conveniently using the functionality implemented in package **multcomp**. The function `glht` just takes a fitted model and a matrix defining the linear functions, and thus hypotheses, to be tested:

```
R> library("multcomp")
R> cars.ht <- glht(lm.cars, linfct = K)
R> summary(cars.ht)
```

Simultaneous Tests for General Linear Hypotheses

```
Fit: lm(formula = dist ~ speed, data = cars)
```

Linear Hypotheses:

	Estimate	Std. Error	t value	p value
(Intercept) == 0	-17.579	6.758	-2.60	0.017 *
speed == 0	3.932	0.416	9.46	<1e-10 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
(Adjusted p values reported -- single-step method)

Simultaneous confidence intervals corresponding to this multiple testing procedure are available via

```
R> confint(cars.ht)
```

Simultaneous Confidence Intervals

```
Fit: lm(formula = dist ~ speed, data = cars)
```

```
Estimated Quantile = 2.13
95% family-wise confidence level
```

Linear Hypotheses:

	Estimate	lwr	upr
(Intercept) == 0	-17.579	-31.977	-3.181
speed == 0	3.932	3.047	4.818

The application of the framework isn't limited to linear models, nonlinear least-squares estimates can be tested as well. Consider constructing simultaneous confidence intervals for the model parameters (example from the manual page of `nls`):

```
R> DNase1 <- subset(DNase, Run == 1)
R> fm1DNase1 <- nls(density ~ SSlogis(log(conc), Asym,
+   xmid, scal), DNase1)
R> K <- diag(3)
R> rownames(K) <- names(coef(fm1DNase1))
R> confint(glht(fm1DNase1, linfct = K))
```

Simultaneous Confidence Intervals

```
Fit: nls(formula = density ~ SSlogis(log(conc), Asym, xmid, scal),
  data = DNase1, algorithm = "default", control = list(maxiter = 50,
    tol = 1e-05, minFactor = 0.0009765625, printEval = FALSE,
    warnOnly = FALSE), trace = FALSE)
```

```
Estimated Quantile = 2.137
95% family-wise confidence level
```

Linear Hypotheses:

	Estimate	lwr	upr
Asym == 0	2.345	2.178	2.512
xmid == 0	1.483	1.309	1.657
scal == 0	1.041	0.972	1.110

which is not totally different from univariate confidence intervals

```
R> confint(fm1DNase1)
```

	2.5%	97.5%
Asym	2.1935	2.539
xmid	1.3215	1.679
scal	0.9743	1.115

because the parameter estimates are highly correlated




Figure 1: `cars` data: Regression line with confidence bands (dashed) and intervals (dotted).

```
R> cov2cor(vcov(fm1DNase1))
```

```
      Asym  xmid  scal
Asym 1.0000 0.9868 0.9008
xmid 0.9868 1.0000 0.9063
scal 0.9008 0.9063 1.0000
```

Example: Confidence Bands for Regression Line. Suppose we want to plot the linear model fit to the `cars` data including an assessment of the variability of the model fit. This can be based on simultaneous confidence intervals for the regression line $x_i^\top \hat{\beta}$:

```
R> K <- model.matrix(lm.cars)[!duplicated(cars$speed),
+ ]
R> ci.cars <- confint(glht(lm.cars, linfct = K), abseps = 0.1)
```

Figure 1 depicts the regression fit together with the confidence band for the regression line and the pointwise confidence intervals as computed by `predict(lm.cars)`.

2 Multiple Comparison Procedures

Multiple comparisons of means, i.e., regression coefficients for groups in AN(C)OVA models, are a special case of the general framework sketched in the previous section. The main difficulty is that the comparisons one is usually interested in, for example all-pairwise differences, can't be directly specified based on model

parameters of an AN(C)OVA regression model. We start with a simple one-way ANOVA example and generalize to ANCOVA models in the following.

Consider a one-way ANOVA model, i.e., the only covariate x is a factor at j levels. In the absence of an intercept term only, the elements of the parameter vector $\beta \in \mathbb{R}^j$ correspond to the mean of the response in each of the j groups:

```
R> ex <- data.frame(y = rnorm(12), x = gl(3, 4, labels = LETTERS[1:3]))
R> aov.ex <- aov(y ~ x - 1, data = ex)
R> coef(aov.ex)
```

```
      xA      xB      xC
-0.2784  0.1496  0.5301
```

Thus, the hypotheses $\beta_2 - \beta_1 = 0$ and $\beta_3 - \beta_1 = 0$ can be written in form of a linear function $\mathbf{K}\beta$ with

```
R> K <- rbind(c(-1, 1, 0), c(-1, 0, 1))
R> rownames(K) <- c("B - A", "C - A")
R> colnames(K) <- names(coef(aov.ex))
R> K
```

```
      xA xB xC
B - A -1  1  0
C - A -1  0  1
```

Using the general linear hypothesis function `glht`, this so-called ‘many-to-one comparison procedure’ [Dunnnett, 1955] can be performed via

```
R> summary(glht(aov.ex, linfct = K))
```

Simultaneous Tests for General Linear Hypotheses

```
Fit: aov(formula = y ~ x - 1, data = ex)
```

Linear Hypotheses:

	Estimate	Std. Error	t value	p value
B - A == 0	0.428	0.653	0.66	0.74
C - A == 0	0.808	0.653	1.24	0.39

(Adjusted p values reported -- single-step method)

Alternatively, a symbolic description of the general linear hypothesis of interest can be supplied to `glht`:

```
R> summary(glht(aov.ex, linfct = c("xB - xA = 0", "xC - xA = 0")))
```

Simultaneous Tests for General Linear Hypotheses

```
Fit: aov(formula = y ~ x - 1, data = ex)
```

Linear Hypotheses:

	Estimate	Std. Error	t value	p value
xB - xA == 0	0.428	0.653	0.66	0.74
xC - xA == 0	0.808	0.653	1.24	0.39

(Adjusted p values reported -- single-step method)

However, in the presence of an intercept term, the full parameter vector $\theta = c(\mu, \beta_1, \dots, \beta_j)$ can't be estimated due to singularities in the corresponding design matrix. Therefore, a vector of *contrasts* γ^* of the original parameter vector γ is fitted. More technically, a contrast matrix \mathbf{C} is included into this model such that $\beta = \mathbf{C}\gamma^*$ any we only obtain estimates for γ^* , but not for γ :

```
R> aov.ex2 <- aov(y ~ x, data = ex)
R> coef(aov.ex2)
```

(Intercept)	xB	xC
-0.2784	0.4279	0.8085

The default contrasts in R are so-called treatment contrasts, nothing but differences in means for one baseline group (compare the Dunnnett contrasts and the estimated regression coefficients):

```
R> contr.treatment(table(ex$x))
```

```
4 4
4 0 0
4 1 0
4 0 1
```

```
R> K %*% contr.treatment(table(ex$x)) %*% coef(aov.ex2)[-1]
```

```
 [,1]
B - A 0.4279
C - A 0.8085
```

so that $\mathbf{K}\mathbf{C}\hat{\beta}^* = \mathbf{K}\hat{\beta}$.

When the `mcp` function is used to specify linear hypotheses, the `glht` function takes care of contrasts. Within `mcp`, the matrix of linear hypotheses \mathbf{K} can be written in terms of γ , not γ^* . Note that the matrix of linear hypotheses only applies to those elements of $\hat{\gamma}^*$ attached to factor x but not to the intercept term:

```
R> summary(glht(aov.ex2, linfct = mcp(x = K)))
```

Simultaneous Tests for General Linear Hypotheses

Multiple Comparisons of Means: User-defined Contrasts

```
Fit: aov(formula = y ~ x, data = ex)
```

Linear Hypotheses:

	Estimate	Std. Error	t value	p value
B - A == 0	0.428	0.653	0.66	0.74
C - A == 0	0.808	0.653	1.24	0.39

(Adjusted p values reported -- single-step method)

or, a little bit more convenient in this simple case:

```
R> summary(glht(aov.ex2, linfct = mcp(x = c("B - A = 0",
+      "C - A = 0"))))
```

Simultaneous Tests for General Linear Hypotheses

Multiple Comparisons of Means: User-defined Contrasts

```
Fit: aov(formula = y ~ x, data = ex)
```

Linear Hypotheses:

	Estimate	Std. Error	t value	p value
B - A == 0	0.428	0.653	0.66	0.74
C - A == 0	0.808	0.653	1.24	0.39

(Adjusted p values reported -- single-step method)

More generally, inference on linear functions of parameters which can be interpreted as ‘means’ are known as *multiple comparison procedures* (MCP). For some of the more prominent special cases, the corresponding linear functions can be computed by convenience functions part of **multcomp**. For example, Tukey all-pair comparisons for the factor *x* can be set up using

```
R> glht(aov.ex2, linfct = mcp(x = "Tukey"))
```

General Linear Hypotheses

Multiple Comparisons of Means: Tukey Contrasts

Linear Hypotheses:

	Estimate
B - A == 0	0.428
C - A == 0	0.808
C - B == 0	0.381

The initial parameterization of the model is automatically taken into account:

```
R> glht(aov.ex, linfct = mcp(x = "Tukey"))
```

General Linear Hypotheses

Multiple Comparisons of Means: Tukey Contrasts

Linear Hypotheses:

	Estimate
B - A == 0	0.428
C - A == 0	0.808
C - B == 0	0.381

3 Test Procedures

Several global and multiple test procedures are available from the `summary` method of `glht` objects and can be specified via its `test` argument:

- `test = univariate()` univariate p values based on either the t or normal distribution are reported. Controls the type I error for each partial hypothesis only.
- `test = Ftest()` global F test for H_0 .
- `test = Chisqtest()` global χ^2 test for H_0 .
- `test = adjusted()` multiple test procedures as specified by the `type` argument to `adjusted`: "single-step" denotes adjusted p values as computed from the joint normal or t distribution of the \mathbf{z} statistics (default), "free" implements multiple testing procedures under free combinations, "Shaffer" implements Bonferroni-adjustments taking logical constraints into account Shaffer [1986] and "Westfall" takes both logical constraints and correlations among the \mathbf{z} statistics into account Westfall [1997]. In addition, all adjustment methods implemented in `p.adjust` can be specified as well.

4 Quality Assurance

The analyses shown in Westfall et al. [1999] can be reproduced using `multcomp` by running the R transcript file in `inst/MCMT`.

References

Charles W. Dunnett. A multiple comparison procedure for comparing several treatments with a control. *Journal of the American Statistical Association*, 50(272):1096–1121, 1955.

- Juliet P. Shaffer. Modified sequentially rejective multiple test procedures. *Journal of the American Statistical Association*, 81:826–831, 1986.
- Peter H. Westfall. Multiple testing of general contrasts using logical constraints and correlations. *Journal of the American Statistical Association*, 92:299–306, 1997.
- Peter H. Westfall, Randall D. Tobias, Dror Rom, Russell D. Wolfinger, and Yosef Hochberg. *Multiple Comparisons and Multiple Tests Using the SAS System*. SAS Institute, Cary, NC, 1999.